Decentralized Time Geography for Ad-Hoc Collaborative Planning

Martin Raubal[1], Stephan Winter[2], and Christopher Dorr[1]

[1] Department of Geography, University of California, Santa Barbara, USA
raubal@geog.ucsb.edu, chdorr@gmail.com
[2] Department of Geomatics, The University of Melbourne, Australia
winter@unimelb.edu.au

Abstract. For an autonomous physical agent, such as a moving robot or a person with their mobile device, performing a task in a spatio-temporal environment often requires interaction with other agents. In this paper we study ad-hoc collaborative planning between these autonomous peers. We introduce the notion of *decentralized time geography*, which differs from the traditional time-geographic framework by taking into account limited local knowledge. This allows agents to perform a space-time analysis within a time-geographic framework that represents local knowledge in a distributed environment as required for ad-hoc coordinated action between agents in physical space. More specifically, we investigate the impact of general agent movement, replacement seeking, and location and goal-directed behavior of the initiating agent on the outcome of the collaborative planning. Empirical tests in a multi-agent simulation framework provide both a proof of concept and specific results for different combinations of agent density and communication radius.

1 Introduction

This paper studies collaboration between mobile social agents, in particular forms of collaboration that require a physical encounter at some point in space and time. Agents communicate in an ad-hoc, peer-to-peer manner, such that each agent has easy access to local knowledge, but no access to global knowledge. There is no central instance in this network that collects data nor does it provide for centralized time-geographic analysis; analysis, where required, has to be done locally by the distributed agents.

Such agents may be mobile social robots, mobile sensor network nodes, or persons with smart mobile devices. Consider, for example, the following problem: A major accident has happened at an industrial site, leading to a spread of chemicals into the surrounding environment. In order to take measurements of several critical variables, a number of agents with different types of sensors are needed at the site within a specified time frame. Traditionally, these agents are activated and managed centrally, but what if the central emergency command center, or the centralized communication infrastructure, was destroyed by the accident? Let us assume that these agents exist somewhere in the environment, potentially off-site, and have the ability to communicate and collaborate in a peer-to-peer manner. Let us further assume that one agent

K. Stewart Hornsby et al. (Eds.): COSIT 2009, LNCS 5756, pp. 436–452, 2009.

discovers the accident and decides that help is needed. The following questions then arise for this initiating agent: How many agents are needed, and how many can be expected at the site by the upper time limit? How large must the search radius for help be to get those agents together? How difficult is it to reach agents within this search radius? What is the risk of limiting a call for help within these boundaries?

Problems such as the chemical spill require *decentralized collaboration* with local interaction of agents to achieve a common goal. Such collaboration takes place under specific space-time constraints in a heterogeneous environment. These space-time constraints limit theoretically which agents can participate in a collaboration requiring a physical encounter. In addition, communication bandwidth and energy resources (in the worst case agents run out of power) are often restricted. Therefore, an efficient management of the collaboration will consider these constraints in limiting the communication to those agents that are relevant for a given task, thus saving communication bandwidth, battery power of mobile agents, and unnecessary message processing overhead of agents out of collaboration range. For this purpose we propose a novel framework based on the *decentralization of time geography*.

Time geography provides a means for accessibility analysis by considering spatio-temporal constraints. However, classical time-geographic analysis [1, 2] is centralized with complete global knowledge and direct communication between the center and individual agents. An emergency management center would identify agents nearby the spill that can reach the site in time, and contact only those. In contrast, decentralized time-geographic analysis operates on limited local knowledge of the mobile social agents. For example, the agent initiating the chemical spill response collaboration does not know where other agents currently are. Our hypothesis is that (1) time geography can be used locally and decentralized to optimize search and (2) its local application makes communications and computations more efficient compared to centralized problem-solving from a global perspective.

Section 2 reviews related work from time geography and decentralized cooperation. Section 3 develops the decentralized time-geographic framework that agents can use to cooperate locally about message spreading and collaboration planning. Section 4 presents an implementation of the framework in a multi-agent simulation, and Section 5 discusses the results of the simulation experiments. The paper concludes with a summary and open questions.

2 Related Work

In the following we introduce the major ideas behind time geography, provide an overview of peer-to-peer communication and agent collaboration, and describe how these areas relate to our research.

2.1 Time Geography

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 [1]. The possibility of being present at a specific location and time is determined by the agent's ability to trade time for space,

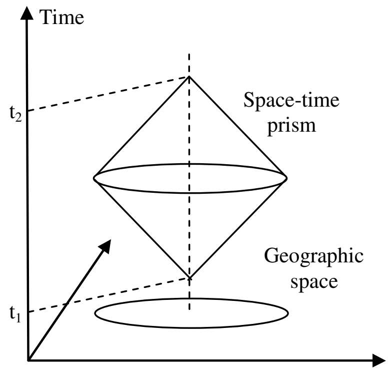

Fig. 1. Space-time prism as intersecting cones

supported by transportation and communication services. *Space-time paths* depict the movement of individual agents in space over time. Such paths are available at various spatial and temporal granularities, and can be represented through different dimensions. All space-time paths must lie within *space-time prisms* (STP). These are geometrical constructs of two intersecting cones [3]. Their boundaries limit the possible locations a path can take (Figure 1). The time budget is defined by $\Delta t = t_2 - t_1$ in which an agent can move away from the origin (the apex at t_1), limited only by the maximum travel velocity and the fact that at t_2 it will be at a certain location (the apex at t_2).

Time geography has been applied to model and measure space-time accessibility in transportation networks [4, 5]. Space-time paths of individuals in networks are limited to movement along edges. The geometry of the STP in a network forms an irregular shape because movement is limited and travel velocity may vary for each edge. Algorithms for calculating the *network time prism* (NTP) can be found in [6] and [7].

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 channel and protocol. In a P2P communication network each node is of equal importance. Nodes can take the role of a communication client, receiving messages from, or sending them to other nodes, but they can also provide message forwarding services for the other nodes. P2P communication networks are transient in nature, with nodes entering and leaving the network freely.

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 typically limited bandwidth and 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. Routing strategies have been studied extensively for static and mobile sensor networks [8, 9], and also for social networks [10].

Parameters defining the connectivity and spread of messages in a peer-to-peer communication network are the communication range and the message forwarding policies. In contrast to classical wireless sensor networks [9], which apply synchronized communication and delayed message forwarding, the types of agents considered in this paper are considered to be 'always-on': broadcasting is possible at any time and message forwarding can happen instantaneously. However, instantaneous (unlimited) message forwarding ('flooding') still reaches only nodes that belong to the connected component of the original sender at the time of broadcasting. Therefore, in mobile networks repeated or opportunistic forwarding are strategies to bridge potential gaps of connectivity, e.g., [11].

Peer-to-peer communication between mobile agents has limited resources, especially bandwidth and on-board battery energy. Accordingly, decentralized algorithms try to minimize the required number of messages to be broadcasted, often at the cost of accuracy.

In the context of the current paper, peer-to-peer communication between mobile agents is complemented by an awareness of the agents of their own position. Agents such as mobile social robots, mobile sensor network nodes, or persons with smart mobile devices are increasingly equipped with positioning technology supporting active or passive locomotion and wayfinding. Tracking positions is an essential factor for time geography, also for decentralized time geography.

2.3 Agent Collaboration

Previous research has focused on technical aspects of P2P collaboration [12] and collaborative, spatio-temporal decision support systems [13]. In [14], a game-theoretic model for the canonical problem of spatio-temporal collaboration was presented with the goal of optimizing individual benefits. Bowman and Hexmoor [15] implemented a simplified agent collaboration system, with boxes that had to be pushed into holes, to investigate the effect of the social network topology on agent collaboration. Agents had to decide whether to collaborate or not based on a payoff criterion. A collaborative, hierarchical multi-agent model integrating knowledge-based communication was implemented in [16] for the RoboCupRescue[1] competition. Global task planning is done in an administrant layer and mobile agents can make their own choices to seek advice from the central agent through an autonomous layer. In general, these robot competitions focus on search and rescue applications: robot capabilities in mobility, sensory perception, planning, mapping, and practical operator interfaces, while searching for simulated victims in unstructured environments.

None of this work explicitly considered the decentralized interaction of mobile agents with both spatio-temporal and communication constraints. These constraints play a role in collaboration within geosensor networks. However, in geosensor networks collaboration has been mostly studied between distributed *immobile* nodes [17, 18], and only recently were mobile nodes allowed to track local movement patterns [19]. Instead we focus here on the coordination of collaboration between different *mobile* agents at a static location and within a given time frame, in some sense generalizing the application-specific approach in [7]. Such spatio-temporal accessibility has been one of the core issues in time-geographic analyses.

[1] http://www.robocup.org/

3 A Decentralized Time Geography Framework

The general problem studied in the following is defined as follows: An agent a_1 at location l_1 and time t_0 needs other agents a_2 ... a_i, with $i \geq 2$, for collaboration at location l_j, with l_j not necessarily l_1, and time t_k, $k > 0$. The agents can only communicate in a peer-to-peer manner—there is neither a central coordinating instance, nor any hierarchic or centralized communication available. The agents' communication efforts required to facilitate the collaboration can be optimized by multiple ways, e.g., by strategies of movement or message routing. Here we are primarily interested in the local application of time-geographic concepts and their impact on limiting the communication and solving the given problem.

The aim of this section is to explore the general dimensions of this problem. These dimensions will form a framework (a set of variables) of decentralized time geography that allows (a) the initiating agent a_1 to specify a particular communication need, and (b) to communicate this need together with the messages to other agents a_2 ... a_i. In this sense, the framework will form an ontology, or a shared vocabulary between the agents, of decentralized time geography.

3.1 Location of the Collaboration

In the above problem, most likely an agent a_1 will initiate a collaboration at its current location l_1, but in principle the location of the collaboration may be different from l_1. From a time-geographic perspective, two other cases can be distinguished. An agent a_1 can invite to a location l_j that is reachable for a_1 before t_k, i.e., a_1 can participate in the collaboration. Alternatively, agent a_1 can invite to a location $l_j \neq l_1$ that is not reachable by a_1 before t_k, i.e., a_1 calls for a collaboration it cannot participate in. As shown in Figure 2, the location l_j and time t_k of the collaboration define a cone, with (l_j, t_k) forming the apex, and the aperture defined by a maximum velocity. Since a_1 knows its maximum velocity, i.e., knows the cone, it can compute the base b_0 of the cone at current time t_0, specifying one out of the following three alternatives:

1. a_1 is at the center of the base of the cone,
2. a_1 is within the base but not at the center, or
3. a_1 is outside of the base.

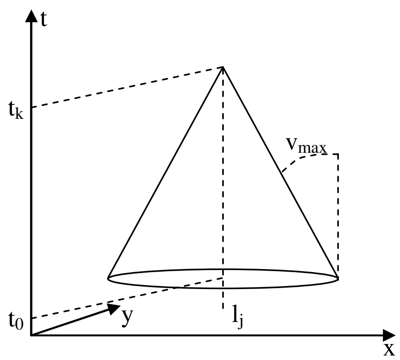

Fig. 2. The cone defined by the location of the collaboration l_j, the time of the collaboration t_k, the current time t_0, and the maximum velocity of agents, v_{max}

The location l_j may even be an extended region instead of a position, $l_j = \int p_k \, dx$. For example, in the introductory scenario a_1 is calling to the spread chemical spill. Such an extension does not change the following principles; it only changes the form of a cone (Figure 2) to a frustum.

Agent a_1 will communicate the cone parameters together with its call for help, such that every agent receiving the call can locally make use of it in two ways: first, it can decide whether it can help, and second, it can decide whether it should forward the message.

3.2 Agent Travel Capabilities

For any agent in the environment confronted with the problem of reconstructing the cone, three of the four parameters are fixed: l_j and t_k are specified in the request message, and the current time t_i, with $0 \leq i \leq k$, can be observed from an on-board clock. For the last one, v_{max}, however, three cases must be distinguished:

1. The agents in the environment are homogeneous. In this case, any agent knowing its own capabilities can safely assume the same capabilities for all other agents. Hence, v_{max} is both constant and globally known.
2. The agents in the environment are heterogeneous, but the environment itself constrains the maximum velocity, e.g., by a traffic code. Such an external behavioral code makes v_{max} again constant and globally known. Still, agents in the environment can have v_{max} below the velocity limit. By including them in the set of relevant agents a communication strategy accepts commissions: these agents may not be able to make it to the meeting point in time.
3. The agents in the environment are heterogeneous, such that an individual agent cannot determine the maximum velocity of any other agent from local knowledge. In this case, an agent can only make an estimate of v_{max}, leading to two potential types of errors:
   - Omissions—the agent's analysis will miss some candidates that have higher velocities than the assumed maximum velocity and are currently outside of the calculated base, but would be inside of the base when applying their true v_{max}.
   - Commissions—the agent's analysis will include some candidates that have lower maximum velocities than the assumed maximum velocity and are currently inside of the calculated space-time prism but would be outside of the base when applying their true v_{max}.

If movements in the environments are restricted to a travel network, the same distinction applies, although the cone becomes discrete and distances are computed on the network [7].

3.3 Message Distribution

Depending on the three location alternatives of agent a_1 in relation to the base (Section 3.1), a_1 may wish to choose different message distribution strategies. The request for collaboration should, according to time geography, be sent to all agents within the

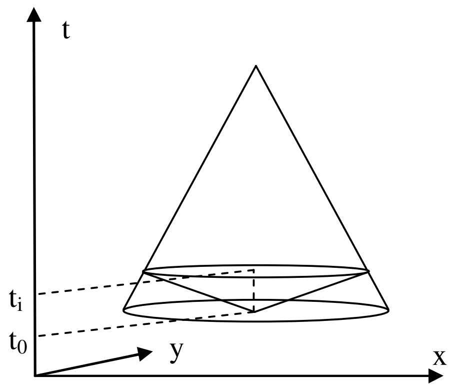

Fig. 3. A call distributed between t_0 and t_i

base b_0 of the cone[2] (l_j, t_k, v_{max}) at t_0. If agent a_1 is inside the base a limited flooding strategy can be applied: all agents receiving the request message, forward (re-broadcast) this message if their radio range intersects with the base of the cone, i.e., if reasonable hope exists to reach relevant agents (agents in b_0). If agent a_1 is outside the base it may wish to geographically route the message first to agents within the base [20] before flooding starts.

The types of agents considered here forward messages instantaneously (in contrast to some geosensor networks that delay each message forwarding by periodical, synchronized communication windows). But this alone does not guarantee that all agents within the base receive the request at t_0. It only guarantees that all agents that are at t_0 also within a_1's connected component of the communication network receive the request at t_0. Since no agent can know the coverage of its own connected component from local knowledge, a_1 can only exploit the dynamically changing connectivity of all agents in its connected component at t_0 by asking these agents to repeat the broadcasting from time to time, $t_i < t_k$, as long as their own radio range intersects with the base b_i of the cone at t_i (the shrinking cone base over time, see Figure 3, can be computed locally). Repeated forwarding increases the probability to reach other relevant agents, although it cannot guarantee reaching all agents that were relevant at t_0.

3.4 Categories of Communication Needs

The type of collaboration may require different communication patterns.

1. In the simplest case, a_1 may only *call* for collaboration, not expecting any other response than agents deciding and physically coming to help. Agent a_1 can construct a cone and determine its base b_0 at t_0 to identify the location of all relevant agents. Then a_1 sends the request for collaboration attaching the cone parameters to the message to enable the recipients of the message to decide whether they can contribute to solving the task.

[2] Note though that the computation of the cone base is an (optimistic) approximation for a space cluttered with features. Real travel distances may be longer than straight lines, but the agents do not know better from local information.

2. Agent a_1 may further specify to receive *offers* for help. Collecting offers might be useful to count and assess the help to be expected as soon as possible, such that if necessary additional measures can be taken (such as relaxing the cone parameters, in particular t_k). This requires suited and prepared agents to respond to a received request, which, if again forwarded repeatedly, adds an additional cone to Figure 3. Since a_1, the target of an offer, is moving, an offer must be sent to the base of a_1's own space-time cone of (l_1, t_0, v_{max}) at the current time. Note that agent a_1 cannot decide whether it has received all offers. Two strategies can be applied: (a) a_1 stops listening after a reasonable time, or (b) a_1 specified in the call for help a time t_i by which it will stop accepting further offers.
3. Alternatively, agent a_1 may have specified that no action should be taken without a specific *booking* of collaboration services. Agent a_1 may want to select a specific number of collaborating agents from incoming offers, or the set of collaborating agents best suited for the task at hand. This third type of message, a booking message, is again addressed to specific but mobile agents, and forms a third cone if forwarded repeatedly. Fragile connectivity means again that booking messages may fail to reach an agent.

3.5 Broadcasting Range

A last distinction is made regarding the communication radius, which is technically a radio range.

1. The radio range can be fixed and the same for all agents, e.g., given by a particular wireless technology.
2. The radio range can be fixed for every agent, but different between individual agents.
3. The radio range can be variable. This case allows broadcasting agents to vary the radio range according to the urgency of the collaboration or the size of the base of the collaboration cone.

4 Simulation

This section presents the implementation of the decentralized time-geographic framework and the specific experiments to evaluate the impact of different parameters on the collaborative task performance. The focus of the current implementation is on the use of the framework for collaborative planning and messages are only forwarded to seek replacements.

4.1 Implementation and Data

NetLogo[3], a cross-platform multi-agent modeling environment, was used for the simulation. It is a programmable environment for simulating natural and social phenomena, and is well suited to model complex systems developing over time. The goal of the following experiments was to demonstrate that time geography can be utilized

[3] http://ccl.northwestern.edu/netlogo/

locally and in a decentralized framework to optimize search, and that this leads to an enhanced efficiency regarding communications and computations compared to centralized problem solving.

In the simulation, one initiating agent a_1 broadcasts a message with a request for help consisting of the task location l_j, upper time limit t_k, and the distribution of different agent types needed. In our case, the simulated environment consists of three different agent types (orange, lime, and magenta) representing different sensing or other capabilities. Every agent that receives the message checks whether it can contribute (i.e., right type of sensor) and reach the location within the specified time limit (i.e., checking whether it is currently located within the space-time cone—compared to the general framework described in Section 3, this only shifts the decision about temporal usefulness from the initiating agent to each individual agent). In addition, there exists a replacement mechanism based on directed messages and 2-way communication within the communication radius, which allows an agent that finds another agent with the same capabilities but closer to the goal to be replaced. For reasons of better comparability the environment is kept constant for all experiments: it

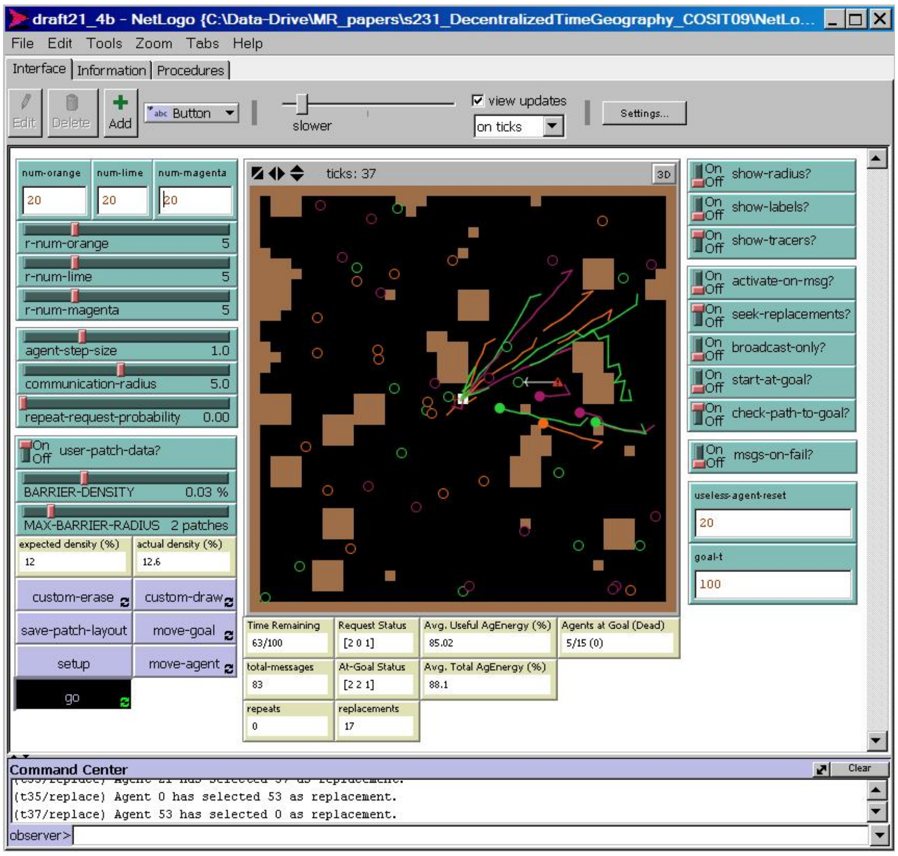

Fig. 4. Screenshot of the NetLogo simulation environment

consists of a 41 x 41 raster grid with a constant barrier structure and density, and the task location l_j is in the center. Furthermore, all agents move with the same velocity. Figure 4 shows the NetLogo simulation environment.

4.2 Experiments

In addition to the general questions of decentralized time geography and its effect on communication and computation, we were particularly interested in the impact of certain variables on solving the given problem. More specifically, we investigated the influence of the following situations:

1. All agents wander randomly in the environment at all times (*AOM false*) vs. agents do not wander without having received a call for help (*AOM true*).
2. Agents seek replacements (*SKR true*) vs. agents do not seek replacements (*SKR false*).
3. Initiating agent a_1 starts at the task location l_j (*SAG true*) vs. initiating agent's location is random within the environment (*SAG false*).
4. Initiating agent a_1 is not at the task location l_j but checks path to the goal first for potential helping agents (*CGP true*) vs. initiating agent wanders randomly in the environment looking for helping agents (*CGP false*).

The experiments were performed for two different numbers of agents (60, 12) randomly placed for each run, with an equal distribution of the three different agent types (20/20/20, 4/4/4), and three different communication radii (5%, 10%, and 20% of total field coverage of the environment). 25% of agents of each type were required at the goal for solving the task, i.e., 5 and 1 of each respectively. Energy consumption was held constant for different activities, i.e., 1.5 to broadcast, 0.2 to seek message, 0.5 to respond, 0.05 to update, and 0.75 to move; the maximum energy was 200 per agent[4]. Every experiment consisted of 100 runs and the individual results were averaged.

4.2.1 Control Conditions

The different behavior mechanisms were compared against control conditions for every combination of agent number and communication radius. This resulted in 6 controls and for each of them the following variables were kept constant: *AOM false, SKR true, SAG false, CGP false.*

4.2.2 Test Conditions

For each test condition we recorded the following variables:

- *Total success rate*: the number of times out of 100 that the simulation was successful, i.e., all agents required to solve the task existed simultaneously at the task location before the upper time limit.
- *Broadcaster success rate*: the number of times out of 100 that the initiating broadcaster succeeded in completing its request, i.e., all agents required started moving towards the task location.
- *Time*: the average time (in ticks) of completion for the successful runs (excludes unsuccessful runs).

[4] These values are arbitrary and can be changed at any time to reflect actual values for a real-time scenario.

- *Total number of messages*: the average number of total messages over all successful runs (excludes unsuccessful runs).
- *Replacements*: the average number of replacement actions per run.
- *Average remaining energy*: the average remaining energy (in %) of all agents.
- *Average used energy*: the average amount of energy used (in units) by all agents. The individual results are given in Table 1 and Table 2.

Table 1. Results of simulation runs for 60 agents and communication radii 5%, 10%, and 20%. Best and worst values (if applicable) within a set of runs are indicated by $^+$ and $^-$ respectively, overall best values for each variable are marked bold.

Scenario (# agents, comm. radius)	Succ.	BCast Succ.	Time	Mess.	Repl.	Remain Energy	Used Energy
control60_10pct	*77*	*91*	*60*	*199.0*	*38.0*	*78.0*	*2639*
E1_60_10_AOMtrue	51$^-$	76$^-$	67$^-$	137.0	18.0	87.0$^+$	1557$^+$
E2_60_10_SKRfalse	64	88	61	57.7$^+$	0.0	81.2	2251
E3_60_10_SAGtrue	79$^+$	93$^+$	55$^+$	160.6	24.1	80.2	2373
E4_60_10_CGPtrue	66	87	64	204.7$^-$	41.0	76.9$^-$	2776$^-$
control60_5pct	*45*	*64*	*72*$^+$	*178.2*$^-$	*43.0*	*73.4*$^-$	*3186*$^-$
E5_60_5_AOMtrue	30$^-$	46$^-$	77$^-$	95.3	13.0	**87.1**$^+$	**1550**$^+$
E6_60_5_SKRfalse	46	68$^+$	75	**57.5**$^+$	0.0	77.9	2650
E7_60_5_SAGtrue	55$^+$	66	72$^+$	155.8	33.6	74.5	3056
E8_60_5_CGPtrue	52	68$^+$	74	174.8	40.5	73.7	3155
control60_20pct	***83***$^+$	*98*	*49*	*248.8*$^-$	*34.5*	*81.7*$^-$	*2192*$^-$
E9_60_20_AOMtrue	58$^-$	95$^-$	57$^-$	172.3	18.8	86.9$^+$	1568$^+$
E10_60_20_SKRfalse	71	96	54	59.7$^+$	0.0	83.7	1957
E11_60_20_SAGtrue	**83**$^+$	**100**$^+$	**38**$^+$	135.0	12.2	86.0	1683
E12_60_20_CGPtrue	75	97	48	238.2	29.2	82.0	2160

Table 2. Results of simulation runs for 12 agents and communication radii 5%, 10%, and 20%. Best and worst values (if applicable) within a set of runs are indicated by $^+$ and $^-$ respectively, overall best values for each variable are marked bold.

Scenario (# agents, comm. radius)	Succ.	BCast Succ.	Time	Mess.	Repl.	Remain Energy	Used Energy
control12_10pct	*65*	*75*	*57*$^-$	*20.0*	*1.6*	*79.4*$^-$	*494*$^-$
E13_12_10_AOMtrue	58$^-$	69$^-$	51	16.4	0.7	90.4$^+$	230$^+$
E14_12_10_SKRfalse	73$^+$	77	52	13.3$^+$	0.0	83.3	400
E15_12_10_SAGtrue	72	83$^+$	48$^+$	19.7	1.3	82.4	424
E16_12_10_CGPtrue	64	76	57$^-$	22.4$^-$	1.9	79.4$^-$	494$^-$
control12_5pct	*56*$^+$	*63*$^+$	*61*$^-$	*15.5*	*1.3*	*77.9*$^-$	*531*$^-$
E17_12_5_AOMtrue	41$^-$	51$^-$	61$^-$	12.8	0.6	89.1$^+$	262$^+$
E18_12_5_SKRfalse	54	62	55$^+$	**12.0**$^+$	0.0	80.5	468
E19_12_5_SAGtrue	52	58	57	16.2$^-$	1.4	78.1	526
E20_12_5_CGPtrue	52	60	55$^+$	15.3	1.2	78.2	523
control12_20pct	*74*	*87*	*48*	*28.3*$^-$	*2.0*	*82.1*$^-$	*430*$^-$
E21_12_20_AOMtrue	72$^-$	78$^-$	43	21.8	0.8	**91.3**$^+$	**208**$^+$
E22_12_20_SKRfalse	81	90	51$^-$	14.2$^+$	0.0	84.8	366
E23_12_20_SAGtrue	**90**$^+$	**97**$^+$	**39**$^+$	22.6	0.7	86.7	319
E24_12_20_CGPtrue	77	89	44	27.7	2.3	84.0	383

5 Analysis and Discussion of Simulation Results

The implementation of the decentralized time-geographic framework demonstrates that time geography can be used locally and optimizes search by selecting only relevant agents, i.e., those agents that can contribute to solving the task and are able to reach the task location within the specified time limit. Furthermore, the replacement mechanism enhances temporal efficiency by substituting helpful agents with other helpful agents of the same type but closer to the goal.

In order to compare our decentralized framework to one with a central instance we calculated an optimal centralized case where the central command has complete knowledge, therefore resulting in the highest success probability. It is based on the assumption that the communication radius of the central instance covers the whole environment and can therefore reach every agent with 1 message. In reality, assuming a flooding strategy utilizing several hops between agents, these numbers will be much higher, and in addition it cannot be guaranteed that every agent will receive the message due to non-connected parts of the network. The total message numbers for the centralized cases are therefore derived by 1 (initiating agent to central instance) + 2 (send from / reply to central instance) * number of agents (60 / 12) + number of requested agents (15 / 3) = 136 and 28 messages respectively. Note that in the centralized case the central instance needs to evaluate whether an agent will make it to the task location on time, whereas in the decentralized framework agents determine this themselves. The results demonstrate that even when compared to the best centralized case, in more than half of the experiments, decentralized time geography leads to a higher communication efficiency (and this will only improve further when compared to flooding strategies). There is a clear dependency between number of agents in the environment and communication efficiency when comparing centralized and decentralized strategies: the lower the agent density the higher the efficiency gain in the decentralized framework. In the case of no replacements, the decentralized time-geographic framework is always superior and there are even scenarios for a large number of agents when this is the case (*60_5_AOMtrue*, *60_20_SAGtrue*).

In the following we discuss the overall results for the control conditions and the individual results for the four test scenarios.

5.1 Analysis of Control Scenarios

To demonstrate the overall picture of the decentralized time-geographic framework, we investigated for each control condition the results for the following variables: total success rate, time of completion, number of messages, number of replacements, and remaining energy. Due to the different units the original values were standardized by calculating their z-values, i.e., the standard deviation functions as the unit of measurement for describing the distances from the mean [21].

The visualization in Figure 5 shows that the total success rate is positively correlated with the length of communication radius. This is independent of the total number of agents. It is important to note that a larger number of agents does not automatically lead to a higher success rate as the *60_5* and *12_5* conditions demonstrate. A significant result is the fact that success rate is inversely related to time: for high success rates the average time of completion for the successful runs is a

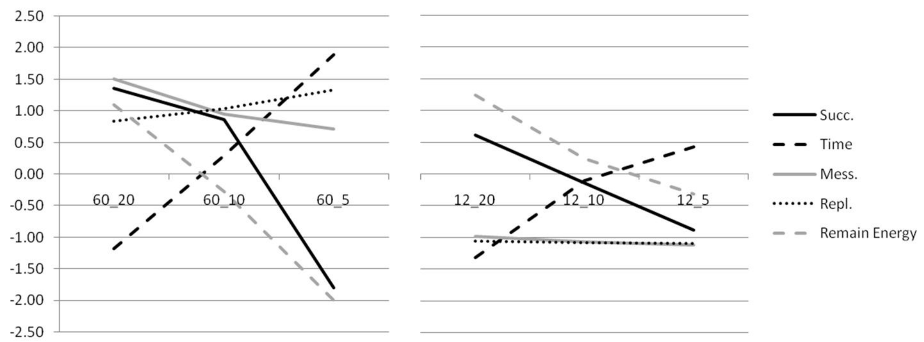

Fig. 5. Comparison of z-transformed values for total success rate, time of completion, number of messages, number of replacements, and remaining energy for the 6 control conditions (numberOfAgents_communicationRadius)

minimum. As expected, more agents in the environment lead to more messages being exchanged and the smaller the communication radius the fewer the messages. The number of replacements stays relatively constant within the 60 and 12 agents conditions and there seem to be no obvious significant overall correlations to any of the other variables. This is surprising because we expected that a large number of replacements will result in more efficiency (in terms of saving time). The most significant counter-example for this is condition *60_5*, which has the largest number of replacements but also uses the most amount of time. Remaining energy correlates positively with success rate and negatively with time, leading to the assertion that high success rates are due to energy-efficient problem-solving.

5.2 Agent Movement

The results clearly demonstrate that (random) agent movement has a large impact on the total and broadcaster success rates, and also on the amount of time needed to reach the task location. Both the total success rate and the broadcaster success rate are lower for all of the scenarios where agents do not wander in the environment without having the message. Except for the *12_10* and *12_20* conditions, such behavior also takes the most time for solving the task. There are fewer messages being exchanged and fewer replacements occur compared to the control conditions. Due to the fact that most of the agents do not move most of the time, the remaining energy is highest for all scenarios.

5.3 Agent Replacements

In the conditions without replacements one could argue both ways, i.e., for a higher success rate because fewer agents are expected to die due to energy loss, and on the other hand for a lower success rate because agent replacements save time (the replacing agent is closer to the goal) resulting in fewer cases where total time runs out. The results indeed demonstrate a mixed effect: On average the conditions for 12 agents show a higher success rate, whereas the conditions for 60 agents show a lower average success rate. The differences to the control conditions are also smaller for 12 agents. The reason for this is that fewer agents imply fewer replacement possibilities and therefore a smaller visible effect. This effect can also be seen regarding time: All scenarios with 60 agents need more time for solving the task without replacements. As expected, without replacements there are fewer messages being sent and these savings in communication costs lead to savings in energy.

5.4 Location of Initiating Agent

Locating the initiating broadcaster at the task location has a significant effect on the outcome of the simulations. Except for the *12_5* condition, the total and broadcaster success rates are on average (*60_20* is a tie for total success rate) significantly higher than for the control conditions. This is not surprising because the broadcaster immediately gets a portion of the requested agents to the goal within a short period of time. A small number of agents and a small communication radius lead to fewer useful agents around the goal and this explains the outlier *12_5*. The condition of having the initiating agent at the goal location also results in the fastest task performance overall (see bold numbers in Table 1 and Table 2), again with *12_5* being the exception. As expected there are fewer total messages (for all but *12_5*) and on average agents use less energy because there is lower cost for travel. Note that for 60 agents there is a significant reduction in replacements because the first portion of useful agents is closest to the goal and therefore does not need replacements. For 12 agents this effect is barely visible.

5.5 Goal-Directed Behavior of Initiating Agent

In this condition, the initiating agent broadcasts the message while moving to the goal itself. While we expected an increase in success rate compared to the control condition, where the initiating agent moves randomly in the environment, the results are mixed. Only the *60_5* and *12_20* conditions result in higher success rates whereas in the other 4 conditions success rates decrease. One possible explanation for this is the extra way of going back from the goal when some of the required agents are still missing. An interesting case is condition *12_10* where the total success rate is lower although the broadcaster success rate is slightly higher. For all other cases total success rate and broadcaster success rate have a positive correlation. Another surprising result is that compared to the control conditions, more time is needed on average for 60 agents to solve the task, whereas for 12 agents this result is reversed. The number of total messages is similar to the control conditions, slightly fewer on average but more for 10% communication radii. There are also only minor differences with regard to remaining energy and number of replacements.

Overall, the simulation results indicate that a high success rate and using as little time as possible, can be achieved with a large communication radius and the initiating agent located at the goal. Scenarios with static agents are very energy-efficient but show the worst results regarding success rates and time. Minimizing the number of messages can be achieved in scenarios without replacement mechanisms but this comes at the cost of lower (though not worst) success rates.

6 Conclusions and Future Work

In this paper we have developed a theoretical framework for *decentralized time geography*. Traditional time-geographic analysis is centralized and assumes complete global knowledge, whereas the move to decentralized time geography allows for the consideration of limited local knowledge. Such framework is important for networks with peer-to-peer collaboration of mobile social agents, where there is no central instance and time-geographic analyses must be performed locally by the distributed agents.

In particular, we investigated the problem of one agent initiating ad-hoc collaborative planning and decision-making among several agents, which eventually leads to physical support at a specific site. We identified as the major components of the decentralized time-geographic framework the location of the collaboration, the agents' travel capabilities, the message distribution strategy, the communication pattern, and the broadcasting range. In order to demonstrate the functioning of decentralized time geography, we performed experiments in a multi-agent simulation framework for different combinations of agent density and communication radius. The results demonstrated that time geography can be used locally and decentralized to optimize search. For more than half of our experiments its local application made communications and computations more efficient compared to problem-solving from a global perspective, i.e., with a central command having complete knowledge of the environment and its agents. We further showed the impact of specific conditions with regard to agent movement, replacement seeking, and location and goal-directed behavior of the initiating agent on the simulation outcomes.

The framework of decentralized time geography has many potential application areas, such as physical robot interaction, search and rescue, environmental analyses, transportation problems, and mobile location-based gaming. Our work serves as a foundation for the theory of decentralized time geography and leads to many directions for future research:

- The theoretical framework developed here needs to be thoroughly tested through extended simulations covering various application domains and additional aspects such as heterogeneous agents with different capabilities and maximum velocities, and different message distribution strategies, such as geographical routing and flooding. A major question that arises for these distribution strategies concerns the communication of the current request state to other agents in the network, i.e., how many agents of different types are still needed at a given point in time.
- In addition to evaluating whether they make it to the task location on time, the spatio-temporal constraints of time geography should enable agents to plan their own communication behavior, such as, in a given situation, deciding how far messages should travel, or where or how often messages should be spread.
- Limited local knowledge, especially about dynamically changing connectivity, loses some agents but this is compensated for by lower energy consumption. It is important to investigate the correlation between such lower consumption and total success rates, as well as how repetitive forwarding strategies may alleviate the problem of disconnected network components.

- In our simulation we have kept the environment fixed, but variations of environmental structure, barrier density, and task location will most likely have a large impact on the outcome of the experiments. Future work should utilize representations of real-world environments, in order to better investigate the impact of such factors, also with regard to choosing optimal spatial search and travel strategies. This may include the possibility of agents joining and leaving a network at any given point in time.
- Extending our scenario to situations with multiple collaborations at different locations or times will help in generalizing the proposed framework for decentralized time geography.

Acknowledgments

The comments from three anonymous reviewers provided useful suggestions to improve the content and clarity of the paper. This work is supported under the Australian Research Council's Discovery Projects funding scheme (project number 0878119).

References

1. Hägerstrand, T.: What about people in regional science? Papers of the Regional Science Association 24, 7–21 (1970)
2. Miller, H.: A Measurement Theory for Time Geography. Geographical Analysis 37(1), 17–45 (2005)
3. Lenntorp, B.: Paths in Space-Time Environments: A Time-Geographic Study of the Movement Possibilities of Individuals. Lund Studies in Geography, Series B (44) (1976)
4. Miller, H.: Measuring space-time accessibility benefits within transportation networks: Basic theory and computational methods. Geographical Analysis 31(2), 187–212 (1999)
5. Kwan, M.-P.: GIS Methods in Time-Geographic Research: Geocomputation and Geovisualization of Human Activity Patterns. Geografiska Annaler B 86(4), 267–280 (2004)
6. Miller, H.: Modeling accessibility using space-time prism concepts within geographical information systems. International Journal of Geographical Information Systems 5(3), 287–301 (1991)
7. Raubal, M., Winter, S., Teßmann, S., Gaisbauer, C.: Time geography for ad-hoc shared-ride trip planning in mobile geosensor networks. ISPRS Journal of Photogrammetry and Remote Sensing 62(5), 366–381 (2007)
8. Nittel, S., Duckham, M., Kulik, L.: Information dissemination in mobile ad-hoc geosensor networks. In: Egenhofer, M., Freksa, C., Miller, H. (eds.) Geographic Information Science - Third International Conference, GIScience, pp. 206–222. Springer, Berlin (2004)
9. Zhao, F., Guibas, L.: Wireless Sensor Networks. Elsevier, Amsterdam (2004)
10. Marti, S., Ganesan, P., Garcia-Molina, H.: SPROUT: P2P Routing with Social Networks. In: Lindner, W., et al. (eds.) Current Trends in Database Technology - EDBT, Workshops, pp. 425–435. Springer, Berlin (2004)
11. Xu, B., Ouksel, A., Wolfson, O.: Opportunistic Resource Exchange in Inter-Vehicle Ad-Hoc Networks. In: Fifth IEEE International Conference on Mobile Data Management, pp. 4–12. IEEE, Berkeley (2004)

12. Schoder, D., Fischbach, K., Schmitt, C.: Core Concepts in Peer-to-Peer Networking. In: Subramanian, R., Goodman, B. (eds.) Peer-to-Peer Computing: The Evolution of a Disruptive Technology, pp. 1–27. Idea Group Inc., Hershey (2005)
13. Jankowski, P., Robischon, S., Tuthill, D., Nyerges, T., Ramsey, K.: Design Considerations and Evaluation of a Collaborative, Spatio-Temporal Decision Support System. Transactions in GIS 10(3), 335–354 (2006)
14. Luo, Y., Bölöni, L.: Children in the Forest: Towards a Canonical Problem of Spatio-Temporal Collaboration. In: AAMAS 2007, Int. Conference on Autonomous Agents and Multiagent Systems, pp. 990–997. IFAAMAS, Honolulu (2007)
15. Bowman, R., Hexmoor, H.: Agent Collaboration and Social Networks. In: International Conference on Integration of Knowledge Intensive Multi-Agent Systems KIMAS 2005: Modeling, EVOLUTION and Engineering, pp. 211–214. IEEE, Waltham (2005)
16. Peng, J., Wu, M., Zhang, X., Xie, Y., Jiang, F., Liu, Y.: A Collaborative Multi-Agent Model with Knowledge-Based Communication for the RoboCupRescue Simulation. In: International Symposium on Collaborative Technologies and Systems (CTS 2006), pp. 341–348 (2006)
17. Duckham, M., Nittel, S., Worboys, M.: Monitoring Dynamic Spatial Fields Using Responsive Geosensor Networks. In: Shahabi, C., Boucelma, O. (eds.) ACM GIS 2005, pp. 51–60. ACM Press, New York (2005)
18. Farah, C., Zhong, C., Worboys, M., Nittel, S.: Detecting Topological Change Using a Wireless Sensor Network. In: Cova, T.J., Miller, H.J., Beard, K., Frank, A.U., Goodchild, M.F. (eds.) GIScience 2008. LNCS, vol. 5266, pp. 55–69. Springer, Heidelberg (2008)
19. Laube, P., Duckham, M., Wolle, T.: Decentralized Movement Pattern Detection amongst Mobile Geosensor Nodes. In: Cova, T., et al. (eds.) GIScience 2008. LNCS, vol. 5266, pp. 199–216. Springer, Heidelberg (2008)
20. Ko, Y.-B., Vaidya, N.: Location-aided routing (LAR) in mobile ad hoc networks. In: Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 1998), pp. 66–75. ACM Press, Dallas (1998)
21. Devore, J., Peck, R.: Statistics - The Exploration and Analysis of Data, 4th edn. Duxbury, Pacific Grove (2001)