

Outline of a Multi-agent System for Regulated Information Exchange in Crime Investigations

Pieter Dijkstra

Centre for Law & ICT,
Faculty of Law
University of Groningen
The Netherlands

p.dijkstra@rug.nl

Floris Bex

Centre for Law & ICT,
Faculty of Law
University of Groningen
The Netherlands

f.j.bex@rug.nl

Henry Prakken

Centre for Law & ICT,
Faculty of Law
University of Groningen
and

Institute of Information and
Computing Sciences,
Utrecht University
The Netherlands

henry@cs.uu.nl

Kees de Vey
Mestdagh

Centre for Law & ICT,
Faculty of Law
University of Groningen
The Netherlands

sesam@rechten.rug.nl

ABSTRACT

This paper outlines a multi-agent architecture for regulated information exchange of crime investigation data between police forces. An architecture is proposed consisting of two agents, a requesting agent and a responding agent, and of a communication language and protocol with which these agents can interact to promote optimal information exchange while respecting the law.

1. INTRODUCTION

In many organisations and societies information has to be exchanged, an activity which is often regulated by law. For instance, in the European Union exchange of any personal data is regulated by privacy law and in the Netherlands exchange of crime investigation data between police departments is regulated by a special act. Typically, organisations must balance the goal to exchange as much information as possible with the aim to stay within the law. It seems natural to investigate to what extent this balancing act can be supported with advanced information technology. Application of regulations can be supported with legal knowledge-based systems but the distributed nature of many organisations suggests that it may be worthwhile to combine knowledge-based technology with multi-agent technology.

This paper provides an initial exploration of this idea in the context of the exchange of crime investigation data between Dutch police forces. In particular, our aim is to outline an architecture of two agents, a requesting agent and a responding agent, and to sketch a communication language and protocol with which these agents can interact to promote optimal information exchange while respecting the law. Being an initial exploration, the analysis in this paper will be mostly conceptual and semiformal, and will leave many details for future research.

This research is part of an ongoing research project ANITA (Administrative Normative Information Transaction Agents), which aims at performing the fundamental research needed to

develop a multi agent system for regulated information exchange in the police intelligence domain [11].

This paper is organised as follows. First in Section 2 we discuss the problem of regulated information exchange and how it manifests itself in crime investigation exchange between Dutch police forces. In Section 3 we present two example interactions between information exchanging police offers, after which in Section 4 we list the main requirements for a multi-agent architecture in this domain. In Section 5 we then outline an architecture that meets these requirements, which we then apply to the examples in Section 6.

2. THE PROBLEM OF REGULATED INFORMATION EXCHANGE

Information exchange is often regulated by legal norms and by the policies of the exchanging institutions. This regulation of information exchange serves several goals. On the one hand, the privacy of the persons who are the subjects of the information must be protected. On the other hand, the legitimate interests of the exchanging institutions must be served. These interests of institutions vary from obtaining as much information as possible from other institutions to further their own objectives, to not providing information to other institutions in order to protect their own objectives.

Several types of conflicts of interest arise from these diverging goals. In most cases there is a central institution (for example, the state, the mother company) that is mainly interested in both optimal and legitimate information exchange, because it has to give account of the effectiveness and lawfulness of its operations to the outside world (for example, the parliament, the shareholders). Besides the central institution there often also are distributed local institutions with their own interests and objectives. The central institutions take the interests at the local level into account by formulating legal norms and central policies;

these give room for fine tuning in local policies and individual decisions by granting discretionary authority to local institutions. In many cases such discretionary authority is again constrained by general obligations (do's and don'ts). Some typical norms resulting from the need to reconcile conflicting interests are: 'Information must be exchanged if this is necessary for the execution of the other's appointed task' and 'It is allowed to refuse to exchange information if such refusal is necessary for the execution of one's own appointed task' (e.g. sections 14 and 13a of the Dutch Police Registers Act).

It can easily be seen that such norms may give rise to interesting negotiation and persuasion dialogues between officials of different local institutions. Ideally, these dialogues guarantee that an optimal and legitimate balance is found in the exchange of information in institutions characterized by differing and in some cases conflicting interests. However, in practice this ideal is not always realised. For example, it is well known that police departments are very reluctant to share crime investigation information with other departments, even if the sharing of information is allowed. One of the ultimate research goals of the ANITA project is to investigate whether such problems can be tackled by providing automated support for information-exchanging police officers.

3. EXAMPLES

As a possible solution to the above-mentioned problems we investigate the use of a multi-agent architecture. The idea is that the overall goals of an organisation (optimal and lawful information exchange) are promoted by the designs of the individual agents and the ways they interact. We illustrate this with a case study in crime investigation in the Dutch police organisation. The Dutch police organisation is divided into separate departments, which each operate in their own region. In order to solve crime cases, departments often need information held by other departments. Information exchange between police departments is governed by national and international privacy regulations and these regulations are supplemented by local rules of the departments. In consequence, when exchanging information with each other, police officers often have to interact in several ways to make sure they conform to the regulations and local interests.

We illustrate such interactions between police officers with two examples from police practice. Although these examples are imaginary, we have been assured by police officers that interactions like these occur in practice. A very typical interaction is about the exchange of information acquired from informants (about 80% of police information on heavy crime in the departments we examined is obtained from informants). Police departments are very cautious about the exchange of this kind of information, since crime suspects who are confronted with information obtained from informants may find out who supplied the information, and this may endanger the safety of the informant and the continuity of the investigation performed by the department that supplied the information. Therefore, in most cases the department that 'runs' the informant will not be willing to supply the information unless the receiving department offers certain guarantees.

Example 1: agent A working in police region A requests information about trading of explosive materials from agent B working in region B.

A: Give me all information about trading of explosive materials.

B: I will not give you this information.

A: Why don't you give me the information?

A: Because I am not allowed to do so.

A: Why are you not allowed to give me the information?

B: Because it is protected.

A: Why is the information protected?

B: Because an investigation may be at risk.

A: You may be right in general but in this case the information is not protected because this is a matter of national importance.

B: Ok, I admit that in this case the information is not protected, so I retract that I am not allowed to give you the information. I will give you the information on the following condition: the given information may not be exchanged with other police officers.

A: I agree with this condition.

Example 2: agent A working in police region A requests information about a suspect from agent B working in region B.

A: Tell me all you know about suspect X.

B: No I won't.

A: Why don't you want to tell me about suspect X?

B: Because I need to protect my informant.

A: Why do you need to protect your informant?

B: I need to protect my informant because your use of my information about suspect X could disclose his identity and also endanger the continuity of several of my investigations.

A: You don't need to protect your informant since I will only use the information about suspect X statistically for policy reasons and therefore not in individual investigations.

B: Ok, in that case I don't need to protect my informant. I will give the information under the following condition: the given information may not be exchanged with other police officers.

A: I agree with this condition.

4. REQUIREMENTS FOR THE MULTI-AGENT ARCHITECTURE

In this section we sketch the requirements for a multi-agent architecture for regulated information exchange that can be applied to the above-sketches police scenario.

Knowledge

As seen in the previous sections the agents must have knowledge of the relevant regulations for information exchange and the information stored in the local database to solve crime cases constrained by those regulations. Also knowledge of regional interpretations of national regulations must be known by the regional agent. This knowledge is only accessible locally and will

not be shared between the agents. Furthermore, the agents must have knowledge of the likely consequences of their communicating acts for the realisation of their goals. Finally, in order to represent knowledge and being able to share information the agents must have an ontology of the domain. Ideally all agents use the same ontology so problems between different ontologies can be avoided.

Reasoning

As is usual in legal domains, most of the available knowledge is defeasible. Also, as could be seen from the examples in the previous section, the interaction between agents often involves argumentation. Therefore, the agents should be capable of generating and evaluating arguments for and against certain claims. Also, to generate conditional offers, the agents must be able to do some form of hypothetical reasoning.

Goals

As described above, the agents in our problem domain have individual goals. In our case there are two agent roles: a requesting agent (denoted by A) who wants to collect as much information as possible for the purpose of his crime investigations, and a responding agent (denoted by B) who wants to protect his own information resources and crime investigations. In addition both agents want to contribute to the overall goals of the police organisation, which are the optimal and lawful exchange of information. The agents and their interactions should be designed in such a way that their behaviour agrees with their goals.

Communication

Of course, the agents should be able to exchange information but other types of interactions should also be possible. Above we noted that the receiving agent's goals sometimes lead him to state conditions under which he is willing to give information. Therefore, the agents must be able to negotiate with each other. Also, the receiving agent may be mistaken in believing that he must or should not give the requested information. Therefore, the agents must be able to engage in persuasion dialogues to reach a better information state. To enable such interactions, a suitable dialogue protocol must be implemented. Also, the agents must be given policies, or tactics, for their behaviour in the dialogues. These policies should be designed to further the agent's goals. Since these goals include those of the overall institution, the agents' policies should induce a fair degree of cooperativeness.

In fact, there are (at least) two ways to model the relation between the three types of dialogues, depending on two ways to interpret the start of a dialogue in this domain.

i. Each dialogue starts as an information-seeking dialogue. It shifts to another type of dialogue if the responding agent B states he will not grant the request since doing so would have negative consequences for his investigations. It then shifts to a persuasion dialogue if the requesting A starts to persuade B that he is wrong about this, while the dialogue shifts to a negotiation dialogue if A promises to do or refrain from doing something on the condition that B gives him the information. Further shifts may occur, for instance, from a persuasion to a negotiation dialogue or vice versa.

ii. Each dialogue starts as a negotiation, viz. as a request to give information about something. Such a dialogue may shift to persuasion if B rejects the request on the grounds that granting it would have negative consequences for his investigations and A tries to persuade B that he is wrong about this. Each terminated persuasion pops up to the interrupted negotiation. If that terminates successfully, a (trivial) information-seeking dialogue starts; its termination also terminates the overall interaction.

We contend that most interactions in our domain will be of type (ii) since usually the requesting agent will not simply ask a question but will inquire whether the other agent is willing to provide him with a certain body of information. This seems more like negotiation than like information-seeking. In the remainder of this paper we will therefore only focus on interactions of the second type.

Figure 1 gives a high-level view of the required multi-agent architecture. Both agents have access to their local information database. If the requesting agent A wants to access information from region B it has to communicate with the responding agent B. Using the communication channel the agents can have dialogues about conditions and when they agree the final step is the exchange of information.

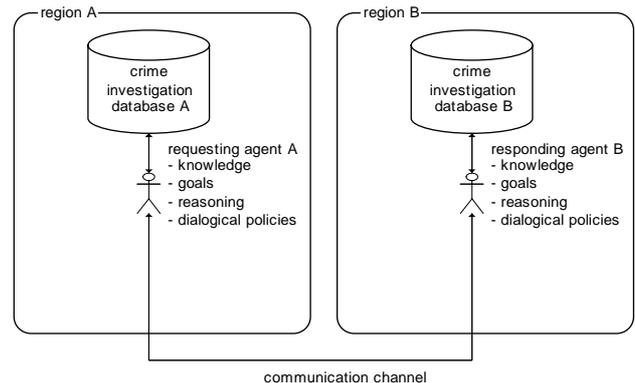


Figure 1: interaction between information-exchanging agents

5. SKETCH OF A COMPUTATIONAL ARCHITECTURE

In this section we sketch a computational architecture that respects the requirements of the previous section. We first sketch a dialogue system for negotiation with embedded persuasion, consisting of a communication language and a protocol. Then we outline the main components of the individual agents: knowledge representation, reasoning, goals and dialogue policies.

5.1. Dialogical interaction

Dialogue systems have a *communication language* C with associated *protocol* P and a *topic language* T with associated *logic* L (possibly nonmonotonic). The communication language consists of *speech acts* l(t) where l is a *locution* and t an element or subset of T or an argument in L. Our language and protocol essentially is that of [9] and [10], which follows the general format of [6] as extended and revised in [7]. The system of [10]

combines a negotiation language and protocol of [12] with a persuasion language and protocol of [7]. The system of [10] seems appropriate for present purposes since, as noted above in Section 3, the agent interactions typically have the form of negotiation with embedded persuasion. Also, this embedding occurs especially when the requesting agent asks why the responding agent rejects the request for information, and this is precisely the kind of embedding modelled in [10]. Finally, the negotiation language of [12] provides the required means to attach conditions to acceptances of request, in a manner explained below.

5.1.1 Communication language

We first present the sublanguages for negotiation and persuasion and then define their combination.

In [10], following [7], a communication language has a reply structure: each speech act replies to one preceding speech act in the dialogue. Moreover, a reply can be of two kinds, being either an attacking or a surrendering reply. How to attack or surrender to another speech act is specified in the following tables.

speech acts	attacks	surrenders
claim p	why p	concede p
why p	p since Q	retract p
p since Q	why q ($q \in Q$) p' since Q'	concede q ($q \in Q$) concede p
concede p		
retract p		

Table 1: a persuasion communication language (Cp)

In table 1 “since” speech acts are arguments in L (which is a logic for defeasible argumentation) and p' since Q' attacks p since Q according to L.

speech acts	attacks	surrenders
offer p	offer q ($q \neq p$) reject p	accept p withdraw
reject p	offer q ($q \neq p$)	
accept p		
withdraw		

Table 2: A negotiation communication language (Cn)

In table 2 the propositional contents of the locutions typically are act descriptions, such as “you tell me all you know about suspect X” or “you do not pass on the information to other police officers”. The idea of table 2 is that an initial request is made as a special kind of offer, namely, an offer that the other party do something. This is for simplicity only: since the replies to a request are the same as to an offer, having separate speech acts for both would be an unnecessary complication. Yet for convenience an initial offer will in examples below be denoted by ‘request’.

The negotiation language is very simple and the distinction between attacking and surrendering replies would seem to make no sense, but later the tables will be extended and combined in a way that makes this distinction sensible for negotiation also.

The propositional variables in table 2 are assumed to conform to the syntax of WP, where propositions are conjunctions of issue-value expressions in a subset of first-order logic. For present purposes what matters most is that conjunctions can be used to state conditions for acceptance. For example, if the responding agent says ‘request tell-me-about-x = yes’ then the responding agent could state conditions in a counteroffer ‘offer (tell-me-about-x = yes \wedge treatment-after-use = destroy)’ thus requiring that the requested information is destroyed after use. We contend that representing conditional offers with a conjunction is more natural than representing them with conditional operators, since the content of a conjunctive agreement allows the agents to infer what they have committed to irrespective of whether the other agent keeps his part of the agreement.

In combining the two combination languages, the idea of [10] is to add to the negotiation language a speech act that triggers a persuasion dialogue: a new attacking reply is added in Cn to a ‘reject p’ move, namely ‘why-reject p’. The only possible reply to this move is claim q, where q is a ground for the rejection. This claim starts a persuasion dialogue.

speech acts	attacks	surrenders
offer p	offer q ($q \neq p$) reject p	accept p withdraw
reject p	offer q ($q \neq p$) why-reject p	
accept p		
withdraw		
why-reject p	claim q	
claim p	why p	concede p
why p	p since Q	retract p
p since Q	why q ($q \in Q$) p' since Q'	concede q ($q \in Q$) concede p
concede p		
retract p		

Table 3: A combined communication language (Cn)

5.1.2. Communication protocol

We first sketch the individual negotiation and persuasion protocols. They have the following rules in common. A *move* is a speech act made by a dialogue participant. If it is not the first move, it replies to a unique preceding move in the dialogue made by the other party, according to the reply structure of C. A dialogue *terminates* if a player is to move but has no legal moves.

The *negotiation protocol* then is very simple. Agent A begins with an offer, and then the agents take turns after each move, replying to the last move of the other party. Thus negotiations can be of arbitrary length and terminate after an accept or withdraw move but they are not guaranteed to terminate.

The *persuasion protocol* is more involved. We sketch the main rules only. Each persuasion starts with a claim made by agent A. The main protocol rule is that each move is relevant. Relevance is defined in terms of the notion of dialogical status of a move. Briefly, a move is *in* if it has a surrendering reply or else all its attackers are *out*; and a move is *out* if it is not surrendered and has an attacking reply that is *in*. As for turntaking, an agent keeps moving until the dialogical status of the initial move has changed, then the turn switches to the other agent. Thus each turn of a player consists of zero or more surrenders followed by zero or one attacker. Also, these rules imply that unlike in negotiation dialogues, in persuasion dialogues postponing replies and making alternative replies to earlier moves are allowed. Thus, if the reply structure of a dialogue is made explicit in a graph, the graph of a negotiation dialogue is a linear structure while that of a persuasion dialogue can be any tree. Finally, all this implies that a persuasion dialogue terminates if a player is to move but has no relevant moves: in the present simple protocol this can happen only if either the player to move is A and he has retracted his initial claim or the player to move is B and he has conceded the initial claim.

As for the *combined protocol*, the main idea is that if a negotiation dialogue shifts to a persuasion dialogue, their relation is one of embedding (cf. [3]): the embedded persuasion dialogue is undertaken until its termination, after which the embedding negotiation dialogue is resumed. So whenever a persuasion move is allowed by the protocol, no negotiation move is allowed. In addition, the structural rules of the persuasion system now also hold for negotiation, especially those of relevance and turntaking. This allows for alternative explanations for rejections and also for accepting an offer (perhaps conditionally) that was first rejected for reasons that could not be upheld in a persuasion dialogue. For a more detailed specification of the combined protocol the reader is referred to [10].

5.2 The agents

We next outline the architecture of the individual agents.

5.2.1 Knowledge Representation

The knowledge representation language must allow for the expression of a suitable crime investigation ontology, the deontic modalities ‘obligatory’, ‘permitted’ and ‘forbidden’ and for the description of actions and their effects. Since knowledge representation is not the focus of this paper, we will not go into further detail

5.2.2. Goals

The agent’s goals will not be represented explicitly inside an agent. Instead, they will be implicitly captured by the agent’s knowledge and by the choices an agent will make in a dialogue (for these choices see the dialogue policies specified below). The agents’ knowledge bases will contain knowledge that is relevant for respecting their goals, such as knowledge on when exchanging information is allowed, obligatory or forbidden and when

exchanging information is likely to endanger informants of crime investigations.

5.2.3. Reasoning engine

We assume that agents are capable of performing argument-based defeasible reasoning with their knowledge. In particular, we assume that they are able to construct arguments for certain propositions and then to verify whether these arguments are justified, defensible or overruled (cf. [8]). We also assume that the agents are able to perform hypothetical reasoning with some form of top-down inference: in situations where standard applications of knowledge-based systems would ask the user for information, our agents can make the required information into a condition of an offer when the condition is an action that can be performed by the other agent.

5.2.4 Dialogue policies

We next specify policies for what an agent will choose to do at various points in a dialogue. These policies can be regarded as partial strategies in a game-theoretic sense. To our knowledge, formal dialogue policies for argumentation dialogues were first studied by [5], who called them “agent attitudes”. The main difference with [5] is that when a policy requires an agent to reason with the available information, in our case the agent only reasons with his initial knowledge base plus the propositions that he has explicitly conceded in the persuasion dialogue, while in [5] the agent must also reason with everything the other agent has said, regardless of whether he has conceded to this or not. We regard the latter as less desirable.

There are two main kinds of dialogue policies: negotiation policies and persuasion policies. Negotiation policies describe, for example, whether the agent is obliged to give the requested information, whether it is sensible to accept a certain offer. Of course these policies are different for the requesting and the responding agent. Also, different responding agents can have different policies. One agent, for example, might easily be persuaded to give information while the other guards its secrets more closely.

Persuasion policies determine, for example, what kinds of arguments (e.g. defensible, justified) an agent should accept and when it is allowed to give a counterargument to an argument moved by the other agent.

In this paper we will constrain ourselves to giving a possible negotiation and a possible persuasion policy for a responding agent. First, we will give a negotiation policy for how to react to a request for information made by the requesting agent. Secondly, we will give a general persuasion policy for how to react to an argument. Policies for other kinds of agents and actions can be developed similarly.

When a responding agent receives a request for information, it must first determine whether it has the obligation and/or the permission to give the information, viz.:

Obligation

Can the responding agent construct a *justified* argument which says that the agent is *obliged* to give the information?

- If there is such an argument, the information should be given to the requesting agent.

- If there is no such argument, the responding agent should find out whether he is permitted to give the information or not.

Note that the agent could also reject the request if there is no justified argument which says he is obliged to give the information. However, this action does not stay true to the overall goal of the system, which says that if at all possible, information should be exchanged.

Permission

Can the responding agent construct a *justified* argument which says that the agent is *not permitted* to give the information?

- If there is such an argument, the request should be rejected.
- If there is no such argument, it can be assumed that the responding agent is permitted to give the information and the agent should find out whether it is sensible to give the information or not.

Obligation and permission usually follow from the relevant regulations which are in force at the police department. The sensibility of giving the information is more dependent of the department's local culture. For example, even if the regulations do not explicitly forbid the agent to give the information, the agent may still decide that giving the information may endanger the investigation at his own department.

Sensibility

Can the responding agent construct a *justified* argument which says that it is *not sensible* to give the information?

- If there is such an argument, the responding agent will have to look if there are additional conditions under which the information *can* be given.
 - If there are additional conditions make an offer to give the information provided the extra conditions are satisfied.
 - If there are no such additional conditions, reject the request.
- If there is no such argument, the responding agent should still look if there are additional conditions under which it is sensible to give the information.
 - If there are additional conditions offer the information provided these conditions are satisfied.
 - If there is no justified argument which says that it is not sensible to give the information and there are no conditions which make the information exchange better, then the responding agent has clearly exhausted all possible possibilities in protecting his information and so it should finally give the information to the requesting agent.

It should be noted that in the above policy the responding agent requires *justified* arguments which steer his decisions. A less careful agent might also accept defensible or even overruled arguments for the fact that he is obliged or permitted to give the information.

We will now turn to some examples of policies which can be used in a persuasion dialogue. The responding agent will have to react to arguments moved by the requesting agent or to *why* questions asked by the requesting agent. We will first explain how an agent can react to an argument.

Responding to arguments

Can the responding agent construct a counterargument which is neither overruled by the opponents argument nor by one of his own, internal arguments?

- If the agent can construct such an argument, it should be moved in the dialogue.
- If the agent cannot construct such an argument and there is a premise *p* of the opponent's argument for which the agent has no justified argument, then the agent should ask a *why p* question.
- If the agent cannot construct such an argument and for all of the premises of the opponent's argument has a justified argument then *concede* to the conclusion of the opponents argument.

Responding to *why* moves

Say that the requesting agent asks a *why p* question in response to a claim or argument move by the responding agent. Can the responding agent construct an argument for *p* which is neither overruled by the opponent's argument nor by one of his own, internal arguments?

- If the agent can construct such an argument, it should be moved in the dialogue.
- If the agent cannot construct such an argument, it should retract it's claim *p* or the conclusion of the argument of which *p* is a premise.

A few things should be noted regarding the above policy. Firstly, the agent uses not only *justified* but also *defensible* arguments to react to his opponent. So in persuasion, the agent is more confident than in negotiation, where he needed *justified* arguments to make decisions. Secondly, the agent is cooperative in that he only moves arguments in the dialogue which are not overruled by his own arguments. So even though the opponent may not have a stronger counterargument to a certain argument, the agent does not move the argument if it is overruled by one of his own arguments. Thirdly, the agent is also cooperative in that it only asks *why p* questions if it does not have an argument for *p*.

Clearly, several other policies are possible. One option we want to explore in future research is to make policies partly domain-specific. For example, the second part of the policy for responding to arguments could be refined such that premises are never challenged when they are about subject X and/or when they are claimed by person Y.

6. ILLUSTRATION OF THE PROPOSED ARCHITECTURE

We next illustrate the outlined architecture with a more formal reconstruction of Example 1. We assume that the language is closed under negation and material implication but make no

further assumptions on the language or logic. Up to B6 each move replies to the immediately preceding move.

The responding agent B has the following commitments:

$\{ (\neg e \Rightarrow \neg i), (\neg i \Rightarrow g), (p \Rightarrow \neg a), (r \Rightarrow p), r, p, n \}$

where e = “the information may be exchanged with other police officers”, i = “the informant is in danger”, g = “you give all information about trading explosive materials”, p = “the information is protected”, a = “agent B is allowed to give information about explosive materials”, r = “the investigation of agent B’s department may be in danger”, n = “the information is a matter of national importance”

The dialogue from example 1 (section 3) will go as follows:

A1: request g (Give me all information about trading of explosive materials).

According to the negotiation policy outlined in section 5.2.4, agent B will first have to look whether it is obliged to give the information or not. There is no rule stating that he is obliged to give the information, so B applies the next step in the policy, which is finding out whether B is allowed to give the information or not. B can construct a justified argument “ $\neg a$ since p ” which states that B is not allowed to give the information, so B rejects the request.

B1: reject g (I will not give you this information)

A2: why-reject g (Why don’t you give me the information?)

B now starts a persuasion dialogue. Until it is terminated no negotiation moves are allowed by the protocol.

B2: claim $\neg a$ (I am not allowed to give you the information)

A3: why $\neg a$ (Why are you not allowed to give me the information?)

B can construct an argument for $\neg a$ so, according to the persuasion policy, it should move this argument.

B3: $\neg a$ since p (I am not allowed since the information is protected)

A4: why p (Why is the information is protected?)

B can construct an argument for p so it should move this argument

B4: p since r (The information is protected because an investigation may be at risk)

A has an argument for the conclusion $\neg p$, which we assume is at least as strong as B’s argument for p

A5: $\neg p$ since n (The information is not protected because this is a matter of national importance)

B cannot ask a “why n ” question here because n is part of B’s commitments, so B has to concede to $\neg p$

B5: concede $\neg p$ (Ok, I admit that in this case the information is not protected)

B’s original reason for $\neg a$ is defeated. B therefore concedes to p ’s negation.

B6: retract $\neg a$ (I retract that I am not allowed to give you the information)

With this move B backtracks to A3, now surrendering to that move. This move terminates the persuasion dialogue so that negotiation moves are allowed again by the protocol.

B now knows he is not obliged to give the information, but he is also allowed to give the information. B now has to look if there are extra conditions under which the information can be given (because it is more sensible). Through backward chaining through the rules $(\neg e \Rightarrow \neg i)$ and $(\neg i \Rightarrow g)$, B arrives at the possible extra condition $\neg e$.

B7: offer $g \wedge \neg e$ (I will give you the information on the following condition: the given information may not be exchanged with other police officers).

With B7 the responding agent backtracks to A1, this time replying with a counteroffer.

A6: accept $g \wedge \neg e$ (I agree with this condition)

7. CONCLUSION

In this paper we have outlined a multi-agent architecture for regulated information exchange and we have illustrated it with an application to information exchange between police forces in the context of crime investigation. The architecture combines and adapts several elements from the literature: a defeasible-argumentation mechanism for the agents’ internal reasoning behaviour, a communication language and protocol for negotiation with embedded persuasion about reasons for rejections of offers, and dialogue policies for one of the agents of our architecture. Our dialogue policies for persuasion arguably improve those of [5] in one respect and we have added dialogue policies for negotiation. Also, we have proposed a novel view on the nature of dialogues in the context of regulated information exchange, viz. as negotiation with embedded persuasion.

As for related research, Trevor Bench-Capon has in [1] proposed an information-seeking protocol where one of the preconditions for answering a question is that doing so is permitted. Thus the question of lawfulness of providing the information is modelled as an aspect of the dialogue protocol. Arguably a drawback of this approach is that usually interaction protocols are meant to promote coherence and rationality of a dialogue; they are not meant to promote lawfulness of dialogues under some arbitrary normative system. We therefore regard discussions about lawfulness of speech acts as a domain matter, to be the topic of separate persuasion dialogues.

Parsons, Sierra & Jennings [4] model negotiation as argumentation: a proposal is the conclusion of an argument, the premises of which are the grounds to make the proposal. We think that whether it is good to communicate the reasons for a proposal is context-dependent: for instance, if a buyer says “please sell me this since I need it badly”, the seller is likely to offer a higher price. Therefore, arguments for or against a proposal can better be exchanged in a separate persuasion dialogue.

Doutre, McBurney and Wooldridge [2] propose a model for regulated information exchange in the medical domain. They model it as information seeking with embedded persuasion about whether providing the requested information is permitted.

As for future research, the present informal and semiformal outline should, of course, be made fully formal and then

implemented. Also, we want to investigate other combination patterns of dialogue types, including also information-seeking. As for the negotiation part of the dialogue system, we aim to investigate whether besides arguing about rejections, other ways to argue in negotiation, such as those studied by [4], occur in our application domain and should therefore be modelled in our architecture.

8. ACKNOWLEDGMENTS

This research was supported by the Netherlands Organisation for Scientific Research (NWO) under project numbers 634.000.017 and 634.000.429. Henry Prakken was also partially supported by the EU under IST-FP6-002307 (ASPIC).

9. REFERENCES

- [1] Bench-Capon, T.J.M. Specifying the Interaction Between Information Sources, *Proceedings of DEXA 98*, Springer LNCS 1460, pp 425-434. Berlin, 1998, Springer Verlag.
- [2] Doutre, S., McBurney, P. and Wooldridge, M. Law-Governed Linda as a semantics for agent interaction protocols (research abstract). To appear in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, Utrecht, The Netherlands, July 2005.
- [3] McBurney, P. and Parson, S. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information* 11 (2002): 315-334.
- [4] Parsons, S., Sierra, C. and Jennings, N. Agents that Reason and Negotiate by Arguing. *Journal of Logic and Computation* 8 (1998): 261-292.
- [5] Parsons, S., Wooldridge, M. and Amgoud, L. An analysis of formal inter-agent dialogues *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, pp. 394-401. New York: ACM Press.
- [6] Prakken, H. On dialogue systems with speech acts, arguments, and counterarguments. *Proceedings of the 7th European Workshop on Logics for Artificial Intelligence (JELIA '2000)*, Springer LNAI 1919, pp. 224-238. Berlin, 2000, Springer Verlag.
- [7] Prakken, H. Coherence and flexibility in two-person dialogue games for argumentation. Technical Report Institute of Information and Computing Sciences, Utrecht University (in preparation).
- [8] Prakken, H. and Vreeswijk, G. Logical systems for defeasible argumentation. In D. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic*, second edition, Vol 4, pp. 219-318. Kluwer Academic Publishers, Dordrecht etc., 2002.
- [9] Van Veenen, J. and Prakken, H. A verifiable protocol for arguing about rejections in negotiation (research abstract). To appear in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, Utrecht, The Netherlands, July 2005.
- [10] Van Veenen, J. and Prakken, H. A verifiable protocol for arguing about rejections in negotiation. Submitted.
- [11] De Vey Mestdagh, C.N.J., Administrative Normative Information Transaction Agents (ANITA): Legitimacy and Information Technology, the best of two worlds. In: *Access to knowledge and its enhancements, Proceedings ToKeN2000 symposium*, Delft University of Technology, February 21, 2003.
- [12] Wooldridge, M. and Parsons, S. Languages for negotiation. *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI-2000)*, pp. 393-400. Amsterdam: IOS Press.