

A Spatio-Temporal Model Towards Ad-Hoc Collaborative Decision-Making

Martin Raubal¹ and Stephan Winter²

¹ Department of Geography, University of California, Santa Barbara, USA
raubal@geog.ucsb.edu

² Department of Geomatics, The University of Melbourne, Australia
winter@unimelb.edu.au

Abstract

For an autonomous agent, performing a task in a spatio-temporal environment often requires interaction with other agents. Such interaction can be initiated by ad-hoc collaborative planning and decision-making, which then leads to physical support on site. On-site collaboration is important for a variety of operations, such as search-and-rescue or pick-up-and-delivery. Tasks are performed through sequences of actions, and agents perceive possibilities for these actions in terms of affordances from the environment. Agent collaboration therefore requires the communication of affordances between agents with different capabilities. This paper introduces a spatio-temporal model for the decentralized decision-making of autonomous agents regarding on-site collaboration. Based on Janelle's time-geographic perspective on communication modes, we demonstrate that different task situations lead to different spatio-temporal constraints on communication, involving both physical presence and telepresence. The application of such constraints leads to an optimized message distribution strategy and therefore efficient affordance communication with regard to maximizing support in performing a given task.

1 Introduction

Performing a task in a spatio-temporal environment often involves collaboration between various agents. Examples include rescue teams in emergency response, intelligent robots supporting humans in elderly care, or ride-sharing with clients seeking hosts for transportation. In order to identify potential helpers, agents need to communicate their tasks and needs to others. Such communication and subsequent decision-making involves at least three different perspectives: The *spatio-temporal* view includes issues such as the distribution of helping agents in space, the urgency of solving a task, and the task location. The *social* point-of-view deals with the willingness of other agents to cooperate, and also involves institutional and legal constraints regarding agent cooperation. The *technical* level tackles issues related to the communication infrastructure. Within the framework of peer-to-peer communication, these perspectives can be viewed as different network levels. Previous research in the area of agent collaboration has focused on the social and technical levels. In this paper we propose a *spatio-temporal model towards collaborative decision-making*, which integrates all three network levels. We hypothesize that the integration of spatio-temporal constraints within a model of agent collaboration makes an optimized message distribution possible and therefore results in efficient communication leading to optimal support in performing a task.

The presented model is based on the communication of affordances between agents in a network of peers. The use of affordances allows us to consider action possibilities, which can be formally represented in a functional framework. The model accounts for the fact that different task situations—with respect to urgency, risk, and location—result in different time-geographic communication constraints. Negotiation between client agents and potential helping agents is represented within a *request-offer-choice* process, which includes affordance-based similarity measurement that takes agent capabilities into account.

Section 2 presents related work in the area of agent planning, peer-to-peer communication, and agent collaboration. Section 3 introduces the individual components and theories, on which the model is founded. In Section 4, we develop the collaborative agent process model, detailing the negotiation process. Section 5 applies this process model to a hypothetical emergency scenario. Section 6 discusses the application. The final section presents conclusions and directions for future research.

2 Related work

2.1 Agent planning

According to the heterogeneity of the involved fields there is no common agreement about a definition of the term *agent* [1]. An agent can be anything, such as a robot that perceives its environment through sensors and acts upon it through effectors [2]. More specifically, agents are considered computer systems that are situated in some environment and can act autonomously [3]. *Multi-agent systems* (MAS) depict systems as a combination of multiple autonomous and independent agents and are therefore well suited to simulate collaboration of different actors.

Agents can be represented as functions that map percepts to actions. Abstract models of agents distinguish between purely reactive agents, agents with subsystems for perception and action, and agents with state. These abstract models can be implemented in different ways, depending on how the decision-making of the agent is realized. Here, we consider *utility-based agents* [2], which have valuation functions that allow them to compare between different action sequences to achieve a goal. Such functions map world states to real numbers, which describe associated degrees of happiness.

Agent-based modeling and simulation has gained much popularity in the field of Geographic Information Science due to the disaggregate nature of agents and their ability to move across different spatial scales and representations [4]. Application scenarios include the modeling of urban phenomena [5], pedestrian movement [6], and shared-ride trip planning [7].

Planning is the development of a strategy for solving a task. For an agent, a plan is an action sequence, where each action to be performed depends on some pre-conditions. Every action causes effects or post-conditions that affect or trigger subsequent actions in the chain. The plan terminates when the goal is reached. A planner in Artificial Intelligence takes three input variables: a representation of the initial state of the world, a representation of the intended outcome (goal), and a set of possible actions to be performed to reach the goal. Formally, a plan is a triple $\langle O, I, AC(p, q) \rangle$ [8] where O is the intended outcome, I the initial state of the world, and AC a set of actions—each defined via pre- and post-conditions p, q . After executing actions the state of the world is changed, which impacts the future plan, therefore planning is a non-linear process. One of the main challenges within dynamic environments is that one can neither assume complete knowledge nor the availability of objects and other agents supporting certain actions.

2.2 Peer-to-peer communication

Peer-to-peer (P2P) communication is ad-hoc communication between distributed agents, without involvement of a dedicated server providing communication services to its clients, or any other hierarchic communication infrastructure. It enables mobile agents to collaborate in an ad-hoc manner provided that they agree on a communication protocol. In a P2P communication network each node is of equal importance. Nodes can take the role of a communication client, receiving services from other nodes, but they are also service providers for the other nodes. P2P communication networks are transient in nature, with nodes entering and leaving the network freely and frequently.

A special class of P2P communication is characterized by mobile nodes. For this class the communication is realized wirelessly by radio, which is short-range, due to the typically limited on-board energy resources. This means connectivity in mobile networks depends on the physical distance of nodes, which is constantly changing. Communication over larger distances relies on message forwarding and routing [9, 10].

2.3 Agent collaboration

In a MAS, autonomous agents may cooperate with others to achieve their goals. How to achieve meaningful coordination is a difficult issue and requires interdisciplinary work within the recently established field of *computational cognitive social science* [11]. Agent collaboration is based on cognitive architectures, representing and explaining cognitive processes of individual agents during the performance of tasks. The most prominent architectures are ACT-R [12] and Soar [13]. The CLARION cognitive architecture [14] extends cognitive modeling to social simulation. It accounts for agents' socially oriented goals and bases agent cooperation on the fact that social interaction between agents is made possible through their understanding of each other's motivations.

Previous research has focused on technical aspects of P2P collaboration [15] and hierarchical multi-agent models integrating knowledge-based communication, such as for the RoboCupRescue¹ simulation [16]. Luo and Bölöni [17] presented a game-theoretic model for the canonical problem of spatio-temporal collaboration with the goal of optimizing individual benefits. A typical application of intelligent agent collaboration is elderly care. For example, the PEIS (Physically Embedded Intelligent Systems) ecology is a network of heterogeneous smart devices that ranges from simple gadgets, such as refrige-

¹ <http://www.robocup.org/>

rators with sensors, to sophisticated mobile robots or even humans [18]. These intelligent entities communicate and collaborate with each other by providing information, and combining physical and virtual functionalities to perform complex tasks, such as supporting human inhabitants in their flats. Shared-ride trip planning is an urban transportation application of ad-hoc agent collaboration: client agents representing customers, e.g., pedestrians, seek transportation by host agents representing vehicles, e.g., public transportation vehicles, in an ad-hoc manner [7, 19]. The goal of collaboration is to bring clients to their destinations, for some host benefits. In typical scenarios multiple clients compete for free capacity and multiple hosts compete for clients. Recently, it was demonstrated how decentralized time geography can be applied to ad-hoc collaborative agent planning [20]. Agents performed a spatio-temporal analysis based on local knowledge in a distributed environment, thereby evaluating whether they can independently contribute to physical support at a specific site by a specified time. Experiments in a multi-agent simulation framework investigated the impact of different combinations of agent density and communication radius, as well as behavioral strategies on task performance.

This paper takes previous work a step further by explicitly considering the interaction of mobile agents in terms of a combination of *spatio-temporal*, *social*, and *communication network* levels. The goal is to develop a general and comprehensive model of spatio-temporal decision-making for ad-hoc agent collaboration. We specifically investigate the cooperation process, which includes communication of affordances between agents. Such communication enables clients and service providers to form better decisions by taking the individual capabilities of agents into account. The model also provides a way to determine the influence of task dimensions and spatio-temporal constraints on the efficiency of communication.

3 Spatio-temporal framework of collaborative decision-making

This section specifies different task dimensions, whose values lead to spatio-temporal constraints for communication. We start by introducing affordances, which are communicated by the collaborating agents, and a recently proposed similarity measure for them. Time geography serves as the basis for modeling the spatio-temporal decision-making process.

3.1 Affordance representation and similarity

The theory of affordances [21] is based on the tenet that agent and environment form an inseparable pair. Affordances have to be described relative to the agent. For example, a chair’s affordance ‘to sit’ results from a bundle of attributes, such as ‘flat and hard surface’ and ‘height’, many of which are relative to the size of an individual agent. Norman [22] recasts affordances as the results from the mental interpretation of things, based on people’s past knowledge and experiences, which are applied to the perception of these things.

In order to supplement Gibson’s theory of perception with elements of cognition, situational aspects, and social constraints, Raubal [23] presented an extended theory of affordances suggesting that affordances belong to three different realms: *physical*, *social-institutional*, and *mental*. This distinction was formally specified in a functional model [24]. The agent is represented through its physical structure, spatial and cognitive capabilities, and a goal. *Physical affordances (Paff)* for the agent result from invariant compounds—unique combinations of physical, chemical, and geometrical properties—and the physical structure of the agent. *Social-institutional affordances (Slaff)* are created through the imposition of social and institutional constraints on physical affordances. *Mental affordances (Maff)* arise for the agent when perceiving a set of *Paffs* and *Slaffs* in an environment at a specific location and time. Affordances offer possibilities for action as well as possibilities for the agent to reason about them and decide whether to utilize them or not, i.e., mental affordances.

We specify an affordance A as a triple $\langle O, E, \{AC\} \rangle$ [25]. The outcome O is the change of world state after executing the actions AC with respect to manipulated entities of type E . Each action is represented by physical (ph) and social-institutional (si) constraints or pre-conditions, AC therefore being defined as a set of actions $\{ac_1(ph_1, si_1), \dots, ac_n(ph_n, si_n)\}$. Constraints are tied to a certain action with respect to an entity, while the outcome is equal for all actions defined for the affordance A . An affordance can be utilized through several actions, e.g., the move-ability affordance of a desk may include the actions *carry* and *push*.

When communicating affordances between agents and evaluating whether an offered affordance is good enough compared to the requested affordance to help in solving a task, it is necessary to determine their similarity. Here, we apply a similarity measure for affordances [26], which uses the action and outcome specification from $\langle O, E, \{AC\} \rangle$. Affordances are more similar the more similar their descriptors are. The overall similarity Sim_A between affordances A_s and A_t is defined as the weighted sum for the individual similarities computed for actions (sim_{AC}) and outcomes (sim_O) (Equation 1). The former depend on the similarity values computed for their physical and social-institutional constraints.

$$Sim_A(A_s, A_t) = \omega_{ac} * \frac{1}{n} \sum sim_{AC} + \omega_o * sim_o; \text{ where } \sum \omega = 1 \quad (1)$$

3.2 Task dimensions

Agents can perform tasks by utilizing various affordances. For example, to change a light bulb, an agent can move a chair below the light, step onto it, and change the bulb [27]. When planning how to solve a task, the agent must take several aspects into account. For the purpose of modeling spatio-temporal collaborative decision-making, we consider the following task dimensions: *collaboration*, *urgency*, *risk*, and *location*. Table 1 describes and explains the possible values for each of them.

Table 1: Task dimensions with possible values and explanations.

Task dimension	Value	Explanation
<i>Collaboration</i>	0	Agent can solve task alone.
	1	Another agent needed to either solve task or support requesting agent in (sub)task.
	n	More than 1 agent needed to solve task.
<i>Urgency</i>	immediate	Immediate help required, values depend on context, e.g., within 10 min.
	flexible	Help required within reasonable time frame.
<i>Risk</i>	high	Describes the inverse probability that the agent(s) will solve the task.
	medium	
	low	
<i>Location</i>	(x, y)	Coordinate pair; default location for the task is the current location of the requesting agent.

3.3 Spatio-temporal communication constraints

Agents and resources are available at a limited number of locations for a limited amount of time. Time geography defines the space-time mechanics of locational presence by considering different constraints [28]. The original time geography framework recognized the possibility of telepresence using electronic communication, although it received much less attention than physical presence. Time geography's focus on time as a resource enabling activity participation has received explicit interest by researchers lately [29, 30]. It fits naturally to views of time as the major scarce resource in information economies and accelerated modern lifestyles [31].

Janelle [32] classified communication modes from a time-geographic perspective. Table 2 summarizes classes based on their spatial and temporal constraints. Spatial constraints are either physical presence or telepresence, while temporal constraints are either synchronous or asynchronous. *Synchronous presence* (SP) is the communication mode of face-to-face (F2F) interaction. F2F requires coincidence in both time and space. *Synchronous telepresence* (ST) requires only coincidence in time. *Asynchronous presence* (AP) requires coincidence in space but not in time. *Asynchronous telepresence* (AT) does not require coincidence in space and time.

Table 2: Spatio-temporal communication constraints, based on [32].

Temporal	Spatial	
	Physical presence	Telepresence
Synchronous	SP Face to face	ST Telephone, instant messaging, radio, teleconferencing
Asynchronous	AP Refrigerator notes, hospital charts	AT Email, fax, printed media, web pages

The spatio-temporal communication constraints for agent collaboration depend directly on some of the task dimension values. Constraints exist if $collaboration \geq 1$ and they vary for different *urgency* values. If *urgency* = *immediate*, then only *SP* and *ST* of other agents lead to potential help in solving a task, because if help comes after some time threshold, the utility for the requesting agent is zero. Take, for example, an emergency scenario where someone who cannot get out of a car that fell into a river, needs to be rescued. There is only a small critical time interval for survival. If *urgency* = *flexible*, then *AP* and *AT* may also be viable, depending on how much time the requesting agent has for solving the task. *AP* may lead to a higher risk because it is assumed that some other agent will come by the requesting agent's location within a certain time interval and react to a posted message. In general, flexible task urgencies result in more choices and less constrained communication.

4 Collaborative agent process model

This section develops a high-level framework for ad-hoc negotiation of on-site collaboration between agents. We allow for autonomous agents that follow their individual goals, and only if they cannot reach them on their own, they ask peers for help.

4.1 Communication for collaboration

Agents use P2P communication to negotiate with each other for collaborative action. One defining parameter for the design of a negotiation procedure is the radio range, which is an issue in all P2P communication and depends on the protocol / platform. To expand the search range for help beyond the immediate radio range, message forwarding strategies can be applied, e.g., within a specified search range pre-calculated by spatio-temporal relevance constraints [7]. We assume here that agents have sufficient energy for movement and physical work on board, and may even utilize internal mechanisms for battery recharge. This allows them to broadcast anytime—in contrast to sensor networks, which are rigidly limited by their energy resources and communicate only in synchronized time windows.

Negotiations within a search range require stable communication links in a potentially fragile communication network over the time of the negotiation process. This requirement makes the problem substantially different from simple message dissemination problems [9]. Robust negotiation strategies require negotiating over short distances and by immediate response. Also, in case a message gets lost, agents must be able to continue their work based on their current information. Negotiations can take different forms. One way is a client agent sending a request, interested agents responding with offers, and the client selecting and booking an offer. Another way is providers advertising their services, clients registering, selecting, and booking when needed. Only the prior form allows for synchronous communication, which facilitates robust negotiations.

Figure 2 illustrates the negotiation process. The client initiating the communication reaches four other agents within its radio range (dark gray). Since these agents are located within the search range (light gray), they re-broadcast the request once. Other agents receiving the request will also re-broadcast once if they are located within the search range. Some agents beyond the search range may have received the request, but since they are outside this range they ignore it. For example, for *Agent 8* being within radio range of *Agent 5* but outside the search range means that it receives the client's request, but does not re-broadcast. Therefore, *Agent 9*, although within radio range of *Agent 8*, will not receive the request. In contrast, *Agent 7* is within the search range but outside the radio range of any broadcasting agent and therefore does not receive the request. The request messages traveling through the communication network keep track of their broadcasting agents. This way, the message remembers the shortest route back to the client.

Agent 1 is sufficiently close to be in *synchronous co-presence* (black circle) with the client and can start collaboration without delay. *Agents 2, 3, and 4* are in *synchronous telepresence*: they are able to communicate directly, but have to approach first before an interaction can take place. Agents that can on-

ly be reached by the client through mediating agents are in asynchronous telepresence, since communication can be delayed by the requirement to re-broadcast sent messages. In situations with $radio\ range \geq search\ range$ asynchronous telepresent agents do not exist.




Figure 2: A client agent (triangle) requesting help via short-range radio and the emerging communication network.

4.2 The negotiation process

This section models the individual steps of the negotiation process. We assume that agents can communicate ad-hoc and make decisions based on utility functions. Accordingly, the focus is on *communication and exploitation of spatio-temporal constraints and affordances*.

4.2.1 Client's request

A *client* initializes a negotiation process as soon as it is confronted with a task beyond its capabilities ($collaboration > 0$). The task focalizes the client's perceived affordances and enables it to create a plan. The client may discover the need for help by learning about the physical and social parameters of the task, or the individual actions involved in the plan. The parameters may be released from the object to be manipulated, or experienced by trial-and-error. At this time the client cannot judge whether a single or multiple collaborative agents are needed to solve the task because it does not know the capabilities of nearby agents. However, the client is able to specify the urgency of performing the task. By setting an upper time limit the client implicitly defines a search range for helping agents. With these parameters a request can be formulated, consisting of:

- *Message type*: 'request' (tells other agents how to treat this message);
- *task location*: client location by default;

- *search range*: a time frame;
 - *task*: a set of $\langle O, E, \{AC\} \rangle$. In this set of affordances the specific subtask the client needs help for has the highest weight. It is parameterized by the difference of learned properties of the subtask and the client's own capabilities. If the client has experienced the properties by trial-and-error, the parameterization takes the form of an inequality, denoting that the value is beyond its own capabilities;
 - list of *broadcasting agents*: agent IDs, initialized by the client's ID.
- The request is broadcasted by the client and re-broadcasted by other agents if they are within the search range. Every forwarding agent appends its ID to the list of broadcasting agents.

4.2.2 Service provider's offer

Recipients of the request take the requested affordances and their parameters into account, as well as their own capabilities, duties (urgency of their own current tasks), and utility functions. An agent can formulate two types of offers: (a) contributing to the specific affordances requested by the client, or (b) suggesting different actions to solve the task. For case (a), the simplest situation occurs if the agent can offer to utilize a requested affordance on its own. However, if the request contains a parameter specification in the form of an inequality or a value exceeding its capabilities, the agent can only offer help within the limits of these capabilities. If the agent cannot utilize the requested affordance, but a similar one, it can still offer this similar affordance. For case (b), the agent can apply the affordance similarity measure (Section 3.1) to its own stored affordances and come up with a suggestion. The agent's offer consists of:

- *Message type*: 'offer';
- *travel time distance* to client;
- offered parameterized *affordance* in terms of actions;
- *similarity value* between offered and requested affordance: may be used by client to estimate the risk with booking this agent;
- list of *agents* leading back to client: reverted list of broadcasting agents;
- list of *broadcasting agents*: agent IDs, initialized by offering agent's ID.

Re-broadcasting of *offers* by other agents is conditional to their ID appearing in the list of agents leading back to the client, and their ID not appearing on the list of broadcasting agents (to avoid multiple broadcasting). This strategy assumes that the communication links available for the request are still intact for the offers. While this is realistic in general, in individual cases an offer may not reach the client due to a recently broken link.

4.2.3 Client's choice and booking

The client will compare the travel time distance specified in the offers with its task urgency, determine the similarity between offered and requested affordances, and the amount of support offered. With its own utility function the client is able to rank all incoming offers. The following cases can be distinguished:

1. *No offer arrived.* The client can enlarge the search radius or change its plan.
2. *At least one offer is made matching the highest ranked affordance in the request.* The client can choose either the nearest offering agent (high task urgency) or the agent with the largest capacity for this affordance (risk reduction).
3. *Offers rank other affordances higher than the requested affordance.* The client can choose the offer with the most similar affordance, assuming compatibility in agent capabilities. Alternatively, it can choose the offer with the highest weight for one affordance, assuming that the weight reflects the offering agent's confidence in being helpful. Accepting other than the requested affordance may require revising the plan.

Once the client has made a decision, it will formulate a booking message, which consists of:

- Message type: 'booking';
- booked *affordance*;
- list of *agents* leading forward to the offering agent: reverted list of the chosen offer's broadcasting agents;
- list of *broadcasting agents*: agent IDs, initialized by client's ID.

Re-broadcasting of requests by other agents is conditional to their ID appearing in the list of agents in the chain forward to the offering agent and their ID not appearing in the list of broadcasting agents. Booked agents will travel to the client's location and help.

5 Application scenario

This section applies the model to a hypothetical emergency scenario involving a car that got hit by a tree (Figure 3). The client *Agent C* (car driver) tried to move the tree without success and is therefore requesting immediate help from other agents² in the communication network. We consider four additional agents (Figure 4): *Agents H₁* and *H₂* are within radio range of *Agent C*; *Agent H₃* is within the search range; and *Agent H₄* is outside the search range.

² Our model is generic and deals with abstract agents. Here, we focus on the description of the collaboration process between software agents, whether the actors represent humans or not.



Figure 3: A car driver finding his vehicle blocked by an obstacle contacts other agents to help remove the obstacle.

The client's request consists of the following parameters (Section 4.2.1):

```
[request; (34.42, -119.70); 15min;
<O: hasPos (e, Pos(y)) & y≠x; E: tree;
AC: carry (ph: hasPos (e, Pos(x)) & WeightKg (e, >30) &
LengthM (e, >2))>;
(C)]
```

It contains the task location in the form of latitude/longitude coordinates and specifies a search range of 15 minutes. The move-ability affordance is represented through an outcome O (entity e must have a different position y compared to current location x), an entity type E , and one action specified by physical aspects ph . Due to its physical capabilities, *Agent C* can only carry trees with a maximum weight of 30kg and a maximum length of 2m. The requested action (as part of the affordance) therefore exceeds these limits. After the outcome O , C initializes the list of broadcasting agents being the first sender.




Figure 4: Client *Agent C* with four potential service providers in the communication network.

All four additional agents receive the client's request. H_1 and H_2 are in synchronous telepresence, and receive the request directly. H_3 and H_4 are in asynchronous telepresence (reached through mediating *Agents* H_2 and H_3). H_2 is occupied with its own task and therefore decides not to make an offer. The

other three agents calculate the shortest path [33] to the task location. Only agents that can reach the task location within the specified time will make an offer. This results in H_4 not making an offer. Thus, only H_1 and H_3 are making an offer to the client. Table 3 compares these offers.

Table 3: Offers from Agents H_1 and H_3 back to the client.

Offer from Agent H_1	Offer from Agent H_3
8min; AC: carry (ph: hasPos (e, Pos(x)) & WeightKg (e, ≤ 40) & LengthM (e, < 1)); 0.84; (H_1 , C); (H_1);	13min; AC: lift (ph: hasPos (e, Pos(x)) & WeightKg (e, ≤ 800) & LengthM (e, [1,15])); 1.00; (H_3 , H_2 , C); (H_3);

The computed travel times from the locations of H_1 and H_3 to the task location are 8 and 13 minutes. With respect to the parameterized affordance, two different actions—*carry* and *lift*—are offered. O and E are equal to the client’s request ($sim = 1$) and therefore omitted in the table. The individual components of each action are used to calculate overall affordance similarity according to the measure introduced in Section 3.1. Because the client’s request specifies minimum agent capabilities for carrying trees in terms of weight and length, every offer equal to or exceeding this limit results in a similarity value of 1, e.g., WeightKg (e, > 30) and WeightKg (e, ≤ 40), and 0 otherwise. Final values are calculated according to Equation 1 with both weights set to 0.5³.

The offer broadcasted by H_1 is received by C directly, and no other agent receiving the message takes an action. In contrast, C is not in the radio range of H_3 , but H_3 specified that H_2 should forward its offer. Agent C then evaluates the incoming offers according to a utility function taking various parameters, such as task urgency and risk, into account. In our example, the client puts a higher weight on agent capacity for solving the task and establishes the ranking (H_3 , H_1). A booking message is therefore sent to H_3 .

³ For H_1 , sim_{AC} results from the similarities of hasPos, WeightKg, and LengthM, i.e., $(1+1+0)/3 = 0.67$. Equation 1 then evaluates to $Sim_A = 0.5*0.67+0.5*1 = 0.84$.

6 Discussion

Compared to an uninformed or brute-force approach such as flooding (every sensor receiving a message broadcasts this message again), the application shows that integrating spatio-temporal constraints within a model of agent collaboration in a P2P network leads to an optimized message distribution among agents and therefore to more efficient support in performing a task. During brute-force search messages are re-broadcasted to every other node within the communication range and this process is repeated consistently. In our case, the *optimization of message distribution* results from the constrained search range based on time-sensitive tasks. Messages are only sent to potential collaborators, resulting in a reduction of overall network traffic and saving bandwidth. Agents (H_i in the application scenario), whose travel time to the task location exceeds a given limit, do not make offers and therefore further reduce the number of messages. *Efficient task support* results from knowing in advance the helping agents' capabilities. In addition to similarity values between requested and offered affordances, and the client's utility-based decision-making, this is a major step towards spatio-temporal efficiency in agent collaboration.

The client found through trial-and-error that its physical capabilities were insufficient for moving the tree. As a result, some of the affordance parameters could only be specified in terms of lower limits, e.g., $\text{LengthM} (e, >2)$. By specifying exact capability values, such as $\text{LengthM} (e, 4.50)$ an interval rather than a Boolean scale could be used to calculate more precise similarity values, e.g., 4.49 is more similar to 4.50 than 4.35. It is important to note though that the final similarity values are not interpreted on an individual basis, but establish an order from most to least similar. In our scenario, social-institutional (si) constraints were only implicitly covered—*Agent H_2* did not want to make an offer—and did not enter the similarity function. As shown in [26], the similarity function for actions sim_{AC} can easily be extended to represent these aspects, such as lower willingness to help during nighttime versus daytime.

Similar to this application, ad-hoc shared-ride services realize the general model of decision-making demonstrated here. The decision model as presented in [7] also relies on a negotiation process of client requests, host offers, and clients' selection and booking. In light of the present general decentralized decision model, their application-specific negotiation can be interpreted as being based on affordances. Clients in the shared-ride scenario formulate their request by specifying their current and desired locations. They perceive the affordance of moving vehicles with free transportation capacity, and accordingly, specify in their request the task 'move me to a specific location'. Hosts can interpret this task directly by their capabilities to offer rides, de-

pending on their free seats and directions. The other aspects of the general model of decision-making are also present: Urgency is known, a search radius is specified by the client, and the clients' utility function consists of an optimal path algorithm for their own trip. Therefore, protocols and algorithms in shared-ride services can be expressed by the presented model.

7 Conclusions and future work

In this paper we have specified a high-level framework for ad-hoc communication between agents that negotiate for collaboration to perform a task. The underlying model accounts for spatio-temporal constraints, leading to efficient communication and task support. The agents' decision-making is based on affordances, to be able to adapt to any context and task. The framework was demonstrated through an application, which gave insight into how the affordance specification enables clients and service providers to form their decisions.

The presented work suggests several directions for future research:

- The process model needs to be implemented and tested in different real-world application scenarios. Decentralized ride-sharing provides one possible scenario, but there are many others, such as emergency response and various interactions between humans and robots, e.g., in elderly care. Agent-based simulations will provide insights into complexity issues and real-world applicability of spatio-temporal communication constraints.
- The demonstrated application includes only a small number of agents, which leads to the question of scalability. We expect that our approach will scale due to its distributed architecture, local processing, and local evaluations of relevance. Future simulations will address this question.
- The similarity measure for affordances needs to be refined and extended. Strategies for combining offers from different service providers (*collaboration* > 1) must be developed, leading to the classic problem of combinatorial optimization, i.e., determining the set of agents with the largest total attribute value. This becomes even more complex when affordances and their parts cannot simply be added up, e.g., *carry* + *lift*. Determining the similarity between affordances specified through different actions will require the use of action / affordance ontologies.
- There are several ways of specifying a client's utility function. Depending on the context, such function may focus on temporal aspects, risk estimations, and social and institutional issues. In addition, economic models will be needed to balance the costs of the service providers with benefits.

Acknowledgments

The comments from four anonymous reviewers provided useful suggestions to improve the content and clarity of the paper. Dr. Winter acknowledges support from the Australian Research Council, DP0878119.

References

1. Sengupta, R. and R. Sieber, Geospatial Agents, Agents Everywhere ... Transactions in GIS, 2007. 11(4): p. 483-506.
2. Russell, S. and P. Norvig, Artificial Intelligence: A Modern Approach. 2nd ed. Prentice Hall Series in Artificial Intelligence. 2003, London: Prentice Hall.
3. Wooldridge, M., Intelligent Agents, in Multiagent Systems - A Modern Approach to Distributed Artificial Intelligence, G. Weiss, Editor. 1999, MIT Press: Cambridge, MA. p. 27-77.
4. O'Sullivan, D., Geographical information science: agent-based models. Progress in Human Geography, 2008. 32(4): p. 541-550.
5. Benenson, I. and P. Torrens, Geosimulation - Automata-based modeling of urban phenomena. 2004, Chichester, England: Wiley.
6. Batty, M., J. Desyllas, and E. Duxbury, The discrete dynamics of small-scale spatial events: agent-based models of mobility in carnivals and street parades. International Journal of Geographic Information Science, 2003. 17(7): p. 673-97.
7. Raubal, M., S. Winter, S. Teßmann, and C. Gaisbauer, Time geography for ad-hoc shared-ride trip planning in mobile geosensor networks. ISPRS Journal of Photogrammetry and Remote Sensing, 2007. 62(5): p. 366-381.
8. Erol, K., J. Hendler, and D. Nau, Complexity results for HTN planning. Annals of Mathematics and Artificial Intelligence, 1996. 18(1): p. 69-93.
9. Nittel, S., M. Duckham, and L. Kulik, Information dissemination in mobile ad-hoc geosensor networks, in Geographic Information Science - Third International Conference, GIScience 2004, M. Egenhofer, C. Freksa, and H. Miller, Editors. 2004, Springer: Berlin. p. 206-222.
10. Zhao, F. and L. Guibas, Wireless Sensor Networks. 2004, Amsterdam: Elsevier.
11. Sun, R., Prolegomena to Integrating Cognitive Modeling and Social Simulation, in Cognition and Multi-Agent Interaction, R. Sun, Editor. 2006, Cambridge University Press: New York, USA. p. 3-26.
12. Anderson, J., D. Bothell, M. Byrne, S. Douglass, C. Lebiere, and Y. Qin, An integrated theory of mind. Psychological Review, 2004. 111(4): p. 1036-1060.
13. Newell, A., Unified Theories of Cognition. 1990, Cambridge, Massachusetts: Harvard University Press.
14. Sun, R., The CLARION Cognitive Architecture: Extending Cognitive Modeling to Social Simulation, in Cognition and Multi-Agent Interaction, R. Sun, Editor. 2006, Cambridge University Press: New York, USA. p. 79-99.
15. Schoder, D., K. Fischbach, and C. Schmitt, Core Concepts in Peer-to-Peer Networking, in Peer-to-Peer Computing: The Evolution of a Disruptive Technology,

- R. Subramanian and B. Goodman, Editors. 2005, Idea Group Inc.: Hershey, PA. p. 1-27.
16. Peng, J., M. Wu, X. Zhang, Y. Xie, F. Jiang, and Y. Liu, A Collaborative Multi-Agent Model with Knowledge-Based Communication for the RoboCupRescue Simulation. in International Symposium on Collaborative Technologies and Systems (CTS'06). 2006.
17. Luo, Y. and L. Bölöni. Children in the Forest: Towards a Canonical Problem of Spatio-Temporal Collaboration. in AAMAS'07, Int. Conference on Autonomous Agents and Multiagent Systems. 2007. Honolulu, Hawai'i, USA: IFAAMAS.
18. Broxvall, M., M. Gritti, A. Saffiotti, B.-S. Seo, and Y.-J. Cho, PEIS Ecology: Integrating Robots into Smart Environments. in IEEE International Conference on Robotics and Automation (ICRA). 2006. Orlando, Florida.
19. Winter, S. and S. Nittel, Ad-hoc shared-ride trip planning by mobile geosensor networks. International Journal of Geographical Information Science, 2006. 20(8): p. 899-916.
20. Raubal, M., S. Winter, and C. Dorr, Decentralized Time Geography for Ad-Hoc Collaborative Planning, in Spatial Information Theory - 9th International Conference, COSIT 2009, Aber Wrac'h, France, September 2009, K. Stewart Hornsby, et al., Editors. 2009, Springer: Berlin. p. 436-452.
21. Gibson, J., The Ecological Approach to Visual Perception. 1979, Boston: Houghton Mifflin Company.
22. Norman, D., The Design of Everyday Things. 1988, New York: Doubleday.
23. Raubal, M., Ontology and epistemology for agent-based wayfinding simulation. International Journal of Geographical Information Science, 2001. 15(7): p. 653-665.
24. Raubal, M. and R. Moratz, A functional model for affordance-based agents, in Towards Affordance-Based Robot Control - International Seminar, Dagstuhl Castle, Germany, June 5-9, 2006. Revised Papers E. Rome, J. Hertzberg, and G. Dorffner, Editors. 2008, Springer: Berlin. p. 91-105.
25. Stoffregen, T., Affordances and events. Ecological Psychology, 2000. 12: p. 1-28.
26. Janowicz, K. and M. Raubal, Affordance-Based Similarity Measurement for Entity Types, in Spatial Information Theory - 8th International Conference, COSIT 2007, Melbourne, Australia, September 2007, S. Winter, et al., Editors. 2007, Springer: Berlin. p. 133-151.
27. Barsalou, L., Ad hoc categories. Memory & Cognition, 1983. 11: p. 211-227.
28. Hägerstrand, T., What about people in regional science? Papers of the Regional Science Association, 1970. 24: p. 7-21.
29. Ren, F. and M.-P. Kwan, Geovisualization of Human Hybrid Activity-Travel Patterns. Transactions in GIS, 2007. 11(5): p. 721-744.
30. Raubal, M., H. Miller, and S. Bridwell, User-Centred Time Geography For Location-Based Services. Geografiska Annaler B, 2004. 86(4): p. 245-265.
31. Miller, H., What about people in geographic information science?, in Representing Geographical Information Systems, P. Fisher and D. Unwin, Editors. 2005, John Wiley. p. 215-242.
32. Janelle, D., Impact of Information Technologies, in The Geography of Urban Transportation, S. Hanson and G. Giuliano, Editors. 2004, Guilford Press: New York. p. 86-112.

33. Dijkstra, E.W., A note on two problems in connection with graphs. *Numerische Mathematik*, 1959. 1(1): p. 269-271.