

Geo-Social Interaction: Context-Aware Help in Large Scale Public Spaces

Nasim Mahmud¹, Petr Aksenov¹, Ansar-Ul-Haque Yasar², Davy Preuveneers²,
Kris Luyten¹, Karin Coninx¹, and Yolande Berbers²

¹ Hasselt University - tUL - IBBT,

Expertise Centre for Digital Media,

Wetenschapspark 2, B-3590 Diepenbeek, Belgium

² Department of Computer Science, Katholieke Universiteit Leuven,

Celestijnenlaan 200A, B-3001 Leuven, Belgium

{nasim.mahmud, petr.aksenov, kris.luyten, karin.coninx}@uhasselt.be,
{ansarulhaque.yasar, davy.preuveneers, yolande.berbers}@cs.kuleuven.be

Abstract. We present an approach to exploit social and spatio-temporal context in order to improve information dissemination in dynamic large-scale public spaces. We illustrate it by applying a proposed measure of geo-social relevance of each individual in a simulated vehicular network and by comparing the performance of different network message passing techniques in an inter-vehicle ‘help-me-best-and-do-it-fast’ communication scenario. We conclude that the use of social networking capabilities of an individual combined with knowledge about their spatio-temporal context information significantly improves purposeful interaction between individuals in terms of both the efficiency of the network data dissemination and the quality of the delivered information.

Keywords: Social network, FOAF, Context awareness, Location awareness, Information dissemination.

1 Introduction

Being able to obtain the right information at the right time has always been a challenge when taking informed decisions. Unfortunately, people who are ‘on-the-move’ often do not have an opportunity to spend a long time looking for what they need. Systems (e.g., intelligent transportation systems) also may require efficient delivery of timely and relevant information, e.g. to optimise traffic flows. The proliferation of novel wireless network technologies has created new opportunities for complex peer-to-peer information dissemination systems, and a key challenge in this area is how to interact, locate and communicate effectively in a large scale public environment.

Context-awareness is often an instrument to decide what information is relevant and as such can help to improve communication and routing efficiency since it allows making informed decisions on the locality and information necessity of different moving objects. In our previous work, we have shown that context

awareness improves network efficiency in a vehicular network by exploiting the user's context in making informed decisions on the data routing [7]. In order to provide improved help and facilitate users to achieve their goals, the system should be able to recognise their context, such as the current place, state, activity, interests, etc. Among the many facets of context, location has been widely recognised as one of the primary context areas and it is the main concern of many modern mobile applications. The social context has also recently taken one of the leading roles in context-awareness. The importance of these two areas have, in turn, led to the concept of location-based social networks and services. Here, social and spatiotemporal branches of context are closely tied together, having given rise to a new type of 'geo-social' interaction, collaboration and information sharing. Some recent examples of the work in this area include the Connected Traveler [4] project that aims at providing relevant information to drivers by means of taking into account their quickly changing location and personalised preferences in the way that only the information the driver explicitly set to be interested in is selected. Connecto [1] is a phone based status and location sharing application, which evolved, throughout use within a small group of friends, from serving for mere location updates into a tool for enriched social interaction through these location-based updates. CityFlocks [2] is a mobile system that allows information seeking visitors to access tacit knowledge from local people about their new community. Commercial applications, such as Foursquare¹ and Gowalla², allow sharing one's current location and recommending (or ranking) favourite places to friends, which proved to be a very attractive addition to the basic social networking capabilities.

These and many other examples show that the area of using location and social context together is very broad. In this paper, we investigate how exploiting a social network, user's personal preferences, information about user's location and movements as well as presence and preferences of other users allows for interaction that helps users in large-scale highly dynamic and populated public spaces to obtain timely and relevant information. In particular, we aim at selecting the most appropriate individual as an assistant in a 'Help me!'-like scenario between members of a vehicular network. Our approach involves a comparison of social and spatial relevance between two peers in terms of the area of expertise and spatial closeness. We have utilised social network analysis to find the most appropriate peer in terms of the possessed knowledge and the distance to cross and have factorised the urgency of the help seeker in terms of time the assistance is being looked for.

2 Three-Leaved Mirror Approach

In our three-leaved mirror approach, the spatial and the social contexts each resemble one of the two side leaves of such a mirror. Like it is with a real three-leaved mirror when the angles of the leaves can be adjusted independently

¹ Foursquare: <http://foursquare.com>

² Gowalla: <http://gowalla.com>

to change the current view of the user from each side, the contribution of the spatial and social context to the overall relevance varies. The central leaf then optimises the information flow in the network by reducing irrelevant information dissemination, thus making the view of the user on the issue at hand more complete, enabling to see ‘the big picture’.

2.1 The Side Leaves: Geo-Social Component

Like network availability, location-awareness has become a ubiquitous property of many mobile devices nowadays. With both networking and location detection technologies readily available in the device itself, we can combine the information available on the user’s location with information provided by the user’s social network. For example, we can easily select a contact from the user’s social network that is in the vicinity and ask this person for assistance. The difficulty lies in dealing with the additional complexity caused by persons in the social network moving and the timeliness of certain requests for assistance (e.g. when no help arrives within 5 minutes the requester will have to manage on their own). We used vehicular networks for evaluating our approach. Since this type of network is highly dynamic, contains fast moving objects and requires a close eye on the timeliness, it serves as the perfect evaluation framework to assess our approach.




Fig. 1. Talking cars. 1(a) Help seeker H sends a help request using an embedded Geo-social Ubiquitous-Help-System (UHS); 1(b) A, B, and C are ready to help H; 1(c) Depending on the returned geo-social relevance score of the helpers, H confirms the offer.

2.2 The Central Leaf: Improved Relevance Backpropagation with Geo-Social Relevance

Geo-social interactions require communication between people, and optimised information dissemination is its vital aspect. There exist a number of information routing strategies, both with and without taking into account the quality of information (QoI) (e.g., broadcasting and backpropagation, respectively). For example, Eichler et al. [3] address the issue of optimal information dissemination in vehicular networks by proposing a framework which integrates many of the existing broadcast-based strategies that deal with reduction of the superfluous transmissions. We propose a QoI-based best-effort mechanism for intelligent

adaptive context dissemination using social interactions with a relevance score function (see section 4) in a large scale vehicular network based on our previous work [5]. In general, the technique can be applied to many types of networks where social interactions matter, and currently we chose vehicular networks since they are a comfortable means for verification as they contain sufficient diversity in terms of both size and dynamics. In our case, each participant has a list of friends (e.g., in a Friend-of-a-friend (FOAF) profile), a score value and properties about the relevant context information they can provide. The algorithm relies on the feedback of neighbours to reduce the number of peers to forward the information to. The information is forwarded to the adjacent nodes who are either friends or friends-of-a-friend having certain degree of score value unless a maximum number of hops is reached. Each forwarding node reduces the hop counter, adds its identifier and marks the message relevancy tag if the information is relevant for its purpose and grades the sending node positively adding it to the friends list. The feedback technique is based on the context information like position, velocity, direction, time-to-live, interest, etc., that decides whether the received data are relevant which helps determine the information relevancy on the intermediate nodes. The feedback to the delivering node is initiated if the context information is relevant, irrelevant, unused, or duplicate information is received. It ensures that the provided information is from a trusted node which is supposed to be accurate and relevant for the receiver. In this mechanism the goal is to efficiently filter and route the relevant information as close to the source as possible in vehicular ad-hoc networks (VANETs).

3 Aspects of Geo-Social Interaction

Our three-leaved mirror approach poses a set of generic aspects to consider for providing context-aware help in large scale public spaces.

Spatial coverage. It is always desirable to know the exact location of an incident for context-aware applications. For instance, in the case of an incident on the road, the authorities should be notified about the exact location to react fast and the information should be delivered only at the right place. Similarly, a context-aware application should be able to sense, manipulate and disseminate context information about direction and velocity of vehicles in the network to predict certain situations like traffic congestions or traffic accidents in specific regions.

Timeliness. It is crucial that the information being disseminated in a large scale network between nodes (or groups of nodes) reaches the destination on time. Timeliness uses time as a relevance criterion for information sharing so that the information received by a node is not older than the 'lifetime of the information' so that the right information can be delivered at the right time.

Completeness. Lack of information can lead to ambiguity. This measure identifies the quality of information (QoI) that is provided by a node in the network and can be comparing the number of attributes received to the total number of attributes to make a well-informed decision.

Trust-worthiness. The information being received at a particular node should be reliable, accurate and trustworthy. Each node involved in a large scale vehicular network should have a social profile listing sharing some common interest with a certain quality of information (QoI).

Significance. It indicates the importance of a certain type of contextual information required by a node in a network. The value of significance will be higher in the case of a life threatening situation.

4 The Geo-Social Relevance Function

The key concept to our approach is the geo-social relevance function that defines how relevant a potential help provider is to the help seeker. It naturally accumulates together certain parameters from each of the three leaves (spatial, social, and network-bound) and results in a sort of weight each involved node of the network has with respect to the help requester. The higher the score, the better the provider. The function is expressed using the following formula:

$$GSR(peer) = A \cdot \sqrt[n]{R \cdot HT} \cdot F_U$$

where the multipliers in the right part are as follows: **Availability A** is a Boolean value that simply indicates whether the corresponding node can be a potential help provider. We assume that if a peer is unavailable ($A = 0$), the help request is still passed further to this peer's friend-list. **Reliability R** of a node in the network is a peer-determined integer value between 1 and 10 indicating how helpful the corresponding node has been in the past. **Help-type HT** measures the requester's and the provider's technical match and is an integer between 0 ('I know nothing') and 10 ('A perfect match'). **Root index n** stands for the number of hops in the network between the requester and provider. The reason for choosing the root-based value for measuring the contribution of the social parameters to the overall relevance is that the level of trust to somebody who is connected to you indirectly decreases significantly. **The spatiotemporal contribution, F_U** , is defined as

$$F_U = \begin{cases} 0 & \text{if } Direction = 0 \\ 0 & \text{if } U \cdot V \leq D_{min} \\ e^C & \text{if } D_{min} \leq U \cdot V \leq D_{max} \\ e \cdot C & \text{if } U \cdot V > D_{max} \end{cases}, \text{ where } C = \frac{U \cdot V - D_{min}}{D_{max} - D_{min}}$$

Here, **Direction** equals 1 if there is a 'movingTowards' relation from the provider to the requester, and 0 otherwise. Urgency **U** is the time interval within which the help is needed; its value is specified by the requester. Velocity **V** is an estimated average velocity of the help provider during period **U**. We assume that the help requester does not move (e.g., their car broke down). The corresponding maximal and minimal distances, D_{max} and D_{min} , between two nodes are calculated at the time when the help request has been received and depend on the spatial topology of the area, such as, for example, the actual length of the

connecting path between the nodes that might be affected by possible repair works, closed or blocked roads, etc. The reasoning behind the expression for F_U is such that for nodes that are far ($0 \leq C < 1$), $e^C > e \cdot C$ and thus e^C has a bigger weight for the score; and for $C \geq 1$, the smaller value of $e \cdot C$ is used thus making the social parameters weigh more than the spatial ones when comparing the scores of two different nodes.

This way, there exists a sort of balance between the social relativity and the spatial closeness so that nodes with different social profiles, but in the direct vicinity, can get equal, or close enough, scores and be chosen by the help requester in accordance with the requester's personal preferences. Besides, the score function makes all members of the network - the requester (urgency U), potential providers (availability A), and all other nodes in the network (reliability R) - collaborate implicitly in finding the fittest solution to the help request.

5 Validation by Simulation

In order to enable intelligent communication between vehicles we need to take into account the social characteristics of the people involved. Since large scale groups are not easily tested with real-life situations, we evaluated our improved relevance backpropagation algorithm using a real time discrete event-based network simulator (OMNeT++)³ to run on a large scale vehicular network using a realistic dataset [6] logged for a period of 24 hours. Individuals in the simulation are distributed over an area of 250 by 260 km. All individuals choose a time to travel and a route in accordance with where they live and current road congestion. The complete dataset contains more than 25,000,000 recorded direction and speed changes of 260,000 vehicles, from which we randomly selected 300 vehicles. Besides, for the sake of simplicity, we normalised the original movement data to fit in the range of [0; 1].

5.1 Data Preparation and Setup

Taking into account the aspects discussed in section 3, this section describes the details of the actual data we used in the simulation.

Friendlist. In general, the number of friends in the friend-list of a node in the network is not limited. However, given the specifics of our simulation and the actual behaviour the simulation is reproducing, we limited each node to have a maximum of 15 friends, the exact pre-simulation initial number of friends being assigned based on the type of help a node can provide (see Table 2). Friend-lists get extended in an asymmetric way: each time a help requester has received help from an appropriate provider, the latter is added to the requester's friend list. In the case of a new help request from the same node, chances are that the same provider can help again, and contacting this provider directly will save time and resources. An extract from the complete assignment showing the initial friend lists of the first five nodes is given in Table 1.

³ OMNeT++: <http://www.omnetpp.org>

Helptype. We introduced nine help-types so that each of the 300 nodes belongs to one of them. The total number of nodes distributed normally among each help-type is as follows using the (help-type - number of nodes of this type) notation: (1-11), (2-26), (3-43), (4-55), (5-55), (6-45), (7-32), (8-21), and (9-12).

The matching table of the HT values corresponding to a pair of help-types is shown in Table 2. Notice the asymmetric nature of the HT values for the corresponding (ReqHT1 to PrHT2) and (ReqHT2 to PrHT1) pairs, meaning that in general, ‘If you are able to help me with my problem, it does not guarantee that I am able to help you with yours.’

Table 1. Initial distribution of friends into friendlists for the first five nodes **Table 2.** Pair-matching of help-types

Node	F1	F2	F3	F4	F5	F6	F7	F8	F9	...	F15
1	79	117	93	0	0	0	0	0	0	...	0
2	10	1	54	74	126	165	184	0	0	...	0
3	10	4	77	44	88	84	189	177	0	...	0
4	11	64	186	0	0	0	0	0	0	...	0
5	74	76	113	97	150	273	0	0	0	...	0

	PrHT1	PrHT2	PrHT3	PrHT4	PrHT5	PrHT6	PrHT7	PrHT8	PrHT9
ReqHT1	10	0	8	8	8	0	0	7	0
ReqHT2	8	10	0	0	0	8	6	8	4
ReqHT3	8	0	10	9	7	0	8	6	8
ReqHT4	9	4	9	10	9	0	4	0	8
ReqHT5	6	6	4	7	10	6	9	8	0
ReqHT6	7	8	0	4	0	10	7	9	9
ReqHT7	4	7	7	8	6	8	10	0	6
ReqHT8	0	9	0	0	4	9	8	10	7
ReqHT9	0	0	6	6	8	7	0	4	10

Reliability. An integer value between 1 and 10 was assigned randomly to each node’s reliability. In general, reliability increases or decreases dynamically based on nodes’ performance as helpers but for the purposes of our simulation, we kept the help-type values static throughout the simulation.

Urgency. Out of the 300 nodes, we picked 10 which would, each at a random point during the simulation time, become a help requester with an individual urgency value the help request is valid for. The urgencies of 600, 300, 120, 180, 120, 120, 120, 90, 180, 540, and 240 seconds, were used by each help requester, respectively.

Since the available measurements of nodes’ movements contained only a timestamp and 2D-coordinates, other required values had to be further derived. Thus, the average velocity of each node over a certain time interval was calculated using the next two subsequent measurements for the node in question which gave an acceptable approximation in terms of the time the help requests remained valid for. There were also no data on the location error of the provided measurements. Therefore we divided the entire area into three sub-areas throughout each of which a measurement had a precision value from a specified interval. The precision of a node’s current position is a characteristic of a localisation system, which tracks the node’s position, and can be defined as the radius of a circle centered in the detected location. The circle then means the node is located somewhere within it. We set the precision intervals to be 0.0005 – 0.0001, 0.0005 – 0.001, and 0.001 – 0.005, respectively, with 0.0001 equal to approximately 25 metres on the original scale. Since we did not have a description of the spatial topology of the area, we assumed the distance between two nodes to

be a simple Euclidian distance. Together with the assigned precisions, it gave the following formulae for computing the D_{min} and D_{max} values between two nodes at the time of calculation:

$$\begin{aligned} D_{min} &= \min(0, D_{mes} - P_A - P_B) \\ D_{max} &= D_{mes} + P_A + P_B \\ D_{mes} &= \sqrt{(x_A - x_B)^2 + (y_A - y_B)^2} \end{aligned}$$

where P_A and P_B are the precision values of the corresponding measured locations of the two interacting nodes. Notice that in the case when the precision circles intersect, D_{min} equals 0.

5.2 Details of Simulation Runs

In our experiments, we let the nodes move around like cars and let connections appear and disappear according to the range to other nodes. The parameters we have taken into account for each node are: (i) Time, (ii) Velocity, (iii) X and Y coordinates, (iv) Number of packets sent, (v) Number of packets received, (vi) Number of forwarded packets and (vii) Time-to-live(TTL). Some nodes acted as context providers and some as context receivers. All nodes forward the information to their peers as long as the maximum TTL has not been reached and all context constraints are met.

We carried out three experiments with (a) our improved relevance backpropagation mechanism with social interactions and the geo-social relevance function, (b) simple relevance backpropagation and (c) state-of-the-art baseline case, plain broadcasting, each for a period of 24 hours.

6 Results

With both the (improved and simple) relevance backpropagation algorithms only relevant context information was sent out to the interested nodes in the network. There are several types of messages in the network: (i) sent (M_s), (ii) unique received (M_{ur}), (iii) unique sent (M_{us}) (iv) forwarded (M_f), (v) duplicate (M_d), and (vi) dropped (M_{drop}). During our simulated experimentation, we measured a set of major network metrics, Network Traffic (NT), Relevancy (R), Trustworthiness (T), Message Distance (MD), and Availability (A), expressed in terms of the above network messages as follows:

$$\begin{aligned} NT &= \Sigma_n(M_s + M_f) \\ R &= \Sigma_n((M_{ur} + M_d) - M_{drop}) / \Sigma_n(M_{ur} + M_d) \\ T &= (1 - (M_f / \Sigma_f(n)(M_{ur} + M_d))) \\ MD &= \Sigma_t Edges / \Sigma_t Nodes \\ A &= \Sigma_n M_{us} / \Sigma_n M_{ur} \end{aligned}$$

For each involved metric, we achieved a significant improvement in its performance using the proposed improved relevance backpropagation as compared

to both simple relevance backpropagation and broadcasting. For example, the utilisation of the Network Traffic is 90% of that using simple relevance backpropagation and as low as 50% of that of simple broadcasting. The five left plots in Figure 2) visualise such comparisons for all five metrics on the percentage scale.




Fig. 2. Improved relevance back propagation technique for routing messages in the network shows better results for each evaluated parameter

We also measured the performance of the geo-social relevance function in three cases of our experimental setup. Some returned scores were quite high in simple relevance backpropagation algorithm (e.g. 96) and low in the improved relevance backpropagation scheme (e.g. 16) for a certain type of help required. This meant that in some cases unexpected or unplanned nodes who just happened to be close enough had a better match than the algorithmically chosen ones, but when averaged over all returned values the improved algorithm outperforms the other (see the right plot in Figure 2).

7 Discussion and Future Work

The simulation results show that our improved relevance backpropagation mechanism achieves a significant improvement in terms of several network quality of information (QoI) parameters like relevancy, message distance, network traffic, availability and trustworthiness as discussed in section 6. The results in Figure 2 show that by eliminating redundant and irrelevant information sources with the help of the geo-social relevance function we can limit the information dissemination to happen within a much smaller number of nodes having a high degree of relevance, reliability and trustworthiness, thus improving the overall performance of a vehicular network.

The dataset we used for simulations contains limited attributes and focuses in the area of vehicular networks. Therefore we plan to conduct a deeper study on the actual values of social and location parameters (such as distribution of individuals' properties, help-type match values, etc.), augment a reward based mobile search in the UHS and later on cross validate our research in different application domains like pedestrian navigation or care-giving in hospitals. We also plan to improve the visualisation of ranking of search results provided to a user by conducting usability tests in a real life setting (e.g. inside a car).

Acknowledgments

Part of the research at EDM is funded by EFRO (European Fund for Regional Development) and the Flemish Government. Part of the research at K.U.Leuven is funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy, and by the Research Fund K.U.Leuven. Funding for this research was also provided by the Research Foundation Flanders (F.W.O. Vlaanderen, project CoLaSUE, number G.0439.08N).

References

1. Barkhuus, L., Brown, B., Bell, M., Sherwood, S., Hall, M., Chalmers, M.: From awareness to repartee: sharing location within social groups. In: Czerwinski, M., Lund, A.M., Tan, D.S. (eds.) CHI, pp. 497–506. ACM, New York (2008)
2. Bilandzic, M., Foth, M., De Luca, A.: Cityflocks: designing social navigation for urban mobile information systems. In: DIS 2008: Proceedings of the 7th ACM Conference on Designing Interactive Systems, USA, pp. 174–183 (2008)
3. Eichler, S., Schroth, C., Kosch, T., Strassberger, M.: Strategies for context-adaptive message dissemination in vehicular ad hoc networks. In: Proceedings of the 2nd Int. Workshop on Vehicle-to-Vehicle Communications (2006)
4. Manasseh, C., Ahern, K., Sengupta, R.: The connected traveler: using location and personalization on mobile devices to improve transportation. In: LocWeb. ACM International Conference Proceeding Series, p. 9. ACM, New York (2009)
5. Preuveneers, D., Berbers, Y.: Architectural backpropagation support for managing ambiguous context in smart environments. In: Stephanidis, C. (ed.) UAHCI 2007 (Part II). LNCS, vol. 4555, pp. 178–187. Springer, Heidelberg (2007)
6. Raney, B., Voellmy, A., Cetin, N., Vrtic, M., Nagel, K.: Towards a microscopic traffic simulation of all of switzerland. In: ICCS 2002: Proceedings of the Int. Conference on Computational Science, London, UK, pp. 371–380 (2002)
7. Yasar, A.-U.-H., Mahmud, N., Preuveneers, D., Luyten, K., Coninx, K., Berbers, Y.: Where people and cars meet: social interactions to improve information sharing in large scale vehicular networks. In: SAC 2010: Proceedings of the 25th ACM Symposium on Applied Computing, USA, pp. 1188–1194 (2010)