ENHANCING COOPERATION THROUGH INTERACTION AND COMMUNICATION IN AGENT-BASED JOINT ACTIVITY-TRAVEL SCHEDULING

Huiye Ma, Urban Planning Group, Eindhoven University *of Technology, h.ma@tue.nl*
Theo Arentze, Urban Planning Group, Eindhoven University of Technology, t.arentze@tue.nl
Harry Timmermans, Urban Planning Group, Eindhoven University of Technology, h.timmermans@tue.nl

1. ABSTRACT

This paper considers making decisions on whether to cooperate with other participants in agent based joint activity scheduling. Several characteristics of the scheduling process must be addressed: participants may not have complete information about the schedules of the others involved; the effectiveness of interaction/communication may not be known a priori; and interaction/communication has some associated cost. The paper adopts a novel probabilistic representation of other agents' beliefs about the actions selected for their own or for the joint activity, given incomplete information. Agents can use this representation to make four kinds of decisions: communicating information relevant to other group member(s), asking for information from other member(s), undertaking actions to cooperate with others, and inquiring actions from others. Several decision-theoretic mechanisms are presented, which include a set of rules for reasoning about the utility and credit of actions and communications, and the cost incurred. It was tested using a multi-agent system on activity travel scheduling with configurations that varied agents' uncertainty about the world and the cost of cooperation. In all cases, agents using the decision-theoretic mechanisms to decide whether to cooperate outperformed agents using other mechanisms in terms of utility.

Keywords: activity travel scheduling, cooperation, multi-agent simulation, decision making.

1. INTRODUCTION

The majority of activity-travel scheduling systems have been designed to generate schedules which will take effect for a certain period of time, *e.g.*, a single day. Among the various activities scheduled, some may require more than one participant, implying that the schedules of the group members (i.e. a family or community) should be synchronized and coordinated. Therefore the involved participants need to communicate and interact so as to schedule the joint activity. Each participant's schedule will be consequently influenced by other(s)'s information and interaction and be rescheduled accordingly.

Cooperation has been widely observed in the activity-travel scheduling area because of joint activities where participants work together toward a shared goal. The participants in a cooperative joint activity form coordinated and mutually supportive schedules (Kamar et al., 2009). They make commitments to the joint activity and perform some constituent tasks of that activity.

The information exchange and the interactive behaviour during the cooperation process trigger dynamics in a variety of ways. For example, person *i* needs to communicate with person *j* to find out the time to meet each other. The information exchange will lead them to adjust their original schedules slightly so that the joint activity will happen at the same time. Another example is about interactions between persons. If person *i* wants to buy some tulips today while person *j* happens to know this and needs to buy some food today in the shopping centre, person *j* may buy the tulips for person *i* since it is very convenient for person *j* to help. This behaviour will bring changes for person *i*'s schedule as she does not need to buy the flowers.

To date, dynamics have been barely incorporated in activity-travel models. Most models are basically concerned with the simulation of daily activity patterns derived from one or two day observed activity-travel diaries and, if they do assume longer time frames, deal with dynamics on the level of activity generation (Arentze and Timmermans, 2007, 2009). More recently, short-term dynamics have drawn most attention, *e.g.*, rescheduling of activities to respond to unexpected events when implementing a planned activity-travel agenda. Nevertheless, dynamics requires more research effort.

To handle dynamics existed in joint activity scheduling in a more practical way, multi-agent systems are an excellent approach because participants can be better modelled and human societies can be better reproduced. Multi-agent-based simulation has been successfully applied in various research fields (Jacyno et al., 2009; Vasirani et al., 2009; Rieser et al., 2009). As in most multi-agent task settings, the joint activity is carried out in a world that is constantly changing; the participants' knowledge about the world is inherently incomplete; individuals have (sensory) access to different parts of the world; and their beliefs about how best to perform an action may differ. All these practical concerns can be readily handled in multi-agent systems.

Although the participants in joint activities have an incentive to cooperate with others by the nature of their commitments to the shared goal and to each others' actions in satisfying that goal, a decision about whether or not to cooperate still requires deliberation. Cooperative actions result in some cost to the agent helping. Usually costs include resources consumed in communicating, lost opportunities to do other activities, and the need for group members to adapt their individual schedules to the action or its effects. On the other hand, arriving at a joint decision in collaboration may also mean that someone has to give in to the other's preference. If the participants interact and communicate continuously in turn, the agent who has given in and lost some utility for cooperating will be paid back by another agent who has received cooperation in previous turns. This has been well observed in real life and motivated the introduction of credit in the literature (Ettema et al., 2007). Thus, in cooperative settings, agents must weigh the trade-off among the potential benefit, its associated costs in the current turn, and the possible payoff in the future.

So far, not much has been said about how interactions and communications affect activity choice in the area of activity-travel scheduling. Therefore, in this piece of work we focus on the decision making of whether to communicate/interact by assuming that:

- Agents are rational, cooperative, and have incomplete information;
- Any communication or interaction has some cost;
- An agent will gain/lose credit when she offers/receives cooperation to/from others through communication or interaction.

This paper addresses the intertwined problems of recognizing when communication / interaction is needed in cooperation and determining whether to communicate / interact, taking into account the cost of it, its possible effect on the beliefs and commitments of group members, the relations between the gain/loss of utility and credit. It is specifically concerned with joint activities that take place in settings in which there is uncertainty about agents' capabilities and about the state of the world. The representation to incorporate agent interaction and communication into activity travel scheduling model is adopted and improved. We propose several mechanisms to enable agents to make decisions on whether or not to start (receive) communication and interaction given the cost, the gain or loss of utility and credit. A multi-agent system is designed to carry out four sets of experiments by considering consecutive scenarios. The experimental results tell that the mechanisms improve agents' performance in terms of perceived utility with various costs considered.

This paper is organized as follows. In section 2, related work is given. In section 3, we present basic definitions of the central concepts for our model. Section 4 describes the communication model and the interaction model. Section 5 shows the multi-agent simulation results. Conclusions and future work are discussed in section 6.

2. RELATED WORK

Several papers in the literature have studied on how to model collaboration and teamwork, all of which recognize communication / interaction as a major requirement for successful and cohesive collaborative activities.

Some prior approaches have axiomatized decisions to communicate or to help in the context of formalizations of collaboration in terms of the intentions, beliefs and mutual beliefs of the participants. Cohen and Levesque's axiomatic approach stipulates that agents communicate to the group whenever a goal is discovered impossible to achieve (Cohen et al., 1997; Levesque et al., 1990). Fan *et al.* extend the set of logical axioms to provide for proactive information exchange (Fan et al., 2005). These approaches neither consider the cost of communication nor provide mechanisms for helpful actions.

The SharedPlan (SP) formalization includes axioms that entail adopting intentions for helpful acts, or lead to communication based on certain kinds of intentions in the SP specification (Grosz et al., 1996). These axioms represent both the benefit of a helpful action to the group activity and the costs to the individual performing the helpful action. However, they do not handle uncertainty regarding the world or agents' capabilities.

STEAM, which drew on both the joint intensions and the SharedPlans theories above, supported the construction of agents able to collaborate in complex, real world domains of military training and robot soccer (Tambe et al., 1997). It included a decision-theoretic mechanism for communication which modelled the cost-benefit trade-off associated with communicating information to the full group. This mechanism constructed a decision-tree for each agent every time a communication action was considered. This had significant complexity costs for agents that needed to consider many such actions.

Works on decentralized approaches to multi-agent planning have provided models that consider the cost-benefit tradeoffs of communication among agents (Goldman et al., 2004). As helpful behaviour can emerge between any agents in the collaborative activity, the helpful behaviour needs to be directly embedded in the joint policy of the whole group of agents, making it exponential in the size of the history of agents' observations. Refining agents' plans in this setting means updating their entire policy every time a helpful action is considered, which is impractical in their application.

Most recently, Kamar et al. (2009) considered communication as a special form of helpful behaviour and then examining the general question of when to help. In addition, it provides a decision-theoretic rather than axiomatic approach to model helpful behaviour decisions. The Colored Trails (CT) system has been applied as a testbed in the evaluation. However, they ignore the difference between that agent who starts the communication and the agent(s) who receive the information and the difference of communication and interaction; the motivation to help has not been handled; selfishness/altruism has not been taken into account.

Interactions for social purposes have been discussed in our previous work (Ettema et al., 2007). The impact of social interaction on decision-making in daily life has been explored. Their work covers short-term activity-travel and more longer-term decisions such as relocation. Moreover, the framework outlined describes how beliefs, preferences and behaviours are adjusted as a result of social interaction. Nevertheless, the reason for people to start interaction is missing and the effect of communication and interaction is not discussed in any detailed level.

3. PRELIMINARY

Throughout the paper, we will illustrate various aspects of this decision-making problem using an example of two agents, G_i and G_j in a community. Each agent needs to make the decision of communicating or interacting with another agent in order to reach a joint re-scheduling given a dynamic activity travel environment.

Assume agent G_i has a fulltime job. Her original schedule for the day is home-work-home. Suppose agent G_j has a part-time job and is free today to decide whether to go shopping or not. At a certain moment, agent G_j is in the mood for joint shopping with G_i. That means that if she can find someone to go together, then she will go. Otherwise, she will not go. Hence agent G_j inquires agent G_i about the possibility of going shopping together.

After agent G_i receives the message from agent G_j, she replies to agent G_j that she is willing to join. Later at the end of agent G_i's work, she realizes that she might be delayed for half an hour. So she tells agent G_j that she may be delayed for half an hour today because of extra work and she would like to pick her up for compensation of the time.

Based on this example, we can see that time is an important issue during communication and interaction. Therefore, we first give definitions on the milestones during the decision making process and then define agent communication and interaction.

3.1 Definitions

A decision point consists of only one communication or interaction for an agent at which point the agent has to make the decision to communicate/interact or not. At each decision point for an agent, there is a corresponding decision tree for the agent to make her decision.

A decision turn consists of at least one communication and/or interaction among agents so that the joint activity of the agents can be fulfilled. A decision turn includes at least one decision point.

Agent communication consists of message passing only.

Agent interaction consists of actions and possible messages which assist action and serve the same goal as the actions.

For each communication or interaction, there is one **initiator**, denoted as G_s, and at least one **respondent**, denoted as G_r.

For example, agent G_j sends a message to agent G_i to ask whether she wants to go shopping together at 5:00 pm today. Agent G_j is the initiator and the message passing will cost her some money. Agent G_i is the respondent and her information set will be updated by the message. This message passing belongs to agent communication. On the other hand, agent interaction refers to the actions where messages might be related but are not the main purpose. When agent G_i picks up agent G_j at home, the pick up action may require a phone call or message passing as well. However, the communication is assisting in finishing the pick up action and may be ignored in the mechanism for simplicity.

3.2 Credit

Credit refers to an existing balance in doing each other favours. Credit is an indicator of how much an agent gives in for doing the favour. Credit can be computed in various ways(Ettema *et al.*, 2007). In this paper, we let the agent being helped evaluate the favour and decide on the amount of credit. Each time agent G_i helps agent G_j, the credit that agent G_i has regarding other agents (including agent G_j) will be increased. This amount of increase is determined by agent G_j and scaled according to agent G_j's total utility. Reversely, if agent G_i receives help from agent G_j, the credit of agent G_i regarding other agents will be decreased because that amount of credit decrease will be transferred and added to agent G_j.

3.3 Utility function

Utility is the subjective value an agent assigns to a choice alternative and tries to maximize when making a choice. Every possible action or communication is evaluated in terms of utility consequences by each agent individually. In general, agent G_i, will be happy if her utility is increased, or the total utility of the group for the joint activity is increased.

Given the joint activity each agent aims to achieve and the assumption of cooperative and truthful behaviours, we postulate that the utility a given agent derives from a set of activities consists of the following components where G_i is agent *i*, GR is the group of agents including agent *i*, t_k^l is the current time assuming the current decision turn is *k* and the current decision point is *l*, and finally $w_1 + w_2 = 1.0$:

$$U(G_i, GR, t_k^l) = w_1 \times U(G_i, t_k^l) + w_2 \times U(GR, t_k^l)$$

The value of $U(GR, t_k^l)$ is computed by:

$$U(GR, t_k^l) = SUM(U(G_j, t_k^l))$$

where G_j represents any agent $j \in GR$; and SUM denotes the total utility of the group computed after traversing the tree taking into account all the participants.

If one puts a larger value on w_1 than on w_2, that implies the agent would like to emphasize her own utility and credit more than the group utility. In the meanwhile, one can deduce that the agent is more selfish than altruistic if she assigns a larger value to w_1 than w_2. The selfishness feature is also reflected in the situation when the agent needs to decide how much percentage of the credits should be returned to another agent, which will be shown in the algorithms in the later section. If the value of w_1 is close to 0, it means that agent G_i is very selfish and will accept all the credits transferred by other agent(s). On the other hand, if it is approaching 1.0, it means the agent is extremely altruistic by asking no credit from other(s) being helped by her.

3.4 The Probabilistic Decision Tree

Key features of the joint activity formalism are that agents' plans may be partial, agents may have incomplete information about the way to accomplish a joint activity, and agents are responsible for different constituent actions. For example, agent G_i may not know the shopping centres that agent G_j prefers, but they still have a joint commitment to go shopping together. However, agents cannot reason about the benefit to the group from engaging in a helpful action when they have no information about the actions that other group members are considering.

To bridge this gap, we adopt a Probabilistic Decision Tree (PDT) which is inspired by and follows part of the Probabilistic Recipe Trees proposed by (Kamar et al., 2009). PDT enables agents to represent their beliefs about the actions that may be selected by group members to complete a joint activity. As an example, we draw a picture of the PDT at certain moment for agent G_i in Figure 1.

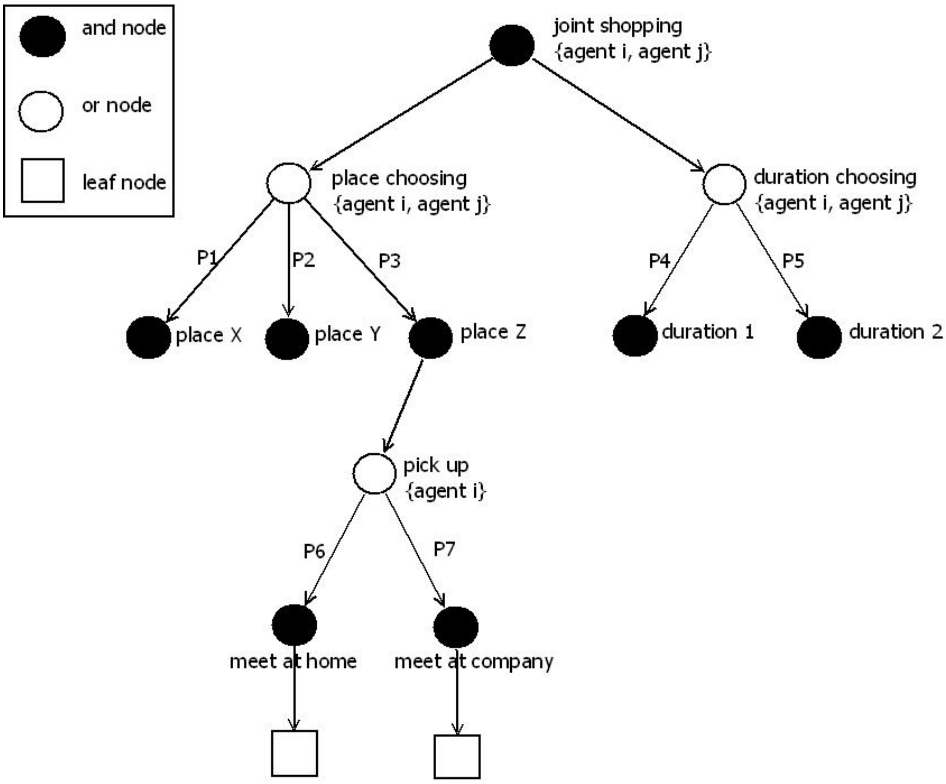

Figure 1. The probabilistic decision tree used in the reasoning.

Assume α represent the joint activity (going shopping together) to be achieved by agent G_i and agent G_j in the figure; β represent the sub-action to be taken by agent G_i or G_j with certain probability. A PDT for α, denoted as PDT_α, is a structured tree representation that defines a complete probability distribution over the possible sub-actions for accomplishing α. Each node in a PDT represents an action and has several properties associated with the action (e.g., the set of agents responsible for carrying out the action). Leaf nodes represent basic-level actions, and intermediate nodes represent complex actions. Intermediate nodes may be either AND or OR nodes. Each child of an AND node represents a constituent sub-action for completing the AND node action. Each child of an OR node represents a possible choice of a sub-action for the OR node action, where the choice is non-deterministic. Each branch from an OR node to one of its children nodes has an associated probability representing the likelihood that the child node is selected as a sub-action for the OR node action.

Figure 1 presents a PDT for going shopping together consisting of agent G_i and agent G_j in agent G_i's point of view at certain moment. The children of the top node of the tree (joint shopping) are agents' place choosing and agents' duration choosing. The likelihood of choosing place X for both agents, P1, in this example is 0.3.

We list all these operators and the relevant notations into Table 1. Among them, three operators have been defined by Kamar et al. (2009) and can be used to restructure a PDT as agents refine their actions with communication and interaction.

Type	Notation	Meaning
Agents	Agent G_i Agent G_j	Agent(s) reasoning about interaction and communication. Partner(s) of agent G_i.
Actions	α β γ	Top level joint activity. Sub-actions included in PDT. The interaction that agent G_i tries to carry out.
Communications	ω	The communication that agent G_i tries to carry out.
PDT	$PDT_\alpha \in PDT$ $PDT_\beta \in PDT$ $PDT_\gamma \in PDT$ $PDT_\omega \in PDT$	PDT selected for joint activity α. PDT selected for sub-action β. PDT selected for action γ. PDT selected for communication ω.
Operators	$PDT_\alpha \cup PDT_\beta$ $PDT_\alpha \backslash PDT_\beta$ $PDT_\alpha \otimes PDT'_\beta$	Add PDT_β as a child of PDT_α. Remove the sub-tree PDT_β from PDT_α. Replace PDT_β by PDT'_β for PDT_α.

Table 1. The list of notations and operators.

The decision-theoretic analysis of interactions and communications requires computing the costs and credits of performing an action or communication based on the selected actions. To this end, we utilize the following set of functions to represent agents' beliefs for the sub-actions selected for an action, and to evaluate the costs, credits, and utilities of these sub-actions.

The function *P*-$CBA(PDT_\alpha, C_\alpha)$ represents the probability of successfully performing action α in context C_α given the sub-actions represented in PDT_α (Kamar et al., 2009). For leaf nodes representing basic actions, the function returns a value that equals the probability that an agent can bring about that basic action in certain context. For AND nodes the function returns the product of the probabilities that the children nodes will succeed. For OR nodes, the function returns an average of the likelihood that the child nodes will succeed, weighted by the probability assigned to the child.

The $(E)Cost(G_i, PDT_\alpha, C_\alpha, t_k^l)$ function represents the (expected) cost to agent G_i if the agents in the group are going to carry out the sub-actions represented in PDT_α in context C_α. For leaf nodes, this function returns the value of the cost incurred by an agent applied to the leaf. For AND nodes this function is a summation of the cost of its children nodes. For OR nodes it is an average of the costs for the children nodes, weighted by the probability assigned to each child.

The $CreditChange(G_i, PDT_\alpha, C_\alpha, t_k^l)$ function represents the credit that agent G_i will transfer to other agent, *e.g.*, agent G_j, because G_i has received help from agent G_j, given that the group is going to carry out the actions represented in PDT_α in context C_α. The value of credit in the current decision turn may be positive, negative, or zero.

$ECreditChange(G_j, PDT_\alpha, C_\alpha, t_k^l)$ function represents the credit that agent G_j expects to receive from agent G_i before G_j goes to help G_i. Please note that the values of these two credits may be different since they reflect the different measurement of the benevolence in two individuals' point of view. The expected credit for agent G_j, $ECreditChange(G_j, PDT_\alpha, C_\alpha, t_k^l)$, will be computed according to G_j's current total utility, which should be a non-decreasing function for $U(G_j, GR, t_k^l)$.

$RCredit(G_j, PDT_\alpha, C_\alpha, t_k^l)$ is the credit that G_j will return to G_i who has transferred the credit to G_j. Taking into account the selfishness or altruism feature of G_j, the credit for agent G_j to return to G_i is equal to $w_2 \times CreditChange(G_i, PDT_\alpha, C_\alpha, t_k^l)$. The part of credit for agent G_j to keep for herself at time t_k^l is computed by: $w_1 \times CreditChange(G_i, PDT_\alpha, C_\alpha, t_k^l)$.

The function $Eval(G_i, GR, PDT_\alpha, C_\alpha, t_k^l)$ represents the difference between the expected utility to the group from agent G_i's point of view and the related cost for carrying out the sub-actions represented in PDT_α in context C_α. It combines the expected utility of the parent node with the expected utilities of its children, denoted as SUM. The expected utility of a node is the value of the action that the node represents multiplied by the success probability (*P*-CBA) of the node. If the node is an OR node, the expected utility of each child node is weighted by its branching probability. In order to compute the result of $SUM(EU(G_i, PDT_\alpha, C_\alpha, t_k^l))$, one needs to traverse the current tree.

$$Eval(G_i, GR, PDT_\alpha, C_\alpha, t_k^l) = SUM\left(EU(G_i, PDT_\alpha, C_\alpha, t_k^l)\right) + ECreditChange(G_j, PDT_\alpha, C_\alpha, t_k^l) + Cost(G_i, PDT_\alpha, C_\alpha, t_k^l)$$

The *Select*-$PDT(G_i, \alpha, C_\alpha, t_k^l)$ refers to the PDT that represents G_i's belief about the possible sub-actions she will select to perform action α in context C_α. The *Predit*-$PDT(G_i, G_j, \alpha, C_\alpha, t_k^l)$

refers to the PDT that represents G_i's belief about the possible actions G_j will select to perform action α in context C_α (Kamar et al., 2009).

4. INTERACTION AND COMMUNICATION MODELS IN ACTIVITY TRAVEL SCHEDULING

4.1 Communication Model

The ability to communicate information allows agents to convey changes in the world or to request information about the world. Assume that an agent is willing to cooperate with others in scheduling the joint activity. When the agent makes the decision to start communication, she will first check whether it is worthwhile to start. We present two pairs of algorithms for handling different situations of starting communication. The first pair handles how to decide conveying information and receiving information. The second pair takes care of the decision on asking for information and offering information.

4.1.1 Conveying Information

In the situations where two agents G_i and G_j are committed to the success of a joint activity, when G_i has a piece of new information, ω, she needs to reason about informing G_j about this information. The decision to communicate may improve the utility of the group from her point of view and may increase her credit as well, but any communication is associated with a cost for herself. Algorithm in Figure 2 specifies the process by which G_i reasons about this trade-off. In particular, G_i reasons about the actions that G_j would adopt for doing sub-actions of α if G_i has communicated information ω. If the utility and credit gain to the group and herself is higher than the cost of communication for her, then agent G_i will start to communicate ω to G_j.

As an example, if agent G_i sees that she will be delayed by work for 30 minutes, and she thinks that agent G_j will not be patient to wait without being notified beforehand, she would conclude that their joint shopping is likely to fail. If G_i informs agent G_j about this information, agent G_j can update her actions so that she can arrange the time to do other things. If agent G_i forecasts that the utility and credit improvement generated by the communication is higher than communication cost, agent G_i will inform agent G_j. However, if agent G_i believes that agent G_j is not likely to mind waiting or the communication cost is very high, then she would not inform agent G_j.

Assume the current decision turn is k and the current decision point is l;
Assume each agent has her information set and her own preference;
Let G_s/G_r be the start/receiver of the communication ω;
Let w_1^i be the variable of selfishness-altruism for agent i;
Let β be the action(s) that intends to do;
PDT_α:=*Predit-PDT*$(G_s, GR, \alpha, C_{GR}, t_k^l)$;
C_β^ω:=*Context-Update*(C_β, ω, t_k^l);
PDT_β^ω:=*Predit-PDT*$(G_s, G_r, \beta, C_\beta^\omega, t_k^l)$;
$PDT_\alpha^\omega := PDT_\alpha \otimes PDT_\beta^\omega$;
if $Eval(G_s, GR, PDT_\alpha^\omega, C_\alpha^\omega, t_k^l) \geq Eval(G_s, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
    G_s conveys ω to G_r;
    $PDT_\alpha := PDT_\alpha^\omega$;
    if G_s receives a credit change, $CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$, from G_r **then**
        $Credit(G_s, PDT_\alpha, C_\alpha, t_k^l) = w_1^s \times CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$;
        Send the credit $w_2^s \times CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$ to G_r;
    end
end

Figure 2. The pseudo-code of the algorithm for conveying communication model by G_s.

4.1.2 Receiving Information

After agent G_i conveys the information to G_j, G_j will update her information set and carry out her actions in the current context. As a consequence, her utility will be increased and she will transfer certain amount of credit to G_i to acknowledge the help. The pseudo-code of the algorithm for receiving information communication model by G_j is shown in Figure 3.

Receive the information ω from G_s;
PDT_α:=*Predict-PDT*$(G_r, GR, \alpha, C_{GR}, t_k^l)$;
C_β^ω:=*Context-Update*(C_β, ω, t_k^l);
PDT_β^ω:=*Select-PDT*$(G_r, \beta, C_\beta^\omega, t_k^l)$;
$PDT_\alpha^\omega := PDT_\alpha \otimes PDT_\beta^\omega$;
$PDT_\alpha := PDT_\alpha^\omega$;
$CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l) = f(U(G_r, PDT_\alpha, C_\alpha, t_k^l))$;
G_r transfers the credit change to G_s;
if G_r receives a credit, $RCredit(G_s, PDT_\alpha, C_\alpha, t_k^l)$, from G_s **then**

$Credit(G_r, PDT_\alpha, C_\alpha, t_k^l) = RCredit(G_s, PDT_\alpha, C_\alpha, t_k^l) - CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$;
end

Figure 3. The pseudo-code of the algorithm for receiving communication model by G_r.

4.1.3 Asking for Information

Agent G_i may need to reason about asking G_j for information ω if this information is beneficial for her and the whole group. To compute the utility, G_i needs to consider how she will adapt her own belief about the actions for the possible answer from G_j. Figure 4 gives the procedure of the communication.

For example, if agent G_i is not patient about waiting with nothing to do and believes that the joint activity will fail as a result of that, she can ask agent G_j if agent G_j will arrive to meet her in time. For each possible answer agent G_i may receive from agent G_j, she updates her belief about actions to select that incorporates that answer. After weighting each possible updated actions with the probability of receiving that answer, agent G_i computes the expected utility for asking and the credit loss. If it is higher than the communication cost, agent G_i considers asking. However, if agent G_i believes that the answer will not improve the actions she selects, or the cost of communication is very high, then she would not consider communicating with agent G_j.

Let $pr(Bel(G_r, \omega))$ be G_s's prediction of the probability of receiving ω from G_r;
PDT_α:=Predict-$PDT(G_s, GR, \alpha, C_{GR}, t_k^l)$;
C_β^ω:=Context-$Update(C_\beta, \omega, t_k^l)$;
PDT_β^ω:=Select-$PDT(G_s, \beta, C_\beta^\omega, t_k^l)$;
$PDT_\alpha^\omega := PDT_\alpha \otimes PDT_\beta^\omega$;
if $pr(Bel(G_r, \omega)) \times Eval(G_s, GR, PDT_\alpha^\omega, C_\alpha^\omega, t_k^l) \geq Eval(G_s, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
    G_s asks for ω from G_r;
    $PDT_\alpha := PDT_\alpha^\omega$;
    $CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l) = f(U(G_s, PDT_\alpha, C_\alpha, t_k^l))$;
    G_s transfers the credit change to G_r;
    if G_s receives a credit, $RCredit(G_r, PDT_\alpha, C_\alpha, t_k^l)$, from G_r **then**

$Credit(G_s, PDT_\alpha, C_\alpha, t_k^l) = RCredit(G_r, PDT_\alpha, C_\alpha, t_k^l) - CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$;
    end
end

Figure 4. The pseudo-code of the algorithm for asking for communication model by G_s.

4.1.4 Offering Information

After agent G_i asks for the information from G_j, G_j will evaluate whether it is worthwhile to answer in the current context. If she sends her answer to G_i, she may get some credit as a reward to her help. Moreover her utility may be increased as well. She may transfer certain amount of credit back to G_i to show her benevolence. The pseudo-code of the algorithm for offering information communication model is given in Figure 5.

PDT_α:=*Predict-PDT*$(G_r, GR, \alpha, C_{GR}, t_k^l)$;
C_β^ω:=*Context-Update*(C_β, ω, t_k^l);
$PDT_\alpha^\omega := PDT_\alpha \otimes PDT_\beta^\omega$;
if $Eval(G_r, GR, PDT_\alpha^\omega, C_\alpha^\omega, t_k^l) < Eval(G_r, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
  return;
end
Offer the information ω to G_s;
$PDT_\alpha := PDT_\alpha^\omega$;
if G_r receives a credit change, $CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$, from G_s **then**
  $Credit(G_r, PDT_\alpha, C_\alpha, t_k^l) = w_1^r \times CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$;
  Send the credit $w_2^r \times CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$ to G_s;
end

Figure 5. The pseudo-code of the algorithm for offering communication model by G_r.

4.2 Interaction Model

According to the way of differentiating various communication models, the interaction models are presented in the same ways. The first pair is to handle the case when one agent cooperates with another agent and when another agent receives the action in Figure 6 and Figure 7. The other pair is to ask for cooperation and how to answer the request in Figure 8 and Figure 9.

PDT_α:=*Predict-PDT*$(G_s, GR, \alpha, C_{GR}, t_k^l)$;
PDT_γ:=*Select-PDT*$(G_s, \gamma, C_\gamma, t_k^l)$;
$PDT_\alpha^\gamma := PDT_\alpha \otimes PDT_\gamma$;
if $Eval(G_s, GR, PDT_\alpha^\gamma, C_\alpha^\gamma, t_k^l) \geq Eval(G_s, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
  G_s implements the action γ;
  $PDT_\alpha := PDT_\alpha^\gamma$;
  if G_s receives a credit change, $CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$, from G_r **then**
    $Credit(G_s, PDT_\alpha, C_\alpha, t_k^l) = w_1^s \times CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$;
    Send the credit $w_2^s \times CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$ to G_r;
  end
end

Figure 6. The pseudo-code of the algorithm for starting interaction model by G_s.

Receive the action γ from G_s;
PDT_α:=*Predict-PDT*$(G_r, GR, \alpha, C_{GR}, t_k^l)$;
C_β^γ:=*Context-Update*(C_β, γ, t_k^l);
PDT_β^γ:=*Select-PDT*$(G_r, \beta, C_\beta^\gamma, t_k^l)$;
$PDT_\alpha^\gamma := PDT_\alpha \otimes PDT_\beta^\gamma$;
$PDT_\alpha := PDT_\alpha^\gamma$;
$CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l) = f'(U(G_r, PDT_\alpha, C_\alpha, t_k^l))$;
G_r transfers the credit change to G_s;
if G_r receives a credit, $RCredit(G_s, PDT_\alpha, C_\alpha, t_k^l)$, from G_s **then**
$Credit(G_r, PDT_\alpha, C_\alpha, t_k^l) = RCredit(G_s, PDT_\alpha, C_\alpha, t_k^l) - CreditChange(G_r, PDT_\alpha, C_\alpha, t_k^l)$;
end

Figure 7. The pseudo-code of the algorithm for receiving interaction model by G_r.

Let $pr(Bel(G_r, \gamma))$ be G_s's prediction of the probability of receiving γ from G_r;
PDT_α:=*Predict-PDT*$(G_s, GR, \alpha, C_{GR}, t_k^l)$;
C_β^γ:=*Context-Update*(C_β, γ, t_k^l);
PDT_β^γ:=*Select-PDT*$(G_s, \beta, C_\beta^\gamma, t_k^l)$;
$PDT_\alpha^\gamma := PDT_\alpha \otimes PDT_\beta^\gamma$;
if $pr(Bel(G_r, \gamma)) \times Eval(G_s, GR, PDT_\alpha^\gamma, C_\alpha^\gamma, t_k^l) \geq Eval(G_s, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
    G_s inquires γ from G_r;
    $PDT_\alpha := PDT_\alpha^\gamma$;
    $CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l) = f'(U(G_s, PDT_\alpha, C_\alpha, t_k^l))$;
    G_s transfers the credit change to G_r;
    if G_s receives a credit, $RCredit(G_r, PDT_\alpha, C_\alpha, t_k^l)$, from G_r **then**
$Credit(G_s, PDT_\alpha, C_\alpha, t_k^l) = RCredit(G_r, PDT_\alpha, C_\alpha, t_k^l) - CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$;
    end
end

Figure 8. The pseudo-code of the algorithm for inquiring interaction model by G_s.

PDT_α:=*Predict-PDT*$(G_r, GR, \alpha, C_{GR}, t_k^l)$;
C_β^γ:=*Context-Update*(C_β, γ, t_k^l);
PDT_β^γ:=*Predict-PDT*$(G_r, G_s, \beta, C_\beta^\gamma, t_k^l)$;
$PDT_\alpha^\gamma := PDT_\alpha \otimes PDT_\beta^\gamma$;
if $Eval(G_r, GR, PDT_\alpha^\gamma, C_\alpha^\gamma, t_k^l) < Eval(G_r, GR, PDT_\alpha, C_\alpha, t_k^l)$ **then**
    return;
end
Offer the action γ to G_s;
$PDT_\alpha := PDT_\alpha^\gamma$;
if G_r receives a credit change, $CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$, from G_s **then**
    $Credit(G_r, PDT_\alpha, C_\alpha, t_k^l) = w_1^r \times CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$;
    Send the credit $w_2^r \times CreditChange(G_s, PDT_\alpha, C_\alpha, t_k^l)$ to G_s;
end

Figure 9. The pseudo-code of the algorithm for offering interaction model by G_r.

5. EMPIRICAL EVALUATION

In this section, we provide an experimental evaluation of the mechanisms described in Section 4. Considering practical scenarios happened in activity travel area, we design several sets of experiments to show how cooperation through communication and/or interaction will affect the utility and credit and how selfishness/altruism will influence the outcome.

The first set of experiments compares the utility and credit outcome when agents are always cooperative, randomly cooperative, or never cooperative. The second set of experiments checks the utility and credit outcome comparing never communication and communication model, on the one hand, and never interaction and interaction model, on the other. The third set of experiments explores consequences when agent's uncertainty about the world increases. We expect the communication will benefit her because other participant(s) will help her with their information. The fourth set of experiments shows how the selfishness or altruism attribute make a difference to agents.

5.1 Experimental setup for all the experiments

In the experiments, two agents *i* and *j* are considered to achieve the joint activity of going shopping. For each agent, at the beginning of a day, the optimal schedule for the day is generated according to an agenda. For example, the original optimal schedule of agent *i* for today is shown in Table 2 and for agent *j* in Table 3. Based on the optimal schedule, the utility for all the activities in the day will be calculated.

Agents will get a utility gain of 100 if the joint shopping is achieved in the end. For the rest sub actions, they will receive a utility gain of 10 for each communication transferred/received and 30 for each interaction done/received. The credit change for communication and interaction will be determined according to the total utility of the agent. For simplicity, we set it the same as the utility change, which means that an agent will receive a credit gain/loss for each information transferred/received and a credit gain/loss of 30 for each interaction done/received. These values are chosen so as to reflect the different levels of satisfaction/effort for a person to achieve a joint activity, to send a message, or to do something helpful to others in real life. They also comply with others' setting in the literature.

In the scheduling process, an agent can cooperate with her partner by responding to her message so that the partner is able to realize which path can bring her the best profit and hence to choose. Moreover, agents can ask for information actively when necessary. There is a cost associated with the message passing and agents need to trade off during their decision making. The cost will be different with our scenarios so as to allow us to check its effect on utility and credit. In the next four sets of experiments, we assume the cost of communication is 5, 10, 20, 30, and 40 while the cost of interaction is 15, 30, 60, 90, and 120.

Serial no.	Title of the activity	Duration

0	Start	9:15am
1	Care contact-2	30 minutes
2	Work	480 minutes
3	Care contact-1	60 minutes
4	Leisure green	30 minutes
5	Care personal	30 minutes
6	Leisure at home	240 minutes
7	Sleep	440 minutes

Table 2. The optimal schedule of agent *i*.

Serial no.	Title of the activity	Duration
0	Start	9:45am
1	Care contact-2	135 minutes
2	Care contact-1	90 minutes
3	Leisure green	90 minutes
4	Care personal	35 minutes
5	House work	105 minutes
6	Leisure at home	385 minutes
7	Sleep	500 minutes

Table 3. The optimal schedule of agent *j*.

At some randomly generated time, agent *j* starts to think about joint shopping and wants to communicate with agent *i* for shopping together. Then both agents will produce their probabilistic decision trees according to their beliefs and knowledge. A sample of PDT by agent *i* for joint shopping is shown in Figure 10 and for agent *j* in Figure 11. With time going on, more information will be combined through communication and interaction. Consequently, the PDT will be adjusted from time to time.

At the beginning of one day; the amount of decision points for all the agents to take in the joint activity is randomly selected from 1 to 10. For each decision point, the time to do it is randomly generated as well. At each decision point, the agent who is going to make her decision at the current point is randomly chosen; and the model for the agent to carry out is also randomly selected from 4 available models in section 4. If all the agents make decisions to start/receive the communication and interaction at all decision points, the joint activity will be achieved successfully.

After all the activities and joint shopping activity of the current day have been finished, another day comes with different optimal schedules for these two agents. The process continues until 100 days have been accomplished. All these random values require 100 days duration to give a stable outcome.

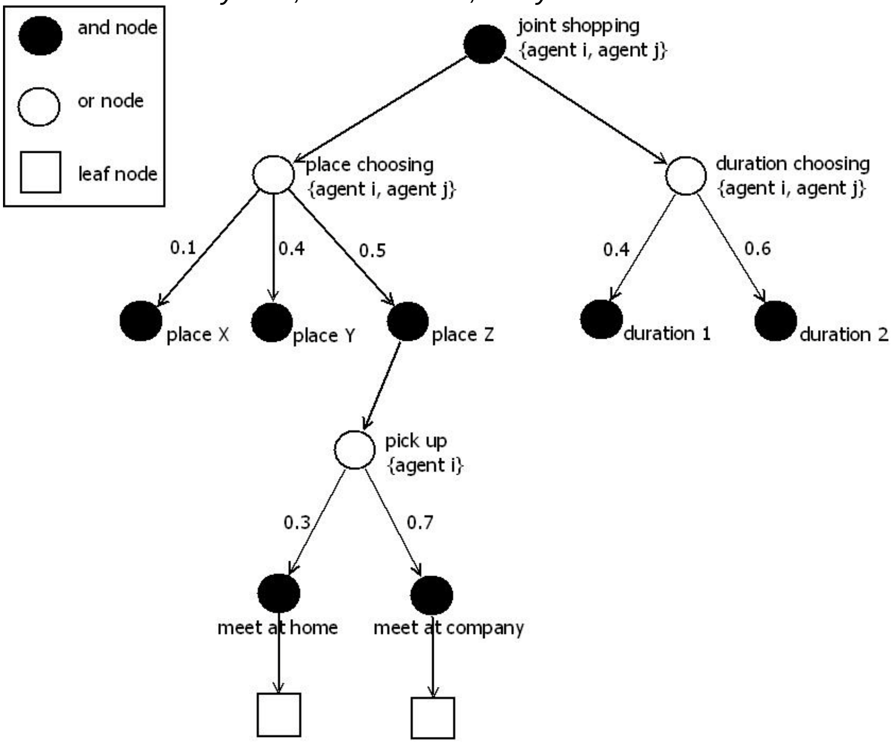

Figure 10. A sample of PDT of agent *i* at certain moment for a joint shopping.

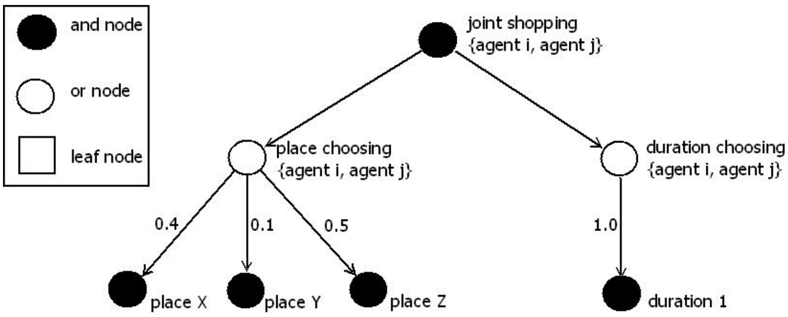

Figure 11. A sample of PDT of agent *j* at certain moment for a joint shopping.

5.2 Experimental setup and result of the first set of experiments

In the first set of experiments, the always cooperation protocol assumes agent(s) will always respond to the request on communication and interaction from her partner in a rational way. An agent will check whether it is worthwhile to cooperate according to the models provided in Section 4. In the random cooperation protocol, a random decision on responding or not is generated before an agent judges whether it is worthwhile. In the never cooperation protocol one will not respond to any request from her partner, which means no cooperation at all. In all these cases, we consider a changing cost of communication and interaction.

The utility of agent *i* when always cooperate, random cooperate, or never cooperate is shown in Figure 12 while for agent *j* in Figure 13. The curve for always cooperation is higher than that for random cooperation which is still higher than that for never cooperation. These curves also tell that when the cost of communication and interaction increase, the utility is going to be worse and worse. Even with different values of cost, rational cooperation in achieving joint activities should be adopted. Figure 14 and Figure 15 show that after credit is added up with utility, the trend of those curves are still the same as without credit.

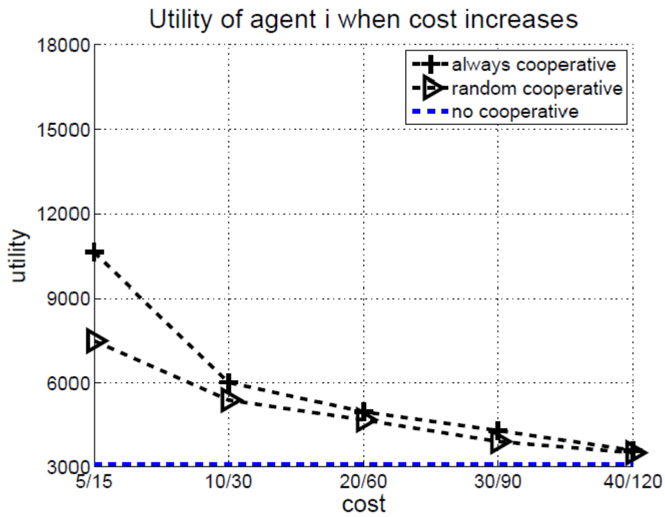

Figure 12. The utility of agent *i* when different cooperation attitude.

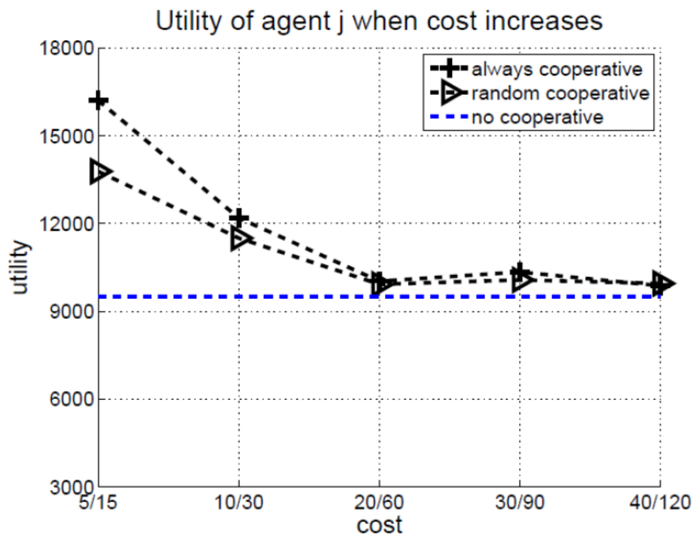

Figure 13. The utility of agent *j* when different cooperation attitude.

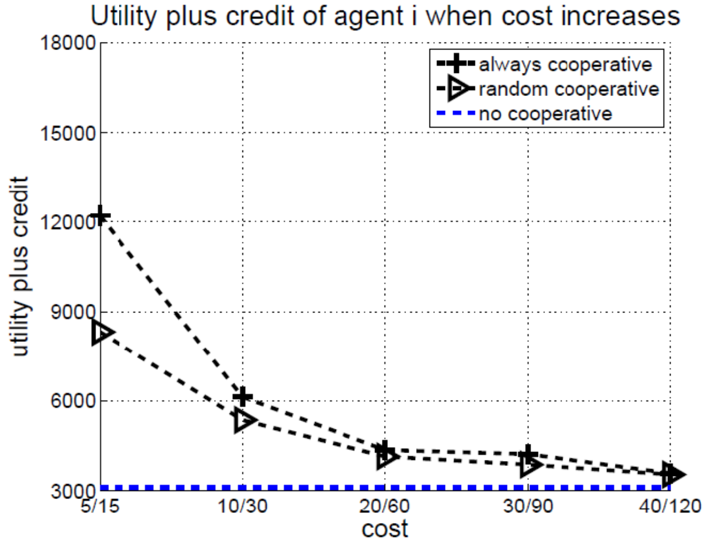

Figure 14. The utility plus credit of agent *i* when different cooperation attitude.

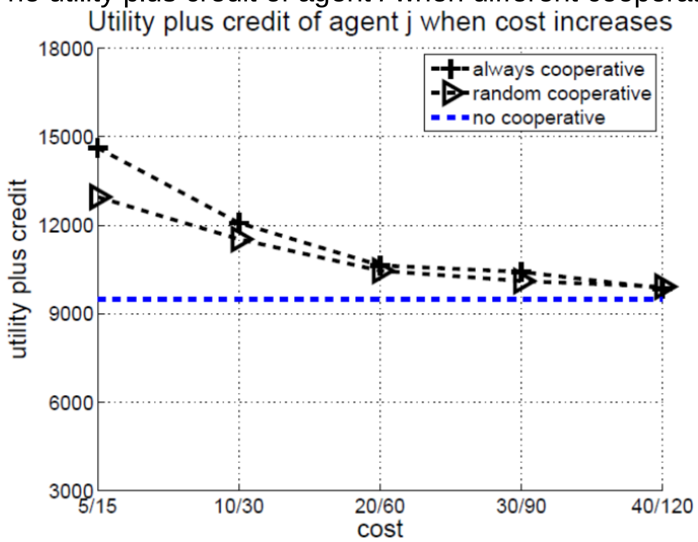

Figure 15. The utility plus credit of agent *j* when different cooperation attitude.

5.3 Experimental setup and result of the second set of experiments

The second set of experiments compares two protocols. One is never communication and always communication protocol. Another is always interaction but no interaction protocol. The bench marks are neither communication nor interaction protocol, and always be rational to cooperate. In this case, the costs of communication and interaction are still changing from 5/15 to 40/120 step by step.

From Figures 16 and 17, we can see that when there is no communication/interaction, the outcomes of both agents are better than when there is neither communication nor interaction. The reason is whenever there is communication/interaction, the agents will get their utility increased from that. However, they are still worse than the case with both communication and interaction most of the time. These confirm that either communication or interaction can already help while both can work the best for the agent.

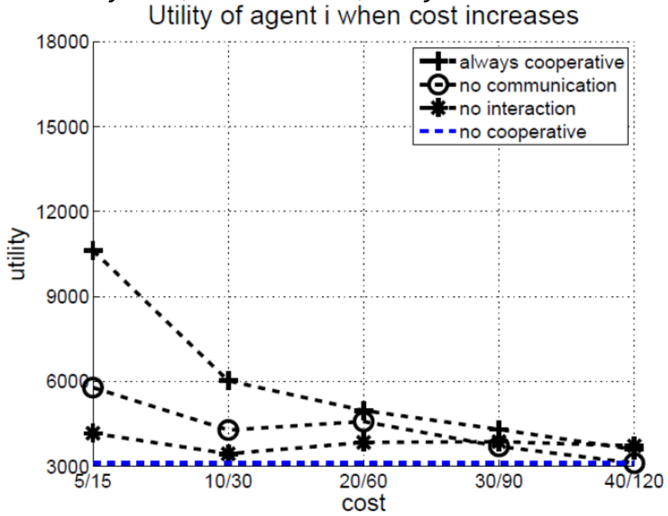

Figure 16. The utility of agent *i* for set 2.

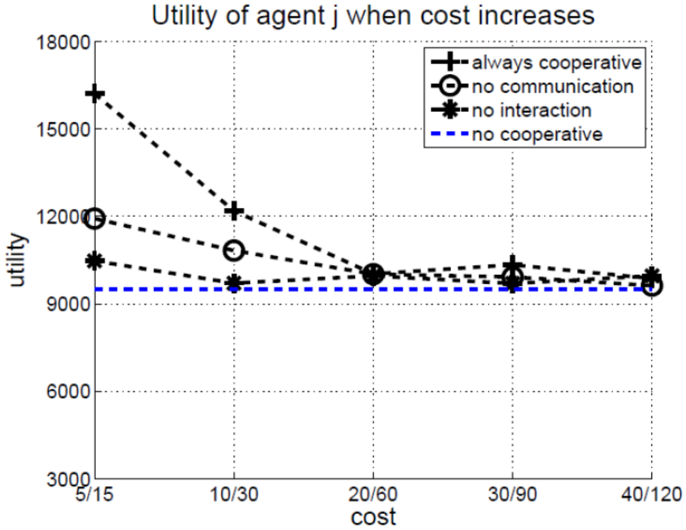

Figure 17. The utility of agent *j* for set 2.

5.4 Experimental setup and result of the third set of experiments

In this set of experiments, we assume that one agent has uncertainty on her expectation from another agent. In particular, when one agent tries to ask for information communication or ask for cooperative action from another agent, she is not certain whether she can get a positive feedback as shown in Figure 4 and Figure 8. We give a variable to express such uncertainty and times it by the previous expected utility from the cooperation.

In this set of experiments, we first set the value of uncertainty of agent *i* to be 0.0 (100% sure), 0.2, 0.4, 0.8, and 1.0 (not sure at all) while keeping agent *j*'s uncertainty always 0.0. All the agents are always rationally cooperative. It can be seen from Figure 18 that when the cost is 5/15, the performance of agent *i* while the uncertainty is small, 0.0, is better than other cases with bigger uncertainties. The reason is that when the agent is sure that she can get a positive feedback and the other agent indeed will provide a positive feedback, she will make a proper decision on whether it is worthwhile to ask for communication or interaction.

However when the cost is going to increase, the performance of agent *i* with various uncertainties is becoming close and fluctuating. The situation is when the cost is bigger and bigger, the difference on uncertainty will affect less and less on whether it is worthwhile to ask for cooperation. That is why the performance will be almost the same when the cost is 40/120. More explanation can be seen clearly in Figure 19 where the amount of successful turns for different uncertainties is similar when the cost is 40/120.

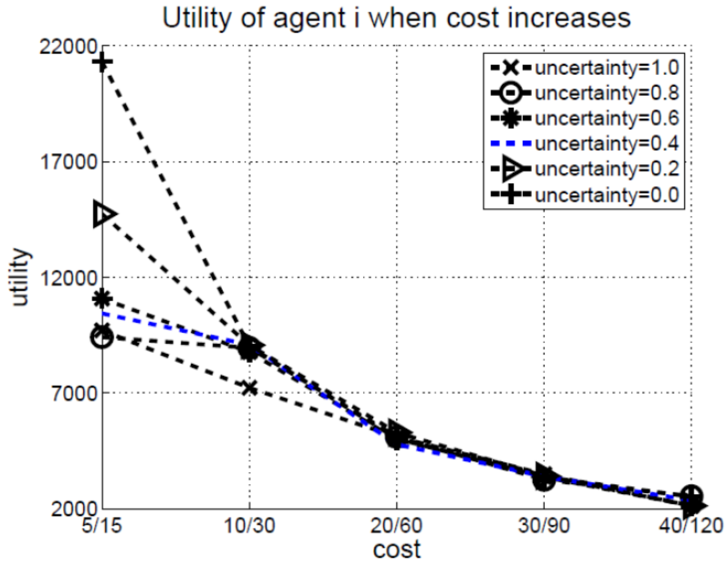

Figure 18. The utility of agent *i* for set 3.

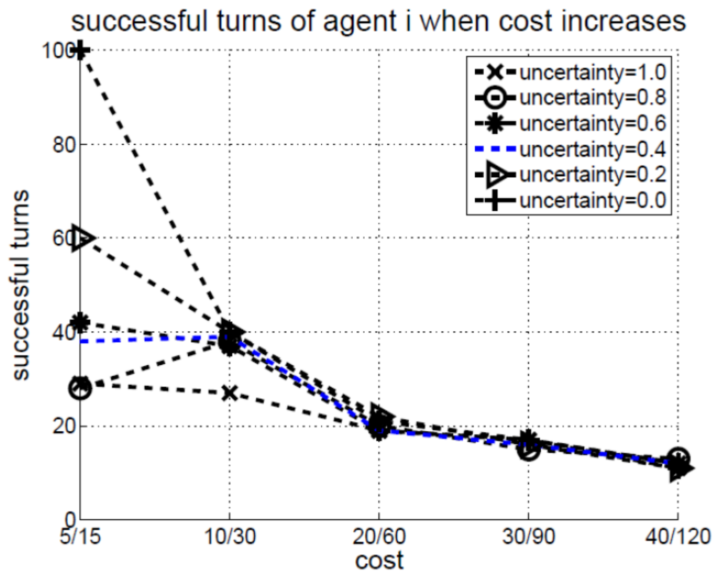

Figure 19. The successfully turns out of 100 turns of agent *i* for set 3.

The performance of agent *j* is shown in Figure 20 where her uncertainty varies but agent *i*'s unchanged. The trend of utility is similar with that of agent *i* in Figure 18. However the difference among curves is quite small for agent *j*. The reason behind is that the amount of successful turns are quite small in this case.

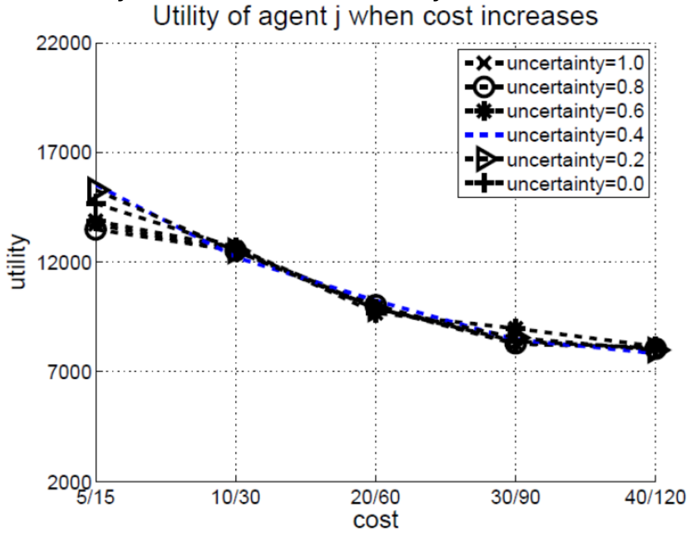

Figure 20. The utility of agent *j* for set 3.

5.5 Experimental setup and result of the fourth set of experiments

According to the results in the previous three sets of experiments, we assume that agents are always rationally cooperative and the uncertainty is equal to 0.0 in the fourth set of experiments. In this set of experiments, we test how the selfishness or altruism will affect the performance of agents. We assume one agent varies her selfishness from 0.0 to 1.0 while another agent keeps her selfishness to be 1.0.

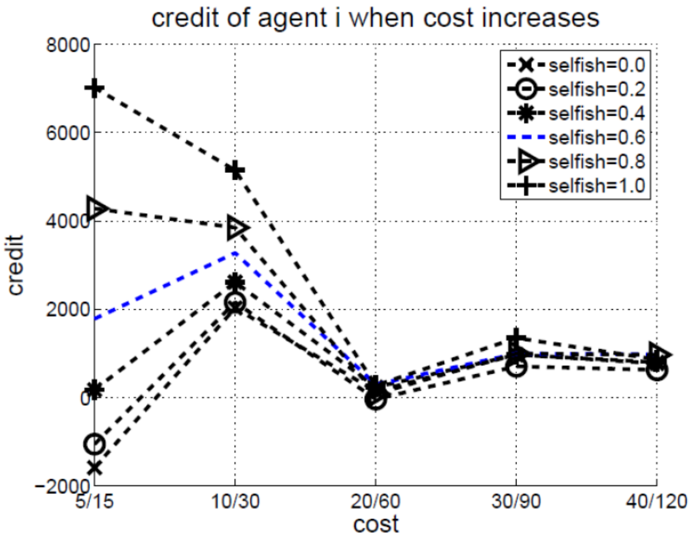

Figure 21. The credit of agent *i* for set 4.

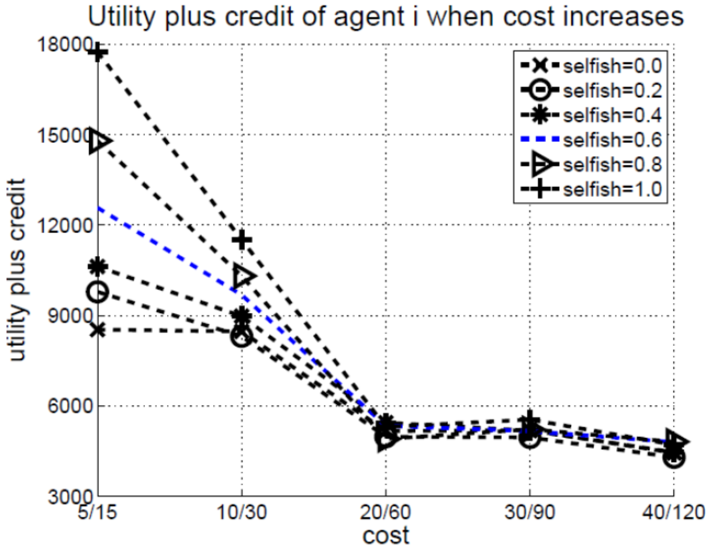

Figure 22. The utility plus credit of agent *i* for set 4.

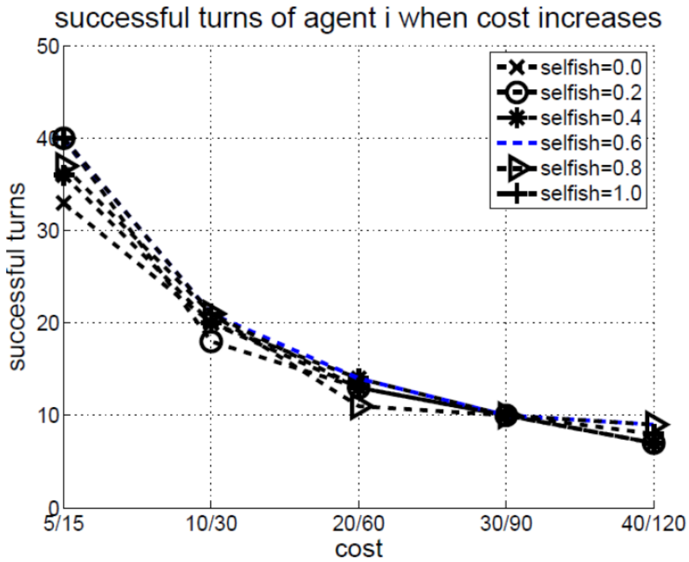

Figure 23. The successful turns out of 100 turns of agent *i* for set 4.

According to Figure 23, at each column of costs, the amount of successful turns is very close no matter how selfish one is. From Figures 21 and 22, we can see that when the costs are small, high value of selfishness will improve agent *i*'s performance. When the costs are higher and higher, selfishness or not will not make a big difference because the amount of successful joint activities is very few.

6. CONCLUSION

In the current activity-travel scheduling systems, communication and interaction are required by joint activities so that agents can form a good cooperation. In the meanwhile dynamics and uncertainties have been introduced through these communication and interaction. To handle the dynamics and uncertainties in a multi-agent setting, agents need to be able to communicate and interact so as to collect information and adjust their behaviour in a timely manner, which brings a shift from pre-specified scheduling to dynamic joint scheduling.

A novel probabilistic representation of other agents' beliefs about the actions selected for their own or for the joint activity is adopted in the paper, given incomplete information. Agents can use this representation to make several decisions: communicating/receiving information relevant to other group member(s), and adding/receiving actions that are helpful to other agent(s) into the cooperative joint scheduling. The corresponding mechanisms are proposed for agents to adopt for reasoning about the utility and credit of interactions and communications, and the cost incurred.

It has been tested using a multi-agent simulation with configurations that vary agents' uncertainty about the world and the cost of action and communication. In all cases, agents using the proposed mechanism to decide whether to interact or communicate have outperformed agents following other mechanisms. Different values of cost, uncertainties, and selfishness/altruism have been shown to affect the performance of agents with various mechanisms.

Therefore the main contributions of the paper are: to integrate agents' interaction and communication into the dynamic (re-)scheduling process; to incorporate credit into decision making mechanisms so as to handle the incentives of cooperating; to propose the strategies for agents to decide whether to interact/communicate with other members of the group; to illustrate the effect of cost and selfishness/altruism on the performance of strategies proposed in a multi-agent simulation on activity travel scheduling.

In the future, we will research on the negotiation problem raised in the joint activity travel scheduling process. Although credit has been slightly covered in the current work, it will be fully discussed from the aspect of incentive compatible mechanism design in our future work.

7. ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Research Council under the *European Community's* Seventh Framework Programme (*FP7/2007-2013)* / ERC *grant agreement* n° *230517* (U4IA project).

The views and opinions expressed in this publication represent those of the authors only. The ERC and European Community are not liable for any use that may be made of the information in this publication.

European Research Council erc

EURATOM

8. REFERENCES

Arentze, T.A. and H.J.P. Timmermans (2007), Modelling dynamics of activity-travel behaviour. In: *Proceedings 12th HKSTS Conference*, Hong Kong (CD-Rom: 30 pp.).

Arentze, T.A. and H.J.P. Timmermans (2009), A need-based model of multi-day, multi-person activity generation, Transportation Research B, 43, 251-265.

Cohen, P.R., Levesque, H.J., Smith, I. (1997). On team formation. In: Contemporary Action Theory Synthese, Kluwer Academic Publishers 87-114.

Ettema, D.F., Arentze, T.A., Timmermans, H.J. (2007). Social influences on household location, mobility and activity choice in integrated micro-simulation models. In: Proceedings of the workshop on Frontiers in Transportation; Social interactions, Amsterdam.

Fan, X., Yen, J., Volz, R.A. (2005). A theoretical framework on proactive information exchange in agent teamwork. Artif. Intell. 169(1) 23-97.

Goldman, C.V., Zilberstein, S. (2004). Decentralized control of cooperative systems: Categorization and complexity analysis. Journal of Artificial Intelligence Research 22, 143-174.Grosz, B.J., Kraus, S. (1996). Collaborative plans for complex group action. Artif. Intell. 86(2) 269-357.

Jacyno, M., Bullock, S., Luck, M., Payne, T.R. (2009), Emergent service provisioning and demand estimation through self-organizing agent communities. In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems, 481-488.

Kamar, E., Gal, Y., Grosz, B.J. (2009). Incorporating helpful behavior into collaborative planning. In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems, 875-882.

Levesque, H., Cohen, P., Nunes, J. (1990). On team formation. In: Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90). 94-99.

Rieser, M., Nagel, K. (2009). Combined agent-based simulation of private car traffic and transit. In Proceedings of The 12th International Conference of the International Association for Travel Behaviour Research.

Tambe, M. (1997). Towards flexible teamwork. Journal of Artificial Intelligence Research 7 83-124.

Vasirani, M., Ossowski, S. (2009). A market-inspired approach to reservation-based urban road tra±c management. In: AAMAS '09: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems 617-624.