

Is it worth arguing?

Nishan C. Karunatilake and Nicholas R. Jennings

School of Electronics and Computer Science, University of Southampton,
Southampton, UK.

{nnc02r, nrj}@ecs.soton.ac.uk

Abstract. Argumentation-based negotiation (ABN) is an effective means of resolving conflicts in a multi-agent society. However, it consumes both time and computational resources for agents to generate, select and evaluate arguments. Furthermore, in many cases, argumentation is not the only means of resolving conflicts. Thus, some could be avoided either by finding an alternative means (evading the conflict) or by modifying the intended course of action (re-planning). Therefore, it would be advantageous for agents to identify those situations and weigh the costs and the benefits of arguing before using it to resolve conflicts. To this end, we present a preliminary empirical analysis to evaluate the performance of a simple ABN system, with respect to other non-arguing approaches, in a particular task allocation scenario. In our experiments, we simulate a multi-agent community and allow the agents to use a combination of ABN, evasion and re-planning techniques to overcome conflicts that arise within the community. Analysing the observed results, we show that, in our domain, ABN presents an effective means of resolving conflicts when the resources are constrained. However, we also show it is a more costly and less effective means, compared to evasion and re-planning methods, when resources are more abundant.

Key words: Argumentation-based Negotiation, Conflict Resolution

1 Introduction

Conflicts are inevitable in multi-agent systems in which autonomous entities pursue their own goals (whether they do so in a self-interested or in a collaborative manner) [1]. They cover physical conflicts arising due to resource limitations (e.g., multiple agents attempting to use a non-shareable resource at the same time) and knowledge conflicts resulting due to discrepancies in viewpoints or opinions (e.g., a contradiction between agents' beliefs about a particular proposition) [1–3]. In either case, however, they present hurdles for the agents to overcome if they are to achieve their goals and actions in a coordinated manner. Against this background, *Argumentation-Based Negotiation* (ABN) is advocated as a promising means of interaction that can allow the agents to resolve these conflicts [4]. In its simplest form, ABN allows agents to exchange proposals that are accompanied by meta-information, which provides support and justification for the proposals. It also allows the exchange of explicit arguments, such as critics, appeals and other forms of persuasive locutions, to influence and persuade the opponent to accept the proposals and come to a mutual agreement [4–6].

Although ABN can be effective at resolving conflicts, there are a number of overheads associated with its use. It takes time to persuade and convince an opponent to change its stance and yield to a less favourable agreement. It takes computational effort for both parties of the conflict to carry out the reasoning required to generate and select a set of convincing arguments, and to evaluate the incoming arguments and reason whether to accept or reject them.

However, not all conflicts need to be resolved. Thus, for example, when faced with a conflict, an agent could find an alternative means to work around the situation; thereby *evading the conflict* rather than attempting to resolve it. By way of an example, consider the case where an agent (A) requires the service of another (B) which is also demanded by a third agent (C). Now if B is unwilling to provide its service, instead of attempting to persuade it to change its conflicting stance, A could simply attempt to find another more willing partner (D) who has a similar capability. The result would still be A overcoming the conflict situation, but not through argumentation. In addition to either evading the conflict or arguing and resolving it, an agent could also attempt to *re-plan and alter the means* by which it intends to achieve the objective so that the conflict situation is removed (e.g., A could delay its task until B becomes available).

Given the overheads of argumentation, and the alternative methods available for overcoming conflicts (evade and re-plan), we believe it is important for agents to be able to weigh up the relative advantages and disadvantages of arguing, before attempting to resolve conflicts through argumentation. This is the main long-term motivation of our research. Specifically, we aim to empirically evaluate the effectiveness and efficiency of argumentation as a conflict resolution mechanism with respect to these other non-arguing alternatives available to the agents. To date, this issue has largely been overlooked in existing literature. Current ABN assumes that the agent has already made the decision to argue (typically without any consideration) and the focus is on the internal mechanisms of argumentation (i.e., how agents can generate, select and evaluate arguments). Our work presents an initial step in this direction.

Against this background, this work advances the state of the art in the following ways. *First*, our main contribution is to evaluate the relative effectiveness and efficiency of using simple forms of ABN, as opposed to evasion and re-planning, to overcome conflicts in a multi-agent system. Specifically, we consider an ABN system in which agents exchange meta-information, alongside their decisions, either to explain the internal constraints that prompt them to make their decisions, or to suggest alternative solutions that satisfy their internal constraints (e.g., I reject this proposal, since I am fully committed at this time or I reject this proposal for the suggested time, but for this price I am willing to perform this service at the following alternative times).¹ Through an empirical evaluation, in an idealised task allocation scenario, we show that such ABN does indeed present a better means of conflict resolution than evasion when the resources are constrained. However, we also demonstrate the diminishing impact (both in effectiveness and in efficiency) of the ABN method as the resource levels increase within the community. *Second*, we demonstrate the superior performance of hybrid strategies (i.e., those that use both ABN and evasion in a combined manner) as opposed to pure strategies that always attempt to use either one or another in conflict resolution. *Third*, to empirically illustrate our concepts, we present a simple, but well-defined multi-agent context, where conflicts occur naturally through interaction of agents with different motivations. Even though, our experimental context embodies a series of simplifying assumptions made for implementation purposes (detailed in Section 3.1), we demonstrate its versatility by replicating both Kraus et al.'s [5] and Jung et al.'s [7] main experimental observations.

The remainder of the paper is structured as follows. Section 2 discusses the related work and establishes our contribution within the current literature. Section 3 details our argumentation context, the conflicts arising within it and presents the different methods and strategies used by the agents to resolve these conflicts. Subsequently, Section 4 details the experimental

¹ Clearly, this is toward the simpler end of the possibilities in argumentation. However, our purpose here is not to exhaustively cover all forms of argumentation. Rather we seek to evaluate the trade-offs involved in engaging in argumentation and concentrating on the simpler models provides an initial point of departure.

setting, presents our results and an analysis of the key observations. Section 5 concludes, and details our future directions.

2 Related Work

Argumentation-based negotiation is fast emerging as an important means of interaction for agents within multi-agent communities [4]. To date, most of the work in this area has focused on the internal mechanisms of argumentation; that is how arguments are generated [6, 8, 9], selected [5, 10, 11] and evaluated [12, 13], and how the process of argumentation can resolve conflicts and achieve agreements [7, 14]. However, no real attention is given to the overall impact of the decision made by the agents to resolve their conflicts by arguing. Rather, it is simply assumed that the agent has already made that decision and the focus is on how the agent can use arguments to resolve the conflict. Thus, unanswered questions remain such as when to use argumentation?, under what conditions does it yield better results than non-arguing strategies?, and what are its implications for the performance of the multi-agent community?

In tackling this problem we draw inspiration from a number of previous efforts in the ABN literature. Specifically, Jung, Tambe and Kulkarni's empirical work [7] acted as an important impetus for our effort. Their work attempts to evaluate the overall impact of using meta-information within a negotiation process to resolve conflicts. To do so, this work models a set of collaborative agents attempting to solve a distributed constraint satisfaction problem (DCSP) [15] and it maps the DCSP into an argumentation context. More specifically, the conflicts are mapped as external constraints affecting the local variables in the DCSP, the pure negotiation process involves the exchange of values for these internal variables, and the meta-information (argument) exchange is mapped as the propagation of internal constraints. Motivated by the desire to resolve the DCSP, the agents can either interact to resolve these conflicts via pure negotiation (without arguments) or ABN. However, the main motivations of our work are quite different from theirs. In particular, their work assumes that all conflicts need to be resolved, and thus they compare ABN to negotiation without argumentation in order to assess the impact of meta-information exchange on the conflict resolution effort. In contrast, we do not believe that all conflicts need to be resolved because they can sometimes be avoided through evasion or re-planning. Therefore, our motivation is to evaluate the importance of ABN as a conflict resolution mechanism as opposed to using other non-arguing means to overcome conflicts.

Kraus, Sycara and Evenchik [5], to a limited extent, consider whether argumentation should be used when faced with a conflict situation. They use a fixed heuristic to enable the agent to decide when to argue and when to stop the argument and re-plan. In their experiments, two self-interested agents are assigned a particular task, which neither has the capability to achieve alone. Thus, the agents must cooperate to achieve the task. The mechanism of achieving cooperation is by using negotiation and persuasion dialogues. According to their heuristic, the agent will *always* first try to argue and reason with the other party and try to achieve an agreement. However, if the agent is unsuccessful in achieving an agreement in a given fixed time schedule, it will stop the argument. In the next time slot it will re-plan, generate a new set of goals and intentions, and will start the process all over again. However, this heuristic is rather rigid and is but one possibility. Moreover, it was tested in a two-agent context where the only option available to an agent was to make the other agent agree (otherwise, it could not complete its task). Generally speaking, when there are only two agents, the alternative options available for the arguer are limited. Thus, the *always argue* approach becomes more viable. Avoiding conflicts is not a possibility, because the agent that wants to achieve the task has to somehow convince the only other agent within the system to provide its services. However, its usage in

a multi-agent context, where there are many other potential alternative agents that might be willing to cooperate, is questionable.

3 The Argumentation Context

To evaluate the overall performance of argumentation as a means of conflict resolution, we require a computational context in which a number of agents interact and conflicts arise as a natural consequence of these interactions. To this end, Section 3.1 presents an overall description of the experimental setting, clearly specifying the task environment, which presents the agents with the motivation to interact. Subsequently, Section 3.2 explains how these interactions give rise to conflicts and then proceeds to explain the three different methods the agents can use to overcome them; namely *argue*, *evade* and *re-plan*. Finally, Section 3.3 details the strategies that agents use to combine these three methods for conflict resolution.

3.1 The Scenario

The scenario simulates a collection of self-interested agents, each with a specific capability and a specific task to achieve. Each task requires a particular series of actions to be achieved in a predefined order, and each action requires a specific capability. However, none of the agents possess the capability to achieve all their actions, thus they need to negotiate for the services of one another. When an agent manages to attain all the capabilities required to execute its actions in the predefined order, the task is completed. Upon completion of the task, the agent receives a specific reward. It is this reward that motivates the self-interested agents to complete their tasks, which, in turn, results in agents interacting within the system.

Time Slot	A (α)	B (β)	C (γ)
	£6,000	£4,000	£10,000
TS0	α	β	β
TS1	β	α	β
TS2	γ	α	α
TS3	α	β	γ

Table 1. A Sample Problem: Presents a three agent society, each having their own capability and their assigned task schedule.

In more detail, Table 1 depicts a sample scenario of a multi-agent community with three such agents; namely A, B and C. Agent A has the capability to perform the action α , while B and C are capable of performing β and γ respectively. Each task is presented as a series of actions. For example, agent A's task involves four actions, which requires capabilities α , β , γ and α respectively. The notion of time is an important parameter in the scenario. Not only must agents achieve their actions in the specified order, but also they need to achieve them in the specified time. Any delays on this time will incur a penalty charge (this penalty calculation is discussed later). All agents operate to a unified clock and an atomic unit of time is termed a time slot. For example, A's task spans four time slots TS0 to TS3. Thus, for A to attain the complete £6,000 reward, it will have to find capable agents to perform α , β , γ and α at TS0, TS1, TS2 and TS3.

How the agents interact to find their task partners is a central issue in this work. In the simplest case, when an agent needs to find a certain capability to achieve some action for a specific time slot, it will first look to see if it possesses the necessary capability to perform the action on its own. If it does so, it assigns that action to itself. However, if it does not possess the required capability, it must attempt to convince another agent to sell its services for that specific time slot.² In the above example, agent A does not have the required capability to perform the action at TS1 (since it does not possess capability β). Therefore, it will attempt to convince another agent B (who has capability β) to sell its services for the time slot TS1.

If an agent does not manage to convince any of its known acquaintances to sell it their service, it has to *delay* that action. Delaying means, it will not accomplish any action within that time slot. Since the agents need to achieve their actions in the strictly prescribed sequence, adding these delays naturally lengthens the time required to accomplish the task.³ As mentioned above, any task completed after the initially assigned time incurs a *penalty*, which, in turn, reduces the task's reward available for the agent upon completion. The amount of penalty is a fixed value per extended time slot and is proportional to the task's initial reward. However, if the agent loses all its initial reward as penalty charges, any further delays will not incur any additional charge. This is an implementation choice made to prevent agents incurring greater penalty charges than their initial allocated reward:

$$\text{Penalty} = \begin{cases} \frac{R_{init}}{T_{init} * mdf} & \text{if } T_{init} < T_{ext} < (T_{init} * mdf), \\ 0 & \text{if } T_{ext} \leq T_{init} \parallel T_{ext} \geq (T_{init} * mdf) \end{cases} \quad (1)$$

where;

- T_{ext} is the extended task duration taken to achieve the task.
- T_{init} is the initial allocated task duration.
- R_{init} is the assigned task reward.
- mdf is the maximum delay factor, which is a constant for all agents in our case.⁴

If a certain agent (in the above example B) agrees to provide its services to a specific agent (A) for a particular time slot (TS1), B will not be able to agree to provide any other action for TS1, unless it cancels its current agreement with A. For example, if C requests B to perform its action, which requires capability β (refer to Table 1) at TS1, it cannot do so unless it reneges on its current contract with A. Our framework allows agents to *renege upon their agreements* if they perceive a more profitable opportunity.⁵ This ability to renege current agreements is important because it promotes opportunities for the agents that seek services later in the scheduling process to achieve agreements if they are willing to pay sufficiently high premiums for these services.

² It is worth noting that in certain situations, even though the agent does possess the capability to accomplish its own action, it may find it more rewarding to find another to perform it. This may occur, if the agent has already agreed to sell its services to another, and it is more rewarding for it to maintain this agreement than to pay another agent to perform its action.

³ Here a delay slot is inserted in place of $TS1$, and the action β at $TS1$ will be scheduled at $TS2$. This process would result in the shift of all subsequent actions by one time slot.

⁴ For example, an agent with a task worth £10,000 spanning 50 time slots, and an mdf set to 4, will incur a penalty of £50 ($\frac{£10000}{(50*4)}$) per each additional time slot taken to complete the task. If the agent takes more than 200 ($50 * 4$) slots its reward would be zero, and, thereafter, it will not incur penalties.

⁵ At this time, the agents do not incur an extra charge for reneging upon their agreements. As explained in Section 5, we aim to investigate these effects in our future experiments.

In this scenario, the main objective of the agents is to maximise their individual earnings. There are two methods of doing so. First, they can complete their assigned tasks. Once an agent completes its task, it will receive the allocated reward (less the penalty charges due to delays). This we term the agent’s *task earnings* (TE). Second, they can sell their services to other agents (which we term the agent’s *service earnings* (SE)). Both components contribute toward the overall *individual earnings* (IE) of the agent.

$$TE = R_{init} - \sum(\text{Penalty}) - \sum(\text{External Service Payment}) \quad (2)$$

$$SE = \sum(\text{External Service Earning}) \quad (3)$$

$$IE = TE + SE \quad (4)$$

Given an overall description of the scenario, we, however, make a number of simplifying assumptions. First, we assume that each agent within the system has complete and accurate knowledge of its own task (i.e., its reward, the actions required, and the sequence in which they need to occur to achieve the task). Thus, during the interaction, the service providers would not be able to give any new information about the task that the buyer would not already know, or be able to convince the buyers on anything contrary about their task specification. For example, the sellers won’t be able to suggest that the actual task is worth less than its initial estimate or be able to recommend different sequences of actions (other than the one specified) to achieve the same task. Second, we assume that the agents are truthful when they communicate information to others, and do not attempt to deceive them into making incorrect decisions. Third, we assume the interactions consist of single encounters, thus, issues such as trust and reputation do not have a material effect within the context.

Even though, all the above are real issues present in multi-agent environments, our motivation for excluding them from the initial experiments is to attain simplicity within the argumentation context. Our desire is to design a context that is simple, yet expressive enough to simulate conflicts and methods of overcoming them (i.e., argue, evade and re-plan as explained in Section 3.2), but not to simulate the most sophisticated simulations of these behaviours. Additionally, excluding these parameters limits the variability present in the system. This allows us to predict more accurate hypotheses about the system, gain a better understanding of the dynamics of the multi-agent interaction, and explain the reasons for the observations with more ease. Given the broad overview of the multi-agent scenario and the assumptions made, we now proceed to explain how interaction within the context leads to conflicts and the three distinct methods used to overcome them.

3.2 Conflicts and Methods of Resolution

The self-interested motivations of our agents give rise to conflicts within the system. Thus, when agents attempt to acquire the services of another, they are motivated to pay the lowest amount they possibly can for that service. This is because the lower an agent’s external service payments are, the higher its own TE will be (equation 2). However, on the other hand, when agents sell their services, they are motivated to attain the highest payment they possibly can to maximise their SE (equation 3). Thus, whenever an agent attempts to convince another to provide its services, it naturally gives rise to conflicts of interests between buyer and seller agents within the system.

The dynamics of interaction become more complicated due to the presence of penalty charges and the ability of agents to renege on their present agreements. Since agents are motivated to maximise their TE, they want to avoid penalties (equation 2). However, if a buyer

is only willing to offer a very low reward for the service, it is more likely to be rejected, and, in turn, stands a higher chance of incurring a penalty. This motivates the agent to make high rewarding proposals. Secondly, because sellers can renege on their present agreements if they receive more rewarding proposals, agreements made at low values are more likely to be revoked than higher rewarding ones. This may also motivate buyers to make higher rewarding offers to ensure their agreements are more secure. Together these opposing motivations dynamically generate conflicts within the system⁶ providing a good context to test the performance of our various methods for overcoming conflicts.

Following presents the three distinct methods we use to overcome these conflicts:

1. **Argue: Use ABN to resolve conflicts**

When an agent requires a capability from an acquaintance, it generates a *proposal* and forwards it to an agent who has that capability. Once received, the agent evaluates the proposal and decides whether to accept or reject it. The agent will then, communicate its decision, either as an *acceptance* or as a *rejection*, to the original agent. If it decides to accept, the interaction ends in an agreement. However, if the decision is to reject, the onus is transferred back to the original buyer agent to generate and forward an alternative proposal. To help this interaction process, the seller agent, will accompany its rejection with two additional forms of meta-information (arguments) that it will convey back to the original buyer agent:

- *Reasons for refusal*: This details the reason that prompted the refusal. In our system, seller agents reject due to two types of reasons. First, the agent may be fully committed to a prior arrangement in the requested time slot, so it returns an argument indicating that the reason for failure is unavailability (rather than the offer price being too low). Second, the offer value may not be sufficiently valuable to the agent, in which case it will return an argument accompanied with its rejection indicating the minimal threshold that must be exceeded before the proposal will be considered. The return of such arguments should assist the buyer in its attempt to choose the next proposal to forward. For example, if the reason is unavailability, the buyer would not make an increased value proposal since doing so would be futile. On the other hand, if the threshold is returned as reason for refusal, the buyer can use this to gauge whether to make another proposal to that agent and if it does then value that should be used in such circumstances. These form of arguments are analogues to the types of meta-information exchanged in Jung et al. [7].
- *Alternative suggestions*: If the seller is willing to work for the suggested value of the offer, but not in the proposed time slot, it will send a number of its neighbouring time slots as alternative suggestions. This meta-information helps the buyer agent in finding agents for those future time slots. For example, assume that in the attempt to find a partner for TS1, agent B indicates to A that it is willing to work for TS2 as an alternative. If agent A requires the same capability for the same price (the price offered when it got the alternative) in TS2, before requesting other random agents, A will first ask B who has already expressed its willingness. Thus, alternatives provide agents with information about their partners' schedules, which they will, in turn, use to selectively choose the order (instead of strictly adhering to a random one) in which they request their partners.

⁶ Here we consider only one form of conflict; namely conflicts resulting from discrepancies of interests. Conflicts of knowledge due to discrepancies of viewpoints or opinions are not considered in this work.

If any such proposal results in an agreement the argue method is said to have succeeded in its objective. However, if all possible proposals fail to make an agreement the argue process ends in failure.

2. ***Evade: Find an alternative method to achieve the same plan***

Unlike the previous method, here the agent does not attempt to use ABN to resolve its conflicts. The agent will only make a single proposal. This is to establish the willingness of the potential partner. If that offer is rejected the agent will not attempt to convince the non-willing partner, but will move on to the next known acquaintance, which has the required capability. However, in this scenario the buyer chooses to offer the maximum price it can in its single proposal. The rationale for this choice is to maximise the chances of success of its single proposal, thus this represents the maximally effective evade strategy. Since the sellers are always motivated to accept higher offers (equation 3), making the highest offer possible maximises the chances of success in its single proposal. If the seller refuses this proposal the evade method fails. On the other hand, if it accepts, then evade method succeeds.

3. ***Re-plan: Change the original plan***

When a conflict arises at a particular time slot, the buyer agent simply places a delay slot in its schedule and tries to arrange for the desired capability to be scheduled to the next time slot. This delay the whole sequence of remaining activities, thus, will extend the task's overall duration by one time slot. While the argue and evade methods remain the main methods in our strategies, re-plan represents the fall back option (refer to Section 3.3). Thus, re-planning through delays (since theoretically an agent can delay forever) will always ensure success in overcoming any specific conflict. However, delays may cause subsequent conflict situations to arise and will render the task less rewarding via penalties.

3.3 Conflict Resolution Strategies

In this section, we presents six different strategies for conflict resolution which differ in terms of the way they order the argue, evade and re-plan methods. These strategies are defined to give a range of different behaviours in resolving conflicts in a multi-agent context. However, they are neither meant to be the most optimal, nor an exhaustive list. Rather their designed purpose is to allow us to perform a comparative analysis of the relative performance of arguing versus evasion in conflict resolution.

- ***Evade 1:*** Randomly select one agent. *Evade* with that agent. If fail, re-plan.
- ***Argue 1:*** Randomly select one agent. *Argue* with that agent. If fail, re-plan.
- ***Always Evade:*** Randomly select one agent at a time and *evade*. Continue *evade* till either an agent agrees or the last agent is reached. If fail with last agent, re-plan.
- ***Evade Finally Argue:*** Similar to *Always Evade*, thus, continue to *evade* till penultimate agent. However, with the *last* agent *argue*. If fail with the last agent, re-plan.
- ***Argue First then Evade:*** Similar to *Always Evade*, but *argue* with the *first* agent. If fail with this agent continue *evade* till either an agent agrees or last agent is reached. If fail with last agent, re-plan.
- ***Always Argue:*** Similar to *Always Evade*, but in *all* encounters *argue* till either an agent agrees or the last agent is reached. If fail with last agent, re-plan.

From the above, *Evade_1* and *Argue_1* only allow the agents to interact with a single partner. Strategies *Always_Evade* and *Always_Argue* allow agents to interact with all potential partners (one at a time). However, they only allow the agents a single method to resolve conflicts (either evade or argue), thus are termed pure strategies. In contrast, *Evade_Finally_Argue* and *Argue_First_then_Evade* are hybrid strategies that selectively use argumentation with evasion; the former gives priority to evasion, while the later gives priority to argumentation. Having introduced our argumentation context, we now turn to our empirical evaluation.

4 Experimental Evaluation

The aim of these experiments is to evaluate the overall effectiveness and efficiency of using a simple ABN, as opposed to evasion and/or re-planning, to overcome conflicts in our chosen scenario. In particular, we simulate a multi-agent context (as per Section 3.1) and endow the agents with different resolution strategies (as per Section 3.3). The observed overall performance of the society is measured and used to carry out a comparative analysis between these strategies.

4.1 Experimental Setting

The experiments are set within a society of 75 agents, each having one out of three capabilities (α , β or γ). These capabilities are equally distributed within the society with 25 agents per capability. All agents are assigned a single task spanning 50 time slots. Each time slot contains a single action that requires a single capability. These actions are randomly distributed within a task. The initial rewards for the tasks are set according to a normal distribution with a mean £10,000 and a standard deviation of £2,500. The *mdf* parameter (equation 1) for the penalty charge is set to 4 (based on initial experiments).

In each experiment, the society differs in terms of its resource settings (RS). In the maximum resource setting (RS25), each agent knows about all the other agents, hence it has maximum access to the resources within the system. On the other hand, in the most constrained setting (RS1), each agent is only aware of the existence of a single (randomly selected) agent per capability. In between we define a series of 12 intermediate settings, where each agent is aware of the existence of 2, 4, . . . , 24 (referred to as RS2, RS4 etc.) other agents per capability. Thus, for example, at RS4, each agent is aware of the existence of 4 other agents with capability α , 4 with β and 4 with γ . We use the following metrics to evaluate the overall performance of the different strategies [7, 10]:

- **Effectiveness of the Strategy**

We use the *total accumulated penalty* incurred by all agents within the society as a measure of effectiveness. If this value is low, the strategy has been effective in handling the conflicts that have arisen.

- **Efficiency of the Strategy**

This reflects the computational cost of interaction incurred by the society, while using a particular strategy to resolve conflicts. As interaction takes longer, more resources are consumed by the agents. On the other hand, these longer interactions also increase the number of messages. Thus, the *total number of messages* provides us a good method to measure computational resources used by the agents during interaction. This covers the messages used to overcome conflicts and reach agreements (including reasons and alternatives exchanged as meta-information), and the messages associated with renegeing from

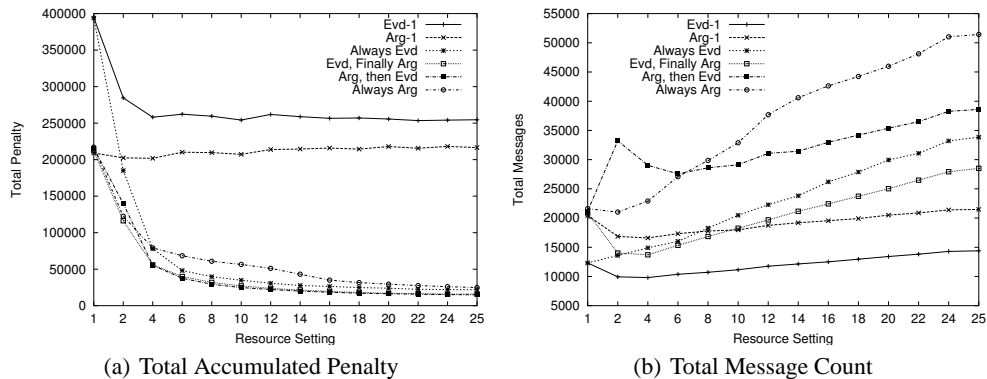


Fig. 1. Variation of Total Penalty and Total Messages with different resource settings

agreements. In this context, a strategy that involves fewer messages is said to have performed more efficiently than one that uses a higher number.

4.2 Results and Observations

Given these experimental settings, we can now turn to the actual results. Here all reported results are averaged over 50 simulation runs to diminish the impact of random noise, and all observations emphasised are statistically significant at the 99% confidence level.

Observation 1: *In highly resource constrained settings, argumentation significantly enhances the overall effectiveness of the society.*

In Figure 1(a), we observe that at highly resource constrained levels (i.e., RS1 and RS2), the strategies that use argumentation to resolve conflicts (namely *Argue_1*, *Evade_Finally_Argue*, *Argue_First_then_Evade* and *Always_Argue*) incur a significantly lower penalty charge than those that merely evade (i.e., *Evade_1* and *Always_Evade*). The impact is most apparent in RS1, where the resources are most constrained. The difference is approximately of a magnitude of 1.84 (i.e., *Evade_1* and *Always_Evade* have an average penalty of £394,250, whereas *Argue_1*, *Evade_Finally_Argue*, *Argue_First_then_Evade* and *Always_Argue* have an average of £213,487). Although slightly reduced, this effect is also observable in RS2 approximately a magnitude of 1.41 between *Evade_1* and *Argue_1*, and 1.47 between *Always_Evade* and *Evade_Finally_Argue*, *Argue_First_then_Evade* and *Always_Argue*. In such scarce resource settings, the number of alternative solutions available to the agent to overcome conflicts is highly constrained. Due to the absence of such alternatives, the evasion techniques (*Evade_1* and *Always_Evade*), tend to fail more as they evade conflicts in search of the non-existent alternatives and thereby incur higher penalties. On the other hand, strategies that attempt to resolve the conflicts through ABN tend to form more agreements and, thus, incur fewer penalty charges.

Further support for this observation can be drawn by comparing the behaviour of strategies *Evade_1* and *Argue_1* over all resource settings. Both of these strategies attempt to overcome conflicts by interacting with a single randomly chosen partner. Although from the outset this does not appear to be a very prudent strategy (constraining oneself to a single partner when there are more potential partners available) these agents were specifically designed to experiment with the relative impact of using argumentation in resource-constrained settings. To this end, Figure 1(a) shows how *Argue_1* continuously incurs low penalties than *Evade_1*. Since these strategies constrain the agents to interact with just a single partner, irrespective of how

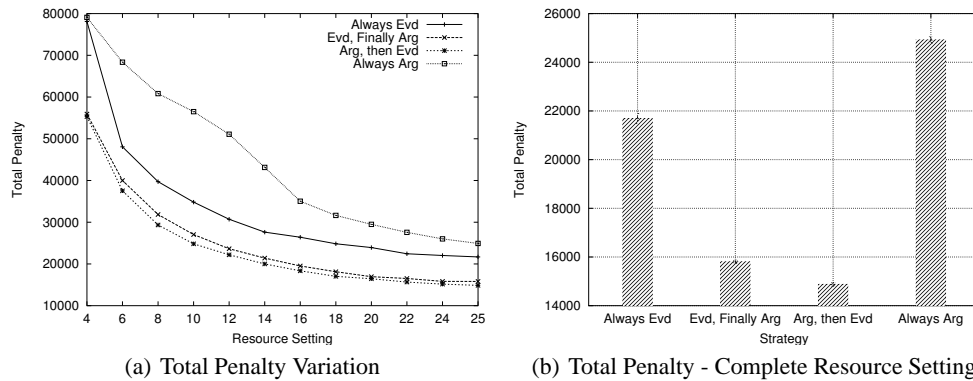


Fig. 2. Magnified Penalty Variations for the high resource settings

much resources are available to them, the agents still operate in limited resource settings. Thus, the alternatives available to them are limited. These observations further justify our conclusion that using ABN to resolve conflicts tends to be a more effective method than evasion in resource-constrained settings. This finding is also consistent with the experimental results observed by Kraus et al. [5], where they presented the benefits of the *Always Argue* strategy in a two agent setting.

Observation 2: *As resource levels increase, both the argue and evade methods become more effective, but the relative difference between them decreases.*

Figure 1(a) also shows that the penalty charges for the strategies *Always Evade*, *Evade Finally Argue*, *Argue First then Evade* and *Always Argue* reduce as resource levels increase. This effect is seen more clearly in Figure 2(a), which presents a magnified view of the penalty variation for these four strategies. The primary reason for these reductions is the increase in resource level. Thus, as resources increase, so does the potential to find an alternative agreeable partner. Thus, a higher number of conflicts can be overcome, which, in turn, reduces the delay. The net result being a reduction in penalty charges for all strategies.

Arguably, a more interesting observation is the differences in the rate of penalty reduction for the strategies that use argumentation and the ones that use evasion. Specifically, the penalty charge of *Always Evade* decreases more rapidly than *Always Argue*. Figure 2(a) shows *Always Evade* surpassing *Always Argue* between RS4 and RS6 and thereafter maintaining its performance. The reason for this difference is as follows. As the potential alternatives increase within the society, the need to convince a non-willing partner decreases. Arguing strategies, which attempt to convince their non-willing partners before attempting to search for these alternatives, do not use these options to the same degree as evasion strategies do, which explains the observable differences in the rate reduction between *Always Evade* and *Always Argue*. Furthermore, Figure 1(b) shows evasion strategies using a lower number of messages than arguing ones. This is because unlike evasion strategies, arguing strategies in their attempt to convince non-willing partners tend to use more messages in their interaction. Thus, even when both arguing and evasion strategies are equally effective, evasion strategies tend to be more efficient. This observation allows us to conclude that as resources become more abundant, evasion increasingly becomes the more preferable option.

Observation 3: *Using argumentation indiscriminately has a negative impact on the systems' overall effectiveness*

Strategy	Total Messages		Total Penalty (£)	
	Mean	Std Div	Mean	Std Div
Evade.1	14397.7	142.95	254634.0	9113.30
Argue.1	21473.4	274.07	216523.0	7913.68
Always Evade	33836.8	1347.78	21688.5	1452.01
Evade, finally Argue	28500.3	361.04	15800.8	439.64
First Argue, then Evade	38607.7	578.20	14873.9	445.52
Always Argue	51425.3	1188.25	24918.7	866.41

Table 2. Summarised Penalty Charges and the Message Counts for the complete resource setting.

Figure 2(a) also allows us to compare the performance of strategy *Always Argue* versus *Evade Finally Argue* and *Argue First then Evade*. Unlike the selective argumentation used by *Evade Finally Argue* and *Argue First then Evade*, *Always Argue* indiscriminately argues in all interactions. However, in both Figure 2(a) and 2(b) it can be seen that *Always Argue* incurs a higher penalty value than those strategies that selectively argue.

To help us analyse the reasons for this effect, Figure 3 presents the number of conflicts for all six strategies in RS25. These conflicts are divided into two sections; namely, the primary conflicts that arise when the agents first attempt to find partners and the secondary conflicts that arise due to agents reneging upon their agreements. It can be observed that the strategies *Always Evade*, *Evade Finally Argue*, *Argue First then Evade* and *Always Argue* incur approximately the same number of primary conflicts. However, the strategies *Evade Finally Argue* and *Argue First then Evade*, which give priority to the argue method, incur a significantly higher number of secondary conflicts. The reason being when agents argue to form agreements, they manage to convince the sellers to make lower price agreements. However, another arguing agent can potentially come forward and, using ABN, negotiate a higher valued contract, which breaks the previous agreement. On the other hand, when agents evade, as they tend to offer the maximum possible reward, they formulate agreements that are difficult to break.

Given the reasons for the discrepancy in the number of conflicts, we proceed to explain the negative impact of indiscriminate argumentation. The differences in the number of conflicts allow us to explain the difference between *Evade Finally Argue* and *Always Argue*. Specifically, Figure 3 shows a lower number of conflicts arise within the society when using *Evade Finally Argue*, which, in turn, results in a lower number of delays (i.e., on average 521 delays were caused by 3545 conflicts with *Evade Finally Argue*, as opposed to 720 delays caused due to 4731 conflicts with *Always Argue*). Even though *Argue First then Evade* caused only a small number of conflicts fewer than *Always Argue* (4442 conflicts as opposed to 4731), most of them got resolved (only 508 (11.5%) delays occurred with *Argue First then Evade* as compared to 720 (15.2%) delays with *Always Argue*). This leads us to conclude that the ABN in combination with evasion is a more effective strategy than indiscriminate argumentation.

Observation 4: *Using argumentation as the last resort tends to produce a higher overall performance.*

Figures 2(a) and 2(b) show a small difference in penalty between strategies *Evade Finally Argue* and *Argue First then Evade* (as per Table 2, £15,800.8 versus £14,873.9). However, Figure 1(b) shows the difference between the number of messages used to achieve this outcome as significantly higher between *Evade Finally Argue* and *Argue First then Evade* (i.e., the difference is of a magnitude of 1.35 times; 28,500.3 message units for *Evade Finally Argue* versus 38,607.7 for *Argue First then Evade*). The reason for this large difference is that when




Fig. 3. Conflict variation over different strategies in complete resource setting.

the agents use *Argue_First_then_Evade*, they always argue with the first agent. In some instances, this argumentation may not yield any agreement. However, since it has already argued with that agent, its message count has already increased. On the other hand, when using *Evade_Finally_Argue*, the agent will attempt to argue only if it gets to the very last encounter. Thus, in many cases, it resolves the conflict before it gets to the last agent. Another observation worth noting is the differences in the number of messages used by *Always_Evade* and *Evade_Finally_Argue*. The former uses more messages than the latter (Table 2 shows that *Always_Evade* use an average 33,836.8 messages, as opposed to *Evade_Finally_Argue* which uses only 28,500.3). Therefore, this shows that selective argumentation not only improves the effectiveness, but also efficiency of the system. Thus, when both efficiency and effectiveness are taken together we can conclude that evading first and arguing as the last resort tends to be the most preferable option among these strategies.

Observation 5: Exchange of meta-information such as reasons and alternatives allow agents to resolve their conflicts more efficiently than using a mere negotiation approach.

Finally, we observe the impact of exchanging meta-information within the negotiation process. To this end, Figure 4 presents the total number of messages used by the society in the complete resource setting (RS25), both when negotiation involves the exchange of meta-information and when it does not. When negotiating without exchanging meta-information, the seller agents do not incorporate reasons and alternatives when they respond to proposals, whereas when they do incorporate them, they argue in the way we have described throughout this paper (refer to Section 3.2). In Figure 4, it is clearly observable that incorporating meta-information into the interaction process allows the agents to reduce the number of messages used to resolve their conflicts. This is most apparent in the *Argue-1* and *Always_Argue* strategies, which predominantly use the argue method to resolve conflicts. The improvement is also present to a lesser degree in *Argue_First_then_Evade*, which gives priority to argue, but only marginally present in *Evade_Finally_Argue* strategy that argues only in the last encounter. The reason for this reduction is due to buyer agents using the additional information provided by the sellers in their proposal selection and partner selection techniques. Specifically, as explained in Section 3.2, when agents receive reasons, either as an *unavailable* message or as *recommended prices*, they, in turn, use this information to decide on their next proposal. On the other hand, alternative suggested by the sellers are used to select the order of contacting potential partners in future interactions. Both of these uses help to reduce the number of un-




Fig. 4. Total Messages - Complete Resource Setting: both with and without Meta-Information.

necessary proposals exchanged within the society so they increase the efficiency of the argue method. This finding is consistent with the experimental results observed by Jung et al. [7], which presents the positive contribution of incorporating meta-information on the negotiation effort. The ability to consistently replicate their observations within our domain, adds further support to our formulated argumentation context.

5 Conclusions and Future Work

ABN has been proposed as a promising means for agents to resolve conflicts in multi-agent systems. However, in many cases, not all conflicts need to be resolved; some can be overcome through evasion or re-planning. In such a context, it is important for the agents to identify the specific situations where arguing is beneficial and those in which it is not. To this end, this paper presented a preliminary empirical evaluation and assessed the efficiency and effectiveness of argumentation as a conflict resolution mechanism with respect to evasion in our particular domain.

Our results can be summarised as three main points. *First*, the relative variation of effectiveness of the methods is very much related to the resources available in the system. The chosen ABN presents a far more effective method of conflict resolution than evasion when the resources are more constrained. However, this effect tends to diminish as resources become more abundant. Furthermore, it is shown that attempting to always argue in high resource settings yields an inferior outcome (both in terms of efficiency and effectiveness) than always using evasion. *Second*, we show that selective use of argumentation is a far more effective and efficient strategy than indiscriminate argumentation. *Finally*, we show the strategy of evading first and arguing as the last resort tends to yield the most favourable overall performance among these strategies. However, this final point needs further investigation to see whether this is an artefact of our domain or is something that is more generally true. Obviously all these results are couched in the context of our particular domain and further investigation is needed to see whether they generalise.

In addition to the generalisation aspect, there are a number of different ways in which the experiments themselves can be extended. To date, our agents only use a handful of simple arguments (reasons and alternatives). It would be interesting to observe the overall effect of incorporating more persuasive forms of locutions such as appeals, threats and promises [5].

Second, in our experiments we maintained the level of commitment for all agreements at zero. This allowed agents to renege without suffering a loss. As a next step we plan to implement the concept of charging a decommitment penalty [16] and observe its impact on the performance of the strategies. Third, in the current implementation, the society has no structure and all agents operate within a peer-to-peer environment. In future developments, we plan to incorporate a social structure governed by roles and relationships [17] within agents and observe its impact on the relative effectiveness of the strategies.

6 Acknowledgements

This research is funded by EPSRC under the Information Exchange project (GR/S03706/01). We thank Chris Reed, Tim Norman, Sarvapali Ramchurn and Partha Dutta for their thoughts, contributions and discussions. We specially thank the two anonymous reviewers for their valuable comments and suggestions, and also express our gratitude to AOS Ltd. for their JACK agent framework and support.

References

1. Tessier, C., Chaudron, L., Mller, H.J., eds.: Agent's Conflicts: New Issues. In: *Conflicting Agents Conflict Management in Multi-Agent Systems*. Kluwer Academic Publishers, Dordrecht, Netherlands (2000) 1–30
2. Castelfranchi, C.: Conflict Ontology. In: *Computational Conflicts, Conflict Modeling for Distributed Intelligent Systems*. Springer (2000) 21–40
3. Walton, D.N., Krabbe, E.C.: Dialogues: Types, Goals, and Shifts. In: *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State Univ of New York Press, Albany, NY, USA (1995) 65–117
4. Rahwan, I., Ramchurn, S.D., Jennings, N.R., McBurney, P., Parsons, S., Sonenberg, L.: Argumentation-based negotiation. *The Knowledge Engineering Review* (2004) to appear.
5. Kraus, S., Sycara, K., Evenchik, A.: Reaching agreements through argumentation. *Artificial Intelligence* **104** (1998) 1–69
6. Sycara, K.: Persuasive argumentation in negotiation. *Theory and Decision* **28** (1990) 203–242
7. Jung, H., Tambe, M., Kulkarni, S.: Argumentation as distributed constraint satisfaction: Applications and results. In: *Proceedings of the Fifth International Conference on Autonomous Agents (Agents'01)*, Montreal, Canada, ACM Press (2001) 324–331
8. Rahwan, I., Sonenberg, L., Dignum, F.: Towards interest-based negotiation. In Rosenschein, J.S., Sandholm, T., Wooldridge, M., Yokoo, M., eds.: *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'03)*, Melbourne, Australia (2003) 773–780
9. Reed, C., Long, D., Fox, M., Garagnani, M.: Persuasion as a form of inter-agent negotiation. In: *Selected Papers of the 2nd Australian Workshop on Distributed Artificial Intelligence (DAI'96)*, Cairns, Australia, Springer-Verlag (1996) 120–136
10. Ramchurn, S.D., Jennings, N.R., Sierra, C.: Persuasive negotiation for autonomous agents: A rhetorical approach. In: *IJCAI Workshop on Computational Models of Natural Argument*, Acapulco, Mexico (2003) 9–18
11. Amgoud, L., Maudet, N.: Strategical considerations for argumentative agents (preliminary report). In: *Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR'02): Special session on Argument, dialogue, decision*, Toulouse, France (2002) 399–407
12. Parsons, S., Sierra, C., Jennings, N.R.: Agents that reason and negotiate by arguing. *Journal of Logic and Computation* **8** (1998) 261–292

13. Sierra, C., Jennings, N.R., Noriega, P., Parsons, S.: A framework for argumentation-based negotiation. In Singh, M.P., Rao, A., Wooldridge, M.J., eds.: Proceedings of Fourth International Workshop on Agent Theories Architectures and Languages (ATAL'97). Volume 1365 of Lecture Notes in Computer Science., Springer-Verlag (1998) 177–192
14. McBurney, P., van Eijk, R., Parsons, S., Amgoud, L.: A dialogue-game protocol for agent purchase negotiations. *Autonomous Agents and Multi-Agent Systems* **7** (2003) 235–273
15. Yokoo, M., Hirayama, K.: Distributed constraint satisfaction algorithm for complex local problems. In: Proceedings of the Third International Conference on Multiagent Systems (ICMAS'98), Paris, France (1998) 372–379
16. Sandholm, T.W., Lesser, V.R.: Advantages of a leveled commitment contracting protocol. In: Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI'96), Portland, OR, USA (1996) 126–133
17. Panzarasa, P., Jennings, N.R., Norman, T.: Social mental shaping: Modelling the impact of sociality on the mental states of autonomous agents. *Computational Intelligence* **17** (2001) 738–782