International Journal of Geographical Information Science
Vol. 00, No. 00, July 2006, 1–21

Ad-hoc shared-ride trip planning by mobile geosensor networks

Stephan Winter
Department of Geomatics, The University of Melbourne, VIC 3010, Australia
winter@unimelb.edu.au
Silvia Nittel
Department for Spatial Information Science and Engineering, University of Maine, USA
nittel@spatial.maine.edu

(Received April 14, 2006)

Recent developments in miniaturization of computing devices, in location-sensing technology and in ubiquitous short-range wireless networks enable new types of social behaviour. This paper investigates one novel application of these technologies, ad-hoc inner-urban shared-ride trip planning: Transportation clients such as pedestrians are seeking ad-hoc shared rides from transportation hosts such as private automobiles, buses, taxi cabs or trains. While centralized trip planners are challenged by assigning clients and hosts in an ad-hoc manner, in particular for non-scheduled hosts, we consider the transportation network as a mobile geosensor network of agents that interact locally by short-range communication and heuristic wayfinding strategies. This approach is not only fully scalable; we can also demonstrate that with short-range communication, and hence, incomplete transportation network knowledge a system still can deliver near-to-optimal trips.

Keywords: route planning, incomplete knowledge, mobile geosensor networks, negotiation strategies, heuristic wayfinding strategies.

1 Introduction

Urban mobility can be greatly enhanced by concepts of ride sharing. Wherever ride sharing has evolved, the process was driven by social conventions more than by technological progress (Resnick 2004), avoiding to tackle the complexities of ad-hoc shared-ride trip planning and assignment. However, with the capabilities of today's technology of small-form, handheld computing devices, location sensing and ubiquitous wireless communication networks—combined to *ad-hoc mobile geosensor networks* (Stefanidis and Nittel 2005)—, new types of multimodal, real-time trip planning and booking systems become possible. We envision a system that integrates the transportation capacities of all types of (volunteering) vehicles in urban traffic in order to identify a trip for persons with an ad-hoc travel demand. The system shall assign persons, or *transportation clients*, to vehicles, or *transportation hosts*, with matching travel plans and free transportation capacities, in an ad-hoc manner.

In this paper, we are looking into the complexities of the trip planning task of an ad-hoc shared-ride system. This task is challenging the current state of knowledge since it copes with spatially and temporally incomplete transportation network knowledge. In principle, ad-hoc shared-ride trip planning requires complete transportation network knowledge for finding optimal solutions. But realistically it is impossible to obtain complete knowledge of all vehicles that are currently and in near future in urban traffic, their travel plans and their current and future utilization. This system is dynamic and non-predictable. This means ad-hoc shared-ride trip planning has to happen with incomplete knowledge regarding the future states of the transportation system. It even has to cope with incomplete knowledge in its current state, given the complexities of tracking large numbers of individual vehicles in real-time. Hence, ad-hoc shared-ride trip planning can only come up with sub-optimal trips.

At the same time, the envisioned ad-hoc shared-ride system shall serve large numbers of concurrent clients, which is also different from traditional shared-ride systems. To cope with scalability we propose a distributed system of autonomous agents solving trip planning *locally*. Technically, this system is an ad-hoc mobile geosensor network, with nodes of transportation clients and hosts that are capable of self-positioning and ad-hoc radio-based peer-to-peer communication. In this system the clients will collect data about the current transportation network, plan a trip, and select hosts. Since communication in a geosensor network is expensive and needs to be minimized for several reasons (e.g., bandwidth, time delay, and potentially battery energy), we go a radical step further and limit the trip planner's knowledge deliberately in the spatial dimension, by contacting only hosts nearby. In this case the research question is whether spatially and temporally limited transportation network knowledge still enables acceptable trips for trip planning clients.

The hypothesis of this paper is that mobile geosensor networks are an *effective* and *efficient* approach to ad-hoc shared-ride trip planning. In this hypothesis we call mobile geosensor networks *effective* if they come up with trips close to the optimal trip according to a chosen cost function. We call them *efficient* if the communication effort for an effective trip in terms of numbers of broadcasted messages in negotiations was significantly lower than for collecting exhaustive transportation network knowledge. We will collect evidence for the hypothesis in three steps: we will show (i) that current optimal trips can be found in geosensor networks, (ii) that trips can be generated with local knowledge only, and (iii) that trip quality and negotiation effort can be balanced by choosing a spatially limited negotiation strategy.

We show these properties by simulation. For this purpose we develop a two-way trip negotiation process, and investigate the implications of different spatial ranges of this process. We develop a protocol that directs messages of the two-way negotiation, and increases its efficiency. Since any trip plan is bound by incomplete knowledge, negotiations be scheduled recursively, such that trip plans can be updated regularly. To study only the implications of negotiation ranges, we choose a fix and simple wayfinding heuristics.

We investigate three different negotiation ranges. First, a spatially unconstrained negotiation range is applied. It yields exhaustive transportation network knowledge of the current situation, and hence, currently optimal shared-ride trips. Next, two spatially limited negotiation ranges are investigated, and their effectiveness and efficiency is assessed by their average shared-ride trip durations, including wait and travel time, and their broadcasting efforts compared to the first strategy. We demonstrate that short negotiation ranges save broadcasting costs and still deliver near-to-optimal trips.

The paper is structured as follows. Section 2 discusses related systems and technologies. Section 3 describes in detail the problem of shared-ride trip planning in the envisioned system. In Section 4, the necessary components of a mobile geosensor network are explained. In Section 5 we formalize the negotiation process between transportation clients and hosts in geosensor networks for the purpose of simulation. The results of the simulation are discussed in Section 6. The paper closes with a summary and an outlook on open questions in Section 7.

2 The position of shared-ride trip planning using geosensor networks

In this section we introduce trip planning services, shared-ride systems and mobile geosensor networks, and relate them to ad-hoc trip planning and local problem solving.

2.1 *Current trip planning systems*

Current approaches for real-time individual trip planning are based on centralized services. This is the case for current commercial solutions, and also for research approaches (Ziliaskopoulos and Mahmassani 1993, Fu 2001, Dillenburg et al. 2002, Chon et al. 2003). A centralized trip planning system typically consists of a database management system that stores a global view of the transportation network. It keeps track of all changes made by continuously moving agents, and either plans optimal trips for all clients (for example, a real-time train trip planner), or it broadcasts traffic conditions to autonomously planning clients (for example, a car navigation system analysing the traffic message channel). Aggregated traffic information is not sufficient for shared-ride trip planning. – In a continuously and unpredictably changing environment, the centralized database system becomes easily the bottleneck just due to the location updates of agents moving in an unconstrained manner. Furthermore, each change in the network potentially requires updating all trip plans and assignments. On top of this, for a shared-ride system the system has to manage real-time communication with clients and hosts. Since any client is potentially related to any host, complexity grows exponentially with the number of clients, which means the system is not scalable.

Using geosensor networks, trip planning becomes a collaborative task in a distributed network of mobile nodes, with ad-hoc peer-to-peer communication (Zhao and Guibas 2004, Stefanidis and Nittel 2005). In this way, the type and design of the communication between the peers becomes the key to dynamic trip planning. This approach can be fully scalable if every new transportation request can be solved locally in the geosensor network. Applications for mobile sensor networks already envision transportation systems (Zhao and Guibas 2004). Sussman (2000) categorized transportation systems by a schema, according to which ad-hoc shared-ride trip planning in mobile geosensor networks can be characterized by *individual travellers*, *urban transport*, and *private operation*. Nijkamp et al. (1996) identify travel information as one of the major functions of transportation systems.

2.2 *Social parameters of shared-ride systems*

Shared-ride systems enjoy some popularity in defined communities, while public shared-ride systems are currently not popular. One of the reasons is the association with hitchhiking. In some cultures hitchhiking has a negative connotation, but not everywhere. Another reason is the inflexibility of current shared-ride systems with real-time travel needs in a dynamic environment. Shared-ride agencies such as the Mitfahrzentrale[1] or RideNow[2] expect that car drivers as well as passengers register their offerings and needs, respectively, well in advance. Institutionalized commercial ad-hoc shared-ride systems such as SuperShuttle[3] operate only from well-known pick-up points and rely on social conventions such as branding. Route planning is still done by the shuttle drivers and is part of the human-human interaction.

However, the current situation is surprising given the enormous potential for shared-ride systems predicted by traffic managers (Dillenburg et al. 2002) or social scientists (Noda et al. 2004, Resnick 2004). Resnick, for example, names some successful shared-ride systems that are ad-hoc without any technological support and function only by social conventions, for example, by waiting for a ride in queues at well-known pick-up points. A more predictable solution, like the proposed one, has therefore the potential for significant social and economic impact.

Hence shared-ride systems have to consider some implications and challenges prior to any realization. They comprise, for example, trust and safety, liability, economic incentives and business models (McCarthy 2001), urban mobility and access, fair share (Naor 2005), and privacy (Monmonier 2002). A particular concern is the change of a potentially negative public perception of shared-ride travelling, and related to that, a change in the proxemics of the involved social beings (Hall 1966). When we look into trip planning we are aware of all these other aspects, but leave them for further work.

---

[1] Mitfahrzentrale™: `http://www.mitfahrzentrale.de`
[2] RideNow™: `http://www.ridenow.com`
[3] SuperShuttle™: `http://www.supershuttle.com`

2.3 *Mobile geosensor networks*

Geosensor networks are a specific type of sensor networks. A sensor network consists of a large collection of individual small computing platforms (nodes) each of which can be equipped with a variety of micro-sensors and is capable of wireless short-range communication (Zhao and Guibas 2004). A geosensor network has at least one positioning sensor node such as a GPS receiver as part of the overall network, so that all other nodes can derive at least their relative geographic position (Stefanidis and Nittel 2005). In mobile geosensor networks, each node is likely to contain a private location sensing capability. Mobile nodes of a geosensor network collaborate in an ad-hoc, task-oriented fashion. In the literature such network topologies are also called mobile ad-hoc networks or MANETs (e.g., Gerla et al. 2005). Today, an ad-hoc mobile geosensor network can be established using hand-held devices as used by pedestrians and automobile drivers; in the near future, hand-held devices will be replaced with cent-size computing nodes that are embedded in cell phones, watches or car navigation systems.

Using wireless sensor networks, several technical solutions or media for wireless communication are possible; they can be classified into short-range and wide-range wireless communication (Zhao and Guibas 2004). To preserve energy, the RF signal strength is kept low. For our type of geosensor network, short-range wireless communication such as Bluetooth or WiFi is of interest. Each client and host is a radio sender as well as a receiver, and broadcasting is used to generate or forward messages to other agents in the reception area of a sender. Since radio range of these technologies is between $3m$ and $100m$, messages can be re-broadcasted by recipients to reach agents in larger distance (*multi-hop*). However, the decision whether an agent will re-broadcast a message, and to whom, influences the spread of the information in the network and the congestion of the network bandwidth so that an optimal trade-off between both has to be found. Furthermore, to minimize energy consumption broadcasting in wireless sensor networks networks can be synchronized, and thus, it might takes place in relatively short and synchronized *communication windows*; the rest of the time the network nodes turn off the radio to preserve energy. The length and frequency of these communication windows depends on the application needs, but it limits the numbers of messages passing through each node.

Other work concentrates on one-way information dissemination about events in mobile geosensor networks (Nittel et al. 2004, Wolfson and Xu 2004). An initial classification of information dissemination strategies was (Nittel et al. 2004):

(i) flooding: each agent that receives a message about a client request passes on the information repeatedly to every other agent within its radio range. Each receiving agent also passed on the information to any other node in the network.
(ii) epidemic: each agent passes on the information to only the first k other agents it encounters. The receiving agents will proceed similarly.
(iii) location-constrained: requests are re-broadcasted by an agent only within the spa-

tially constrained proximity of the original request, and then no longer passed on.

However, the present problem of trip negotiations needs an appropriate two-way communication between the communication originator and hosts within the chosen negotiation range. This needs to be studied in more detail.

3 The problem definition

In this section we study in detail the ad-hoc shared-ride navigation problem in an unpredictably dynamic transportation network, and compare it to current trip planning systems. For that purpose we also introduce a scenario of ad-hoc shared-ride travelling in an urban environment.

3.1 *A shared-ride planning and assignment system*

Consider the following scenario. Hillary has just missed her bus to work today. Around Hillary is heavy traffic. Now, she is glad to have subscribed to a transportation service that mediates between her current travel needs to her destination, and those buses, trams, taxis and subscribed car drivers who are going in her direction. She switches on her device, which immediately starts to communicate with devices of vehicles close by and starts trip planning and booking. Soon after, Hillary sees a friendly car driver stopping to give her a ride. The ride takes her on the first leg of her trip. During the ride, her device still runs in the background. It looks up in the network for appropriate transfers, and books them for Hillary. Hillary will be on time for work today.

In contrast to current real-time route planning services, Hillary's service has no central communication and planning component. Instead, all negotiations happen directly between Hillary's device and the devices of vehicles close by. In this way, the data for trip planning is always current, but local. Spatial proximity of clients and hosts is dictated by the limited radio range of the devices, where larger ranges can be accomplished by message forwarding.

3.1.1 ***The transportation client agents.*** In the scenario people like Hillary are looking for rides from their current position to a particular destination. We call these travellers, or more precise, their devices, *transportation client agents*, or clients for short, and denote them by C_i. Clients are mobile agents that *sense* their own current location, communicate with near-by agents, *plan* a shared-ride trip, and *act* by taking a ride or moving autonomously. The negotiation between transportation agents encompasses that the clients can broadcast a request (which may be forwarded), collect offers, and book specific transportation hosts.

3.1.2 ***The transportation host agents.*** In the scenario *transportation host agents*, or hosts for short, are the devices of all sorts of vehicles, such as private cars, buses, taxis, or subways. Hosts are denoted as H_j. The travel plans of hosts form the links of a transportation network. These links are spatially bound to the street network, but temporally highly irregular. Some types of hosts follow pre-defined routes, other types have no routes, or can change their routes ad-hoc.

In general, future states of the transportation network cannot be seen though. All vehicles are at least to some extent autonomous. Taxis and car drivers do not follow any schedule, and buses and trams frequently run out of schedule. At any time a new vehicle can enter the traffic and offer rides, current vehicles can get occupied and are temporarily not available—in particular private cars offer a rigorously limited transportation capacity—, and other vehicles reach their destination and withdraw from the network.

3.1.3 ***The shared-ride transportation network.*** In our scenario, traffic is bound to the physical street network. From a network perspective, the dynamic provision of transport along network edges (street segments) forms a time-dependent cost function for these edges. Clients travelling along a street segment have to wait until a host with free capacity comes along.

Not all future transportation opportunities are known at a time t_k. Consequently there is no guarantee for any connected sequence of host segments to the client's destination in any trip planning process, and new knowledge might only emerge over time during travelling. We can assume, however, that every client finds transportation to his/her destination sooner or later.

The relevance of transportation hosts for trip planning decreases with spatial distance from a client. A distant host becomes only relevant if the client does not find a nearer host, i.e., departs sooner anyway. If the distant host is outside of the client's search range, she will wait until suited hosts appear in her range.

In all these aspects the transportation network differs fundamentally from classical multi-modal networks, which are assumed to be scheduled and always connected. In multi-modal networks the time-dependent cost functions are predictable and known in advance. Additional components might consider real-time information on delays and updated schedules (Ziliaskopoulos and Wardell 2000). Other time-dependent route planning algorithms assume static (street) network with dynamically changing weight functions, for example, according to the current traffic situation (Chon et al. 2003). Schedule-based algorithms can be found in the literature (Cooke and Halsey 1966, Klafszky 1972, Peng and Tsou 2003, Orda and Rom 1990). Implemented in central services, for example in Web services for public transportation planning, they rely on comparatively small concurrent user numbers.

3.2 *A shared-ride wayfinding heuristics*

Hillary is looking ad-hoc for a ride in a dynamic transportation network. The system on her device, the client agent, can only gain temporally and spatially limited knowledge of the actual transportation network. Hence the client agent needs wayfinding strategies to deal with the incomplete network knowledge, especially with gaps.

Client agents can select different wayfinding strategies, which vary the navigation result. For example, clients can apply the *least-angle strategy*, choosing from the available hosts the first that goes in their direction, or they can apply a *longest-leg strategy*, looking for the host that brings them closest to their destination (Hochmair and Frank 2002). However, investigating different navigation strategies is beyond the scope of this paper. Instead, we choose one strategy, and focus only on the effects of different communication strategies. Other wayfinding strategies will be affected by different communication strategies in similar ways.

We assume that client agents know the street network, but have limited knowledge of the actual traffic and ride opportunities. In this case, a client agent can choose to stick to the shortest distance route, or one of them if there are several, and look for transport along only this route. This wayfinding strategy is conceptually related to the least-angle strategy: the selected route is the graph geodesic. In contrast to the least-angle strategy, the shortest route strategy is not burdened with the danger of running into dead-ends.

Applying this wayfinding strategy requires no route planning after the initialization. The client's and hosts' devices only need to match sequences of street network edges to find overlaps between demand and supply. This aspect makes the wayfinding strategy computationally cheap. Furthermore, the information needs of the client can be specified straightforward: they concern transportation along the edges of the chosen route. Offers consist of subsets of these edges, attached with time stamps. This means, with this strategy the message lengths are manageable (linear with the length of the route), and the agents' internal main memories are not burdened much, only by strings and pattern matching.

The chosen wayfinding strategy is heuristic, which can lead to suboptimal results: the shortest distance route is not necessarily the fastest overall. However, in this paper we are only interested in the effects of different communication strategies, and compare therefore the trips travelled by the clients with the trips the client would make with exhaustive network knowledge. In contrast, a comparison with the overall fastest route would assess the wayfinding strategy, which can be done in future work.

4 Ad-hoc mobile geosensor networks for trip planning

In this section we consider the acting agents in the trip planning process as nodes in an ad-hoc mobile geosensor network, and we introduce some relevant communication concepts of geosensor networks for shared-ride trip planning.

4.1 *A mobile geosensor network of clients and hosts*

In our context, each geosensor node runs a local agent which is either a transportation client or host. The collaborative task of the geosensor network is shared-ride trip planning with clients becoming their own trip planners. They communicate with nearby hosts to learn about currently available transportation means. They select some of the hosts, book them, and travel with them.

Roussopoulos et al. (2004) have developed criteria to decide whether a problem is a 'peer-to-peer problem'. Referring to these criteria, ad-hoc shared-ride trip planning is clearly a case for peer-to-peer approaches (as realized by mobile geosensor networks):

- low-budget decisions: transportation information is a penny business.
- relevance: local communication in a geosensor network reaches the relevant agents, and directing messages will further reduce any unnecessary communication.
- trust: there is low motivation for giving false transportation information.
- rate of change: the rate of change in a mobile geosensor network is high. While this may be a disadvantage in a distrustful environment, we even argue that in our case the high rate of change is a motivation for a peer-to-peer solution.
- criticality: transportation information is uncritical; if the optimal trip cannot be detected the second optimal will do.

4.2 *Local communication for trip planning in dynamic networks*

To determine an optimal shared-ride trip, a client needs to maintain information about all transportation hosts that are relevant to the planned trip. Given the dynamics in the transportation network it seems promising for clients to achieve partial trip planning with local knowledge, and to update trip plans in intervals in order to achieve an overall optimum for the entire trip.

From a trip planning perspective, the probability is higher that nearby transportation hosts contribute to optimal (fastest) trips, because clients will wait less long for them than for hosts far away (assuming that travel speeds are homogeneous). Also, hosts that reach the client sooner will likely be selected by the client, because the network has not changed much since booking, and there might not be much new evidence for changes in bookings.

From a geosensor network perspective, energy is one of the scarce resources, and the most energy-consuming activity of a node is using the wireless communication medium. Another scarce resource in the network is physical communication bandwidth, which is likely the more relevant bottleneck in this particular problem. For both reasons the number of messages has to be minimized.

Hence, the question arises, by which ways and at which costs (in terms of increasing trip duration) the spread of messages can be focused.

4.3 *Negotiation ranges for ad-hoc shared-ride trip planning*

In our scenario all transportation agents communicate with neighbouring agents in synchronized communication windows. Within these communication windows negotiations for trip planning and booking have to be accomplished.

In contrast to the problems of information dissemination in mobile sensor networks, negotiations for ad-hoc shared-ride trip planning require two-way communication. This negotiation process consists of three steps: (i) the client sends a request into the network, (ii) the hosts having relevant information return offers, and (iii) the client books the host with the optimal offer. Negotiation needs some kind of transactional protocol that makes clear to both clients and hosts that they created a contract. Requests of clients form messages that are addressed to everyone (no addressee in particular), and disseminated into the network. Offers from hosts and booking messages from clients, however, are directly addressed, passing them through a reinforced, preferred chain of communication hops between the client and the host. This path was established in the phase of the dissemination of the request.

Negotiations for ad-hoc shared-ride trip planning require communication windows being long enough to accomplish the full negotiation procedure. This means that a communication window has to allow multiple hops. At the same time communication windows have to be short enough to guarantee a stable communication network topology for directed messaging. As a rough estimate, if urban traffic flows with $30km/h$ a window of two seconds would allow nodes to move $16m$, or less than 20% of a radio range of $100m$. The movements are small enough to not (much) change the network topology, but the two seconds will technically limit the maximal number of hops. We further assume that no message survives a communication window, i.e., with the end of the communication window all for our purposes relevant communication processes shall be completed.

Therefore the communication strategies for information dissemination (Section 2.3) have to be replaced by a strategy that allows for two-way negotiations in limited ranges:

(i) *unconstrained* (closest match to the flooding strategy). Within one communication window each node of a geosensor network broadcasts every message it receives if it did not forward this message already. Clients can expect to get offers from all reachable hosts, and hence, they get the most complete knowledge of the current transportation network.

(ii) *short-range proximity* (closest match to the location-constrained strategy). Client requests are communicated only to agents within their radio range (single-hop), and offers and bookings are not forwarded either. The communication traffic in the network is drastically reduced compared to the previous strategy. Energy savings will be significant. However, the client reaches a much smaller number of hosts, and hence, will find suboptimal offers only.

(iii) *mid-range proximity* (another match to the location-constrained strategy). Client requests are passed on to a proximity defined by a number of hops. Compared to short-range proximity the communication traffic is increased, but the hosts that are reached are still in some proximity to the client. Thus the requests might reach more *relevant* hosts than the unconstrained communication strategy.

5 Formalization and design of a simulation system

As we are interested in the effectiveness and efficiency of the negotiation ranges for shared-ride trip planning, we develop a formal model of a street network with clients and hosts that can be implemented for simulation purposes. We carefully observe the criteria of credible simulation specifications (Pawlikowski et al. 2002, Kurkowski et al. 2005).

5.1 *The simulation parameters*

The simulation happens in a regular grid 'street' network, and is structured by clock cycles. Each cycle consists of two phases: one (instantaneous) phase of negotiations between clients and hosts (a few seconds in the real world), and one phase of moving. After each cycle all agents are located at (or allocated to) intersections. Furthermore, the radio range is assumed to be limited to one street segment, i.e., to the four-neighbourhood of each intersection.

Clients know their current position and destination. They are immobile and can only travel with hosts. Furthermore they apply a simple heuristic wayfinding strategy: they travel only along the route of the graph geodesic to their desired destination. They look for the fastest trip along this route; other cost factors are neglected. The simulation knows a single client, and competing clients are modelled by a parameter to specify average booking rates of hosts. The single client's route is located in the central part of the grid to avoid boundary effects in the simulation.

Hosts have randomly chosen trips of constant duration, which realizes a typical random walk mobility model (Camp et al. 2002) of finite trips. Furthermore, hosts are generated staggeringly: in each cycle, some hosts reach their destination, and new ones are constructed. Hence, the host density is constant over time, and there is some degree of surprise for every negotiation process. In our simulation all hosts are moving with the same speed of one segment per cycle.

The remaining parameter in this process is the negotiation range, which can take three values: unconstrained, short-range, or mid-range. In this system we are interested in two output parameters: the duration of the client's trip, and the number of broadcasted messages in all negotiation cycles during the travel. For the latter we study next the number of broadcasts in one negotiation process. The algorithm developed will then run in each negotiation process; the numbers of broadcasted messages simply add up. The simulation stops when the client reaches its desired destination.

This type of simulation is also called steady-state.

5.2 *The negotiation process*

To model a negotiation process, we first switch from the street network view (Figure 1) to a communication network view (Figure 2). Then we specify the messages to be exchanged and study their exchange.

Figure 1 shows a client C and seven hosts H_1-H_7 in the street network. With a radio range of one segment, Figure 2 shows the corresponding communication network. Two agents are connected by a link if they are in direct communication range to each other. We call this graph a neighbourhood graph.

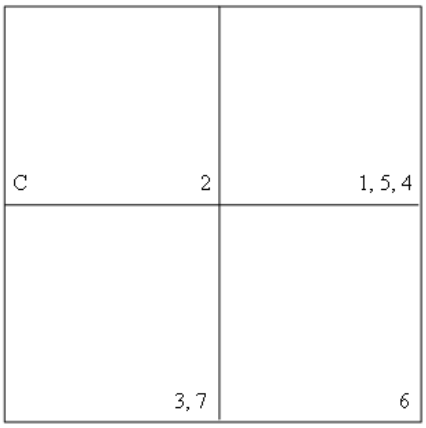

Figure 1. The locations of a client and seven hosts in a transportation network (snapshot).

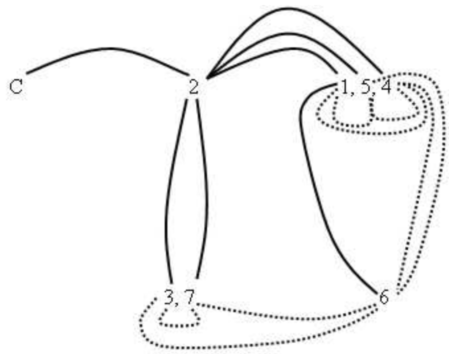

Figure 2. The communication network of the agents, and, as the subset of solid lines, the shortest path tree from C.

On this communication network we can demonstrate the three phases of each negotiation: sending requests r, sending offers o, and sending booking messages b. For the demonstration we apply the unconstrained negotiation range, and we will discuss in Section 5.3 the modifications for other negotiation ranges.

5.2.1 ***Requests.***

A client sends a request r specifying the sequence of street segments of their route ahead. In our example, the client C's request is broadcasted through the paths shown in Table 1. In this table, the agents that receive the request for the first time (i.e., on the shortest communication path) are printed bold; the other agents are printed in brackets. Only when agents receive a request for the first time they broadcast it. That means, in this situation each agent in the connected network broadcasts once. In other words, with an unlimited negotiation range the number of broadcasts of a request is equal to the number of agents in the client's communication network.

Table 1. Paths of request messages.

sender	message	receiver
C	r	**2**
2	r	**1**, **5**, **4**, **3**, **7**, [C]
1	r	[2], [5], [4], **6**
5	r	[2], [1], [4], [6]
4	r	[2], [5], [1], [6]
3	r	[2], [7], [6]
7	r	[2], [3], [6]
6	r	[1], [5], [4], [3], [7]

Furthermore, we introduce a message protocol that generates a history of hops. Each broadcasting agent attaches its name to the request r, as shown in Table 2. In that way, each recipient knows the shortest path back to the client sending the request. This information can be exploited for directing the offer and booking messages.

Table 2. The history of hops attached to each request.

agent	received request	broadcasted request
C		r, C
2	r, C	r, C, 2
1	r, C, 2	r, C, 2, 1
5	r, C, 2	r, C, 2, 5
4	r, C, 2	r, C, 2, 4
3	r, C, 2	r, C, 2, 3
7	r, C, 2	r, C, 2, 7
6	r, C, 2, 1	r, C, 2, 1, 6

5.2.2 ***Offers.*** Any host receiving a request that matches in some parts with its own travel plans, and having still free capacity, will respond by an offer o. The offer specifies the identified street segments and their time stamps in the host's schedule. An offer is addressed and directed by reversing the history of the request. Only agents on this list will forward the message.

In our example hosts H_6, H_3, and H_2 are going to make an offer to C (o_6, o_3, o_2). The set of broadcasts for these offers is shown in Table 3. In the table, the hosts in parenthesis are receiving a message, but are not on the address list, and hence, do not forward the offer. Clients do not forward offers addressed to them. In other words, each offer causes a number of broadcasts equivalent to the length of the shortest path branch between the offering host and requesting client. For illustration, Figure 2 shows the client's shortest path tree of the neighbourhood graph.

Table 3. The paths of the offers.

sender	message	receiver
6	o6	(3), (7), **1**, (5), (4)
1	o6	[(5)], [(4)], **2**, [6]
2	o6	[1], [(5)], [(4)], (3), (7), **C**
3	o3	(7), (6), **2**
2	o3	(1), (5), (4), [3], [(7)], **C**
2	o2	(1), (5), (4), (3), (7), **C**

5.2.3 ***Bookings.*** The requesting client collects all offers, and selects the optimal one(s). Within our specifications the optimal offer is the one that promises the earliest start. This choice has to be booked with the offering host(s).

In our example client C is going to accept an offer o_3 from host H_3. The set of broadcasts for the booking message b_3 is listed in Table 4. The table shows that each booking causes a number of broadcasts again equivalent to the length of the shortest path branch between the client and the offering host.

Table 4. The paths of the booking messages.

sender	message	receiver
C	b3	**2**
2	b3	(1), (5), (4), **3**, (7), [C]

Client C would also like to cancel a previous booking with host H_7 (cancellation message c_7). Note that C currently has offers only from H_6, H_3, and H_2 in hand, and hence, does not know where H_7 is. Host H_7 may even be disconnected (it is connected in our example). Because cancellation messages cannot be guaranteed to reach their addressee, an alternative way of cancelling is used: previous bookings, if not confirmed in this negotiation process, will time out automatically before the next negotiation cycle.

5.3 *Counting messages in a negotiation process*

The negotiation process discussed above gives reason for the following algorithm to count the broadcasted messages (Algorithm 1). First the algorithm computes for a neighbourhood graph (line 3), and on this graph the shortest path tree (Dijkstra 1959) from the client (line 4). Then, of particular interest are the lines where the counter for the number of broadcasts of messages $no\_of\_messages$ is increased.

The counter is set first to the number of connected agents since all agents broadcast a request (line 7). Subsequently, the counter is increased by the lengths of shortest path tree branches of offering hosts since an offer is broadcasted by all agents along the shortest path tree branch from host to client (line 11). Finally, the counter is increased by the same amount of broadcasts for the booking of a host by the client (line 14).

Algorithm 1. Counting the number of messages broadcasted in one negotiation cycle.

```
   Data: Location and travel interest of client, location and travel plans of hosts.
   Result: Number of messages in a realization of a full negotiation cycle.
1  no_of_messages = 0;
2  foreach client do
3      construct neighborhood graph;
4      calculate shortest path tree from client;
5      no_of_agents = number of nodes in neighborhood graph;
6      generate request;
7      no_of_messages+ = no_of_agents ;
8      foreach host in neighborhood graph do
9          if going to make an offer then
10             generate offer ;
11             no_of_messages+ = length of shortest path branch (host, client) ;
12     choose optimal offer ;
13     generate booking with host ;
14     no_of_messages+ = length of SP branch (client, host) ;
```

This algorithm was developed so far for the unconstrained negotiation range. However, it needs only a small modification to work for spatially constrained negotiation ranges as well. For this purpose we introduce an additional parameter m specifying the radius of the range. If $m = 1$ the simulation realizes a short-range negotiation, and if m is larger the simulation realizes a mid-range negotiation. The unconstrained negotiation can be considered as the special case of $m = \infty$. The parameter m can be part of the request message. Each agent receiving such a request determines the number of the previous hops p (length p of the request history), and forwards the request only as long as $p < m$. The rest of the negotiation process remains unchanged. With other words, all that it needs is to cut the shortest path tree at level m in Algorithm 1.

6 Simulation results

The simulation was implemented in Java, and then observed for varying simulation parameters. In the simulation, two characteristics are observed: the total number of broadcasted messages and the number of time intervals the client C is travelling. The reported results are average values and confidence intervals for large numbers of simulations. The source code can be obtained from the first author to study the simulation in detail or to repeat our experiments.

For all experiments, the simulated world consisted of a 10×10 street grid, and the client's route in the center of the world was of length 5. Hosts were generated at random locations and with random travel plans of constant length; the density of hosts was kept constant over the duration of each experiment. Competition for seats is introduced by a chance of a host being booked of 33%, and the radius m of the mid-range communication strategy is set to 3. For each parameter pair of negotiation range and host density we ran 1000 simulations.

The first insight is demonstrated in Fig. 3. Independently for the particular negotiation range it shows the consequences of the variation in the density of hosts on the probability of getting a ride. The higher the host density becomes the shorter the trip durations. The relationship goes asymptotically to the route length since ideally there are always hosts that offer a ride for the next street segment.

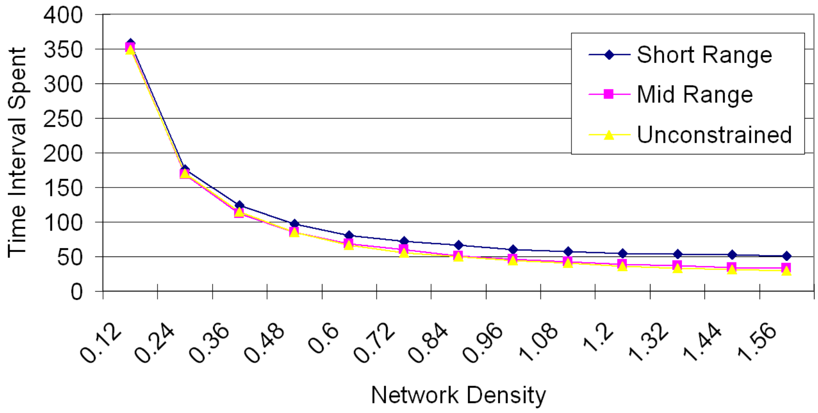

Figure 3. The average travel time depending on the negotiation range for various host densities.

The next step to investigate is the quality of the found trips depending on the chosen negotiation range. Remember that the chosen optimization criterion is trip duration; the quality of the trip increases (only) as the trip duration decreases. Figure 3 shows three curves, one for each negotiation range. All curves behave similarly by decreasing asymptotically towards the route length. However, the short-range negotiation does not come down as fast as the other two. For example, with a host density of 1.56 hosts per street node the client needs on average 50 time intervals to reach its destination with the short-range negotiation, but only 33 with mid-range and 30 with the unconstrained negotiation (see Table 5 for details). Thus the short-range negotiation is significantly less effective, but mid-range and unconstrained negotiation are nearly not distinguishable in effectivity.

The last question to be investigated concerns the number of messages sent by the different negotiation ranges. Figure 4 shows the steep increase of messages created by an unconstrained negotiation range. This range is by far the least efficient, and this effect is the stronger the higher the host density. For a host density of 1.56 the

Table 5. Results for a host density of 1.56, with standard deviations and 95% confidence intervals.

	trip duration		σ	ci	messages		σ	ci
	cycles	%			no.	%		
short range	50	152	20	1.2	77	21	25	1.5
mid range	33	100	13	0.8	369	100	134	8.3
unconstrained	30	91	12	0.8	3986	1078	1651	102.3

unconstrained negotiation range produces on average 3986 messages, the mid-range strategy 369, and the short-range strategy only 77 (see again Table 5 for details).

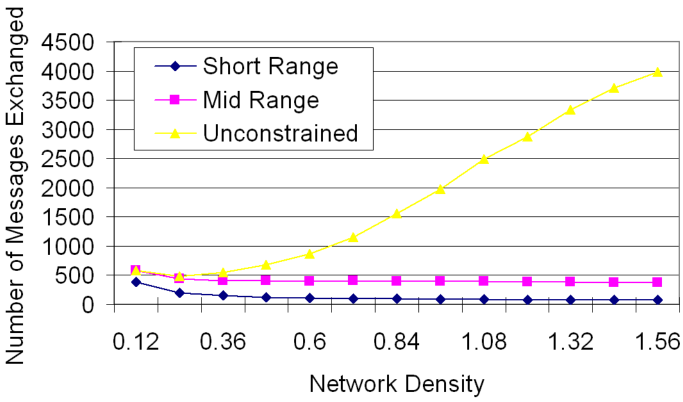

Figure 4. The number of messages exchanged with the three negotiation ranges for different host densities.

For the unconstrained negotiation range, the number of broadcasted messages will increase continuously with the density of hosts. Since for the densities beyond the right end of Figure 4 on average all hosts are connected, broadcasts of requests increase linearly with the number of hosts. Broadcasts of offers do increase much less significantly since far hosts frequently do not contribute to the requested trip. However, the number of hops of these messages can grow, as of an eventual booking message from the client.

For the short-range negotiation, the number of broadcasted requests is constantly 1 for each negotiation cycle, since requests are not forwarded. Eventual offers and bookings are also broadcasted only by the original senders. While numbers of requests are a function of the trip duration, numbers of offers are a function of route length and host density, and numbers of bookings are a function of the route length only. Trip duration and host density are negatively correlated, which lets approach the curve asymptotically a constant.

The mid-range negotiation mixes the two behaviours discussed before. It appears to be limited in its growth because the communication range is limited. Hence, the number of broadcasted requests does no longer increase with the total number of hosts, but only with the number of hosts in the chosen range. Since this set of hosts contains on average most of the for the request relevant hosts (compare Figure 3), the reduction of numbers of messages does not reflect in an increase of the duration of the trip. Hence, this strategy is efficient.

7 Conclusions

In this paper we demonstrated by simulation that shared-ride trip planning can be accomplished by the availability of an ad-hoc mobile geosensor network. Furthermore we show that this solution is *effective* and *efficient*. While it is most effective with an unconstrained negotiation range, this strategy is inefficient from an energy and bandwidth standpoint. This strategy is also not feasible: the necessarily short communication windows limit practically the number of hops of messages. In contrast, the short-range negotiation is the most efficient, but least effective. Compared to the two, the mid-range negotiation proves the hypothesis: it is effective (e.g., for 1.56 hosts per node 10% longer trips than with unconstrained negotiations (on average), but 66% shorter trips than with short-range negotiations) and it is efficient (e.g., for 1.56 hosts per node 5 times more messages than with short-range, but 9% of the messages with unconstrained negotiations—and this number is steeply decreasing with an increase of host density). It can be expected that the trends reflected in these numbers hold for different street network forms and mid-range thresholds.

The results relate to a specific wayfinding strategy, travelling along the shortest distance route. It can be expected that for other wayfinding strategies the results will be in principle the same since the relationship between the nearness of the agents and their relevance for each other holds universally.

In future work we will investigate the following open questions.

(i) In this paper, we have chosen an inflexible wayfinding strategy. Strictly following the shortest distance route might result in longer trip durations in many contexts. In real transportation networks the shortest distance route is not necessarily the fastest, e.g., when the network is hierarchic. One extension of this paper is investigating different wayfinding strategies with flexible route choice.

(ii) In this paper, we have only optimized travel time. Other optimization functions can be chosen, and especially multiple-criteria optimization has real applications in multi-modal transportation. For example, our simulation could be extended to minimize the number of transfers during a trip, or to find an adjusted optimum between travel duration and travel convenience.

(iii) In this paper, we assumed an equal probability distribution for random booking, which is a sufficient first approximation. But a consequence of conservative booking is an unequal booking distribution over the time intervals ahead. Thus, with conservative booking there are less hosts available for mid-range planning. Clients will find this counterproductive, and might restrict themselves to less greedy booking strategies for a common benefit. Investigating rigidly the consequences of conservative booking, and comparing it with the effects of other booking strategies, is an interesting question for the future.

(iv) One assumption in our simulation is equal behaviour of all transportation hosts. In real urban traffic this might be a sufficient first approximation. However, integrating different modes of transportation, especially different speeds and different pricing, requires a relaxation of this condition. A related extension of our simulation would allow client agents to walk, at least single segments, to bridge small gaps in the transportation network.

(v) In this paper we consequently considered updating as an active search process of the client agents. A central service would prefer event-triggered messaging, reducing the planning tasks to times when needed. This passive process can be investigated for geosensor networks as well. Controlling the revision of trip plans by events (Worboys and Hornsby 2004) means that client agents act only when they approach a gap in their bookings, or when a new host appears in their field of observation. And hosts act if they enter the traffic, or if they find bookings dissolved.

This paper focused on route planning with local, incomplete knowledge. Other issues of an ad-hoc shared-ride system, like social, economic, or privacy issues, still need to be investigated.

Acknowledgements

Many thanks to Trang Cao and Lin Jie Guang, University of Melbourne, who have implemented the simulation and computed the results. The idea of local trip planning was born in a discussion with Monika Sester, University of Hannover. This work has been supported by the Australian Academy of Science with a travel grant for the first author. The work of the second author was partially supported by the National Science Foundation under NSF grant number EPS-9983432.

REFERENCES

Camp, T., Boleng, J., Davies, V., 2002. A survey of mobility models for ad hoc network research. Wireless Communication and Mobile Computing 2 (5), 483–502.

Chon, H. D., Agrawal, D., Abbadi, A. E., 2003. Fates: Finding a time dependent shortest path. In: Chen, M. S. (Ed.), MDM 2003. Vol. 2574 of Lecture Notes in Computer Science. Springer, Berlin, pp. 165–180.

Cooke, K. L., Halsey, E., 1966. The shortest route through a network with time-dependent internodal transit times. Journal of Mathematical Analysis and Applications 14, 493–498.

Dijkstra, E. W., 1959. A note on two problems in connexion with graphs. Numerische Mathematik 1, 269–271.

Dillenburg, J. F., Wolfson, O., Nelson, P. C., 2002. The intelligent travel assistant. In: 5th International Conference on Intelligent Transportation Systems. IEEE, Singapore, pp. 691–696.

Fu, L., 2001. Adaptive routing algorithm for in-vehicle route guidance systems with real-time information. Transportation Research Part B 35 (8), 749–765.

Gerla, M., Chen, L.-J., Lee, Y.-Z., Zhou, B., Chen, J., Yang, G., Das, S., 2005. Dealing with node mobility in ad hoc wireless network. In: Bernardo, M., Bogliolo, A. (Eds.), Formal Methods for Mobile Computing. Vol. 3465 of Lecture Notes in Computer Science. Springer, Heidelberg, pp. 69–106.

Hall, E. T., 1966. The Hidden Dimension. Doubleday & Company, Garden City, NY.

Hochmair, H., Frank, A. U., 2002. Influence of estimation errors on wayfinding-decisions in unknown street networks analyzing the least-angle strategy. Spatial Cognition and Computation 2 (4), 283–313.

Klafszky, E., 1972. Determination of shortest path in a network with time-dependent edge-lengths. Mathematische Operationsforschung und Statistik 3 (4), 255–257.

Kurkowski, S., Camp, T., Colagrosso, M., 2005. Manet simulation studies: The incredibles. Mobile Computing and Communications Review 9 (4), 50–61.

McCarthy, P. S., 2001. Transportation Economics. Blackwell Publishers Inc., Malden, MA.

Monmonier, M., 2002. Spying with Maps. University of Chicago Press, Chicago.

Naor, M., 2005. On fairness in the carpool problem. Journal of Algorithms 55, 93–98.

Nijkamp, P., Pepping, G., Banister, D., 1996. Telematics and Transport Behaviour. Advances in Spatial Science. Springer, Berlin.

Nittel, S., Duckham, M., Kulik, L., 2004. Information dissemination in mobile adhoc geosensor networks. In: Egenhofer, M., Freksa, C., Miller, H. J. (Eds.), Geographic Information Science. Vol. 3234 of Lecture Notes in Computer Science. Springer, Berlin, pp. 206–222.

Noda, I., Ohta, M., Shinoda, K., Kumada, Y., Nakashima, H., 2004. Evaluation of usability of dial-a-ride systems by social simulation. In: Hales, D., Edmonds, B., Norling, E., Rouchier, J. (Eds.), Multi-Agent-Based Simulation III. Vol. 2927 of Lecture Notes in Artificial Intelligence. Springer, Berlin, pp. 167–181.

Orda, A., Rom, R., 1990. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. Journal of the ACM 37 (3), 607–625.

Pawlikowski, K., Jeong, H.-D. J., Lee, J.-S. R., 2002. On credibility of simulation studies of telecommunication networks. IEEE Communications Magazine 40 (1), 132–139.

Peng, Z.-R., Tsou, M.-H., 2003. Internet GIS. John Wiley & Sons, Hoboken, NJ.

Resnick, P., 2004. Impersonal sociotechnical capital, ICTs, and collective action among strangers. In: Dutton, W., Kahin, B., O'Callaghan, R., Wyckoff, A. (Eds.), Transforming Enterprise. MIT Press, Cambridge, Mass., pp. 399–412.

Roussopoulos, M., Baker, M., Rosenthal, D., Giuli, T. J., Maniatis, P., Mogul, J., 2004. 2 P2P or not 2 P2P? In: Voelker, G. M., Shenker, S. (Eds.), Peer-to-Peer Systems III: Third International Workshop IPTPS 2004. Vol. 3279 of Lecture Notes in Computer Science. Springer, Berlin, pp. 33–43.

Stefanidis, A., Nittel, S., 2005. GeoSensor Networks. CRC Press, Boca Raton, FL.

Sussman, J., 2000. Introduction to Transportation Systems. Artech House, Boston, MA.

Wolfson, O., Xu, B., 2004. Opportunistic dissemination of spatio-temporal resource information in mobile peer-to-peer networks. In: First International Workshop on P2P Data Management, Security and Trust. DEXA, Zaragoza, Spain, pp. 954–958.

Worboys, M., Hornsby, K., 2004. From objects to events. In: Egenhofer, M., Freksa, C., Miller, H. J. (Eds.), Geographic Information Science. Vol. 3234 of Lecture Notes in Computer Science. Springer, Berlin, pp. 327–343.

Zhao, F., Guibas, L. J., 2004. Wireless Sensor Networks. Elsevier, Amsterdam.

Ziliaskopoulos, A., Mahmassani, H., 1993. A time-dependent shortest path algorithm for real-time intelligent vehicle highway system applications. Transportation Research Record 1408, 94–104.

Ziliaskopoulos, A., Wardell, W., 2000. A intermodal optimum path algorithm for multimodal networks with dynamic arc travel times and switching delays. European Journal of Operational Research 125 (3), 486–502.