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

Pieter Dijkstra[1], Floris Bex[1], Henry Prakken[1,2] and Kees de Vey Mestdagh[1]

[1]Centre for Law & ICT, Faculty of Law
University of Groningen
The Netherlands
{p.dijkstra, f.j.bex, c.n.j.de.vey.mestdagh}@rug.nl
[2]Department of Information and Computing Sciences, Utrecht University
The Netherlands
henry@cs.uu.nl

Abstract. This paper outlines a multi-agent architecture for regulated information exchange of crime investigation data between police forces. Interactions between police officers about information exchange are analysed as negotiation dialogues with embedded persuasion dialogues. An architecture is then proposed consisting of two agents, a requesting agent and a responding agent, and a communication language and protocol with which these agents can interact to promote optimal information exchange while respecting the law. Finally, dialogue policies are defined for the individual agents, specifying their behaviour within a negotiation. Essentially, when deciding to accept or reject an offer or to make a counteroffer, an agent first determines whether it is obligatory or permitted to perform the actions specified in the offer. If permitted but not obligatory, the agent next determines whether it is in his interest to accept the offer.

1 Introduction

In many organisations 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 obligation to stay within the law. It is important 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 explores this idea in the context of the exchange of crime investigation data between Dutch police departments. In particular, we propose an architecture of two agents, a requesting agent and a responding agent, and we define a communication language and protocol with which these agents can interact to promote optimal information exchange while respecting the law. We model the overall structure of their interactions as a negotiation dialogue about whether a certain body of information will be exchanged which may shift to a persuasion dialogue about whether such exchange is allowed or sensible. Finally, we define dialogue policies for the individual agents, specifying their behaviour within a negotiation. Essentially, when deciding to accept or reject an offer or to make a counteroffer, an agent first reasons about the law and then about the interests that are at stake: he first determines whether it is obligatory or permitted to perform the actions specified in the offer; if permitted but not obligatory, the agent next determines whether it is in his interests to accept the offer. The policies are specified within the overall assumption that the agents take the overall goal of optimal and lawful information exchange into account.

The architecture proposed in this paper assumes and combines several formalisms and ideas from the literature. The main original contributions are an account of regulated information exchange as negotiation with embedded persuasion and a specification of dialogue policies for negotiation within this application domain. In order to focus on these original contributions, the elements taken from the literature will be summarised or referred to only; for full descriptions of the technical details the reader will be referred to the literature.

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 propose 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 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, that each operate in its 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 at the same time serve local interests.

We illustrate such interactions between police officers with two examples from police practice. We have been assured by police officers that interactions like these are exemplary. 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 in explosive materials from agent *b* working in region *b*.
A: Give me all information about trading in explosive materials.
B: I will not give you this information.
A: Why don't you give me the information?
B: Because I am not allowed to do so.
A: Why are you not allowed to share the information?
B: Because it is protected.
A: You may be right in general but in this case it is allowed to share the information because this is a matter of national importance.
B: Ok, I admit that in this case it is allowed to share the information, 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 my informant's identity and though endangering 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 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-sketched 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. Knowledge of regional interpretations of national regulations must also be available to 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 with different ontologies can be avoided.

Reasoning
As usual in legal domains, most of the available knowledge is defeasible. Also, as can 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. Finally, 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 (information-seeking, negotiation and persuasion) dialogues, depending on two ways to interpret the start of a dialogue in this domain.

i. The first interpretation is that each dialogue starts as an information-seeking dialogue. It then shifts to another type of dialogue, either persuasion or negotiation. It shifts to a persuasion dialogue if the responding agent *b* states he will not grant the request since doing so would have negative consequences for his investigations and the requesting *a* starts to persuade *b* that he is wrong about this. The dialogue instead shifts to negotiation 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. The second interpretation is that 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. After the persuasion terminates the interrupted negotiation resumes. 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* he 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 which is the exchange of information.

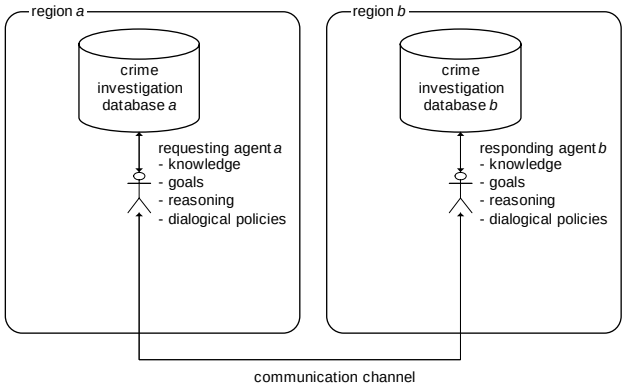

Fig. 1. Interaction between information-exchanging agents

5 Outline of a computational architecture

In this section we outline 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 [10], which follows the general format of [7] as extended and revised in [8]. The system of [10] combines a negotiation protocol of [12] with a persuasion language and protocol of [8]. 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].

5.1.1 Communication language

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

In [10], following [8], 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.

Table 1. A persuasion communication language (C_p)

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		

In table 1 "since" speech acts state arguments in L (which is a logic for defeasible argumentation) and p' since Q' attacks p since Q according to L, that is, the latter is a counterargument of the former.

Table 2. A negotiation communication language (C_n)

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		

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 (possibly conjoined with an action to-be performed by the offeror). 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 be denoted by 'request' in the example below.

The negotiation language is very simple and the distinction between attacking and surrendering replies would seem to make no sense, but the tables will be extended and combined in a way that makes this distinction sensible for negotiation also. In table 2 the propositional contents of the locutions typically are (possibly negated) 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". Since the logic of reasoning about actions is not the main concern of this paper, we simply assume that these act descriptions are expressed in first-order logic. More precisely, formulas from T_n are assumed to be conjunctions of first-order literals (atomic formulas or their negations) without free variables. For present purposes what matters most is that conjunctions can be used to state conditions for acceptance. For example, if the requesting agent says 'request: you tell me all you know about suspect X' then the responding agent could state conditions in a counteroffer 'offer: I tell you all I know about suspect X $\wedge$ you do not pass on the information to other police officers' thus requiring that the requested information is not passed on to other police officers. 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 communication 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 C_n 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.

Table 3. A combined communication language (C_{np})

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		

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* 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 elaborate. 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 is it 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. [4]): 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. Since knowledge representation and reasoning are not the primary focus of this paper, we will be brief and informal on these aspects, making only those assumptions about the architecture that are required by the specification of the dialogue policies.

5.2.1 Representation of knowledge and goals

The knowledge representation language must allow for the expression of a suitable crime investigation ontology, for the deontic modalities 'obligatory', 'permitted' and 'forbidden', for the description of actions and their effects and for the expression of an agent's goals. These goals are assumed to be represented in the agent's internal knowledge base, together with knowledge that is relevant to respecting the agent's goals, such as knowledge on when exchanging information is obligatory, allowed or forbidden and when exchanging information is promotes or violates the agent's interests.

5.2.2. Reasoning engine

We assume that agents are capable of performing argument-based defeasible reasoning with their internal 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. [9]). In the present model the agents especially reason about whether something is obligatory, permitted or forbidden, and about whether something is in their interest. With respect to the latter it is useful to remark that we assume that any conflict between arguments on this issue can be resolved in the defeasible reasoning process. For instance, if the agent has constructed a justified argument that a particular action does not violate his interests, then we assume that he has considered and refuted all counterarguments that, say, the action detracts from one of his goals, or that there is another way to satisfy the same goal. For a similar approach to practical reasoning based on argumentation schemes see Atkinson et al. [1].

We also assume that the agents are able to perform hypothetical reasoning. This is especially useful when the agent has to reason about whether to make a counteroffer to an offer: hypothetical reasoning provides the candidate conditions to be added to such a counteroffer.

5.2.3 Dialogue policies

We next specify policies for what an agent will choose to do at various points in a dialogue. In our application, two main kinds of dialogue policies need to be specified, viz. for negotiation and for persuasion. To our knowledge, formal dialogue policies for argumentation dialogues were first studied by [6], who called them "agent attitudes". They only defined policies for persuasion, while we are particularly interested in policies for negotiation and largely adopt [6]'s persuasion policies. An important difference with [6] is that when a policy requires an agent to construct arguments, in our case the agent only reasons with his initial knowledge base plus the propositions that he has explicitly conceded to in the persuasion dialogue, while in [6] 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 realistic.

Our negotiation policies consider two issues: the normative issue whether accepting an offer is obligatory, permitted or forbidden, and the teleological issue whether accepting an offer is in the agent's interest. Of course these policies can be 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 and thus agree with every request for information he receives, while another agent may guard his secrets more closely, adding extra conditions to the information exchange to make sure the information does not fall into the wrong hands.

For simplicity we have chosen to specify just one set of policies which we think best suits our domain (although we will comment on other design choices when relevant). As remarked above, in the police domain all agents must on the one hand protect their own interests, such as their own investigations and informants, but on the other hand they must cooperate to

ensure optimal exchange of information within the bounds of the law. Our set of policies is meant to implement these general considerations. Since the negotiation policies are the main novel contribution of this paper, we will especially focus on their specification.

The following notation will be used in the policies. Recall first that T_n is the negotiation topic language and T_p is the persuasion topic language. The content of an offer is a conjunction of literals from T_n. For any formula p of T_n and agent a the notation p^a denotes the conjunction of all conjuncts in p that describe actions performed by a. Also, $KB_d(a) \subseteq T_p$ denotes the set of beliefs of agent a at dialogue stage d. Note that these beliefs may change during persuasion if the agent concedes or retracts propositions and then brings his internal beliefs in agreement with his (public) commitments. Finally, we assume that each offer content p is divided into two parts $q \wedge c$, where q are the essentials of an offer, which the offeror does not want to be changed in counteroffers, and where c are the conditions of the offer, which the offeror does not mind to be changed in counteroffers. For instance, an offer "send me a file containing all you know about mr. X" can be decomposed into an essential element "inform me about all you know about mr. X" and a condition "send me the information in a file". The offeror thus indicates that he does not want to receive counteroffers "I will tell you how to find out everything about mr. X" but that he does not mind receiving a counteroffer "I will tell you everything I know about mr. X over the phone but I will not send you a file". We will now give the different policies, starting with the negotiation policies.

5.2.3.1 Negotiation policies

Responding to an offer

We now specify the policy agent a should apply in responding to an offer in a dialogue. Note that the offer can be the initial one (i.e., a request) or a counteroffer.

When agent a receives an offer $p = q \wedge c$, where q are the *essentials* of the offer and c the *conditions*, he should perform the following actions:

- Determine whether $KB_d(a)$ supports a (justified/defensible) argument for the conclusion that p^a is obligatory. If there is such an argument, then accept the offer. Otherwise,
- Determine whether $KB_d(a)$ supports a (justified/defensible) argument for the conclusion that p^a is forbidden. If there is such an argument, then reject the offer. Otherwise,
- Determine whether $KB_d(a)$ supports a (justified/defensible) argument for the conclusion that p^a violates a's interests. If there no such argument, then accept the offer. Otherwise:
  - Find a subset-minimal set $c' \subseteq T_n$ such that $KB_d(a) \cup c'$ supports a (justified/defensible) argument for the conclusion that p does not violate a's interests and $KB_d(a)$ does not support a (justified/defensible) argument for the conclusion that $(q \wedge c')$ is forbidden.
  - If there is one such set c', then make a counteroffer $q \wedge c'$.
  - If there are more than one such sets, make a counteroffer $q \wedge c'$, where c' is the subset-minimal set that has the most elements in common with c.
  - Otherwise reject the offer p.

The negotiation policy contains a number of design choices. Of course, in other domains a choice for other types of agents can be made. For example, the agent could also reject the request for information if there is no justified argument which says he is obliged to give the information. However, this action is not in line with the overall goal of the system, which says that if at all possible, information should be exchanged. Furthermore, in our design of the policy we have not yet specified whether the agent needs justified, defensible or just well-formed arguments for his conclusions. This can have a significant impact on the behaviour of the agent. For example, an agent who accepts well-formed arguments that say that it is not forbidden to give a certain piece of information will be much easier persuaded than an agent who only accepts justified arguments for the conclusion that giving the information is not forbidden. So the first agent, who only needs a well-formed argument, will not be as protective of his information as the second agent, who needs a justified argument. Different policies are also possible with respect to making a counteroffer. We have described a policy where having as few extra conditions as possible is more important than making a counteroffer which includes the conditions offered by the opponent. In our policy it is thus possible to delete some or all conditions of the opponent's offer. Another policy would be a policy where an agent can only add conditions to the original offer. Another choice option is in how an agent should internally reason about whether an offer is in the agent's interests. Is the reasoning only about his own actions or also about the actions of the other agent? The different design options mentioned here are mostly empirical questions, to be answered by domain analysis.

Responding to a reject
The next policy describes responding to a reject.

- First respond with a *why-reject* move. If the resulting persuasion subdialogue is won, then it is the other agent's turn. Otherwise,
- if the reject move responded to the initial offer then reply with a withdraw, while
- if the reject move responded to a counteroffer, backtrack to the target of the reject move.
  - If an alternative counteroffer exists that satisfies the policy for responding to an offer then make it.
  - Otherwise reply with a withdraw.

5.2.3.2 Persuasion policies
We now turn to a less formal specification of persuasion policies, which determine how an agent should respond to arguments and *why* moves. Other persuasion policies can be developed similarly. Recall that our policies assume that an agent *a* reasons with $KB_d(a)$, which is his initial internal knowledge base, possibly modified by his *concede* and *retract* moves during the dialogue. We will first explain how an agent can react to an argument.

Responding to arguments
Can the agent construct a (justified/defensible) counterargument?

- 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/defensible) 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 opponent's 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 a (justified/defensible) argument for *p*?

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

A few things should be noted regarding the above policies. Firstly, as in the negotiation policies, we have not specified what kind of arguments (justified or defensible) an agent needs for his decisions. Secondly, the agent is cooperative in that he only asks *why p* questions if he 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, who is considered to be a reliable source of information concerning X.

6 Illustration of the proposed architecture

We now 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 we make no further assumptions on the language or logic. We also assume that agent *b* needs defensible arguments. The knowledge of the agent are all the rules and facts known by him about a requested piece of information. Up to step 9 each move replies to the immediately preceding move.

The responding agent *b* has the following knowledge:

$\mathbf{KB_0(b) = \{ (\neg e \Rightarrow \neg i),\ (\neg i \Rightarrow g),\ (p \Rightarrow \neg a),\ p,\ n\}}$

where $\boldsymbol{e}$ = "the information may be exchanged with other police officers", $\boldsymbol{i}$ = "the informant is in danger", $\boldsymbol{g}$ = "you give all information about trading in explosive materials", $\boldsymbol{p}$ = "the information is protected", $\boldsymbol{a}$ = "agent *b* is allowed to give information about trading in explosive materials", $\boldsymbol{n}$ = "the information is a matter of national importance".

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

1 - A: request *g* (Give me all information about trading in explosive materials).

According to the negotiation policy outlined in section 5.2.3.1, agent b will first have to look whether he is obliged to give the information or not. He cannot construct a defensible argument from his knowledge stating that he is obliged to give the information, so b applies the next step in his policy, which is finding out whether b is allowed to give the information or not. b can construct an argument "¬a since p" which states that b is not allowed to give the information, so b rejects the request.

2 - B: reject *g* (I will not give you this information)

3 - A: 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.

4 - B: claim ¬*a* (I am not allowed to give you the information)

5 - A: why ¬*a* (Why are you not allowed to give me the information?)

According to agent b's persuasion policy, he has to move defensible arguments so he moves argument ¬a since p.

6 - B: ¬*a* since *p* (I am not allowed to give the information since the information is protected)

Say that agent a rebuts p ⇒ ¬a and defeats argument p ⇒ ¬a.

7 - A: *a* since *n* (The information is allowed to be shared because this is a matter of national importance)

b cannot ask a "why n" question here because b can construct a (trivial) defensible argument for n since n is part of his knowledge base, so agent b has to concede to a.

8 - B: concede *a* (Ok, I admit that in this case the information is allowed to be shared)

B retracts ¬a from his commitments because *his commitments have to stay consistent.*

9 - B: retract ¬*a* (I retract the fact 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 that he is allowed to give the information. b now has to look if there are extra conditions serving his interests under which the information can be given. Through backward chaining through the rules (¬i ⇒ g) and (¬e ⇒ ¬i), b arrives at the possible extra condition ¬e and internally verifies that this extra condition does not make the offer he has in mind violate his interests. With B6 the responding agent then backtracks to A1, this time replying with a counteroffer.

10 - B: offer *g* ∧ ¬*e* (I will give you the information on the following condition: the given information may not be exchanged with other police officers).

We assume that the local interests of agent a are not violated and therefore he accepts the offer.

11 - A: accept *g* ∧ ¬*e* (I agree with this condition)

7 Conclusion

In this paper we have proposed a multi-agent architecture for regulated information exchange and we have illustrated it with examples of 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 [6] 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 [2] 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 [5] 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 [3] 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 informal and semiformal parts of our architecture specification should, of course, be made fully formal and then implemented. With respect to the negotiation policies we intend to investigate alternative design choices on the various aspects discussed in Section 5.6.3. 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 [5], occur in our application domain and should therefore be modelled in our architecture.

8 Acknowledgements

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] Atkinson, K. Bench-Capon, T.J.M. & McBurney, P. Arguing about cases as practical reasoning, *Proceedings of the Tenth International Conference on Artificial Intelligence and Law (ICAIL-05)*, 35-44. New York: ACM Press, 2005.

[2] 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.

[3] Doutre, S., McBurney, P. and Wooldridge, M. Law-Governed Linda as a semantics for agent interaction protocols (research abstract). In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)*, Utrecht, The Netherlands, July 2005: 1257-1258.

[4] McBurney, P. and Parsons, S. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information* 11 (2002): 315-334.

[5] Parsons, S., Sierra, C. and Jennings, N. Agents that reason and negotiate by arguing. *Journal of Logic and Computation* 8 (1998): 261-292.

[6] 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.

[7] 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.
[8] Prakken, H. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, to appear.
[9] Prakken, H. and Vreeswijk, G. Logical systems for defeasible argumentation. In D. Gabbay and F. Guenthner (eds.), *Handbook of Philosophical Logic*, second edition, Vol 4, pp. 219-318. Kluwer Academic Publishers, Dordrecht etc., 2002.
[10] Van Veenen, J. and Prakken, H. A protocol for arguing about rejections in negotiation. *Proceedings of the Second International Workshop on Argumentation in Multi-Agent Systems (ARGMAS 2005)*, Utrecht, The Netherlands, July 2005.
[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.