

RoadSpeak: Enabling Voice Chat on Roadways using Vehicular Social Networks^{* †}

Stephen Smaldone, Lu Han, Pravin Shankar, and Liviu Iftode
Department of Computer Science, Rutgers University, Piscataway, NJ
{smaldone,luhan,spravin,iftode}@cs.rutgers.edu

ABSTRACT

A great number of people spend one or more hours each day driving between home and the office. These daily roadway commutes are highly predictable and regular, and provide a great opportunity to form virtual mobile communities. However, even though these commuters are already physically present in the same location, they are limited in their ability to communicate with each other. This paper presents a framework for building such communities, which we call Vehicular Social Networks (VSNs), to facilitate better communication between commuters driving on roadways. As a proof of concept, we present the design of RoadSpeak, a VSN-based system that allows drivers to automatically join VSNs along popular roadways, and to communicate with each other by means of voice chat messages.

1. INTRODUCTION

Social Networks allow people with common interests to come together and form virtual communities. Mobile Social Networks connect people who are already together (sharing the same locality), enabling them to take advantage of their close proximity to form more tightly-coupled, possibly even ad-hoc, virtual communities. In this socially-driven virtual world, people can form “instant villages” that mirror and facilitate real-world interaction.

Today, services such as Dodgeball [6] utilize mobile phones to exchange information between users regarding their locality to identify opportunities to meet. Friends can find each other when they happen to be close by, and people can find others they have never met before who share their interests as well as their location. Mobile social networking services are successful primarily because they directly take advantage of the, often periodic, physical locality patterns of the users to improve their knowledge about opportunities to socialize.

An unexplored source of physical locality, which can be

taken advantage of for mobile social networking, is the daily roadway commute. A great number of people, especially in densely populated locales, spend one or more hours each day driving between home and the office. Since these commutes typically take people over highways and other frequented corridors, they are highly predictable and regular. Day after day, the same people travel along the same roadways at the same time. Therefore, the opportunity exists to form periodic virtual mobile communities. We call these virtual communities *Vehicular Social Networks* (VSNs).

The key property is that roadways, and the VSNs overlaid, provide a sufficient and regular concentration of people who may wish to socialize. People partake in all forms of entertainment, while commuting to and from work. They listen to music, listen to and participate in talk radio, chat with friends over cell phones, etc. In fact, in analogous commuting situations, e.g., on buses or trains, it is common for people to socialize and communicate with each other. We believe that a substantial opportunity exists for people to organize into social groups and expand their social connectedness while traveling on roadways, for the purposes of entertainment or to pass the time during their commute, or even to ask for help in case of emergency.

We envision three major reasons for VSN formation: (i) for entertainment, (ii) for utility, and (iii) for emergency purposes. Entertainment-based VSNs are formed for people to share common interests, for example, to discuss recent social and political issues, sports, or any other interesting and entertaining topics. Entertainment-based VSNs primarily serve to occupy people during their long and boring commutes, similar to the way radio programs provide entertainment. Utility-based VSNs are formed to facilitate non-essential, yet helpful connections to occur. For example, local towns or states might establish VSNs to discuss interesting local events or points of interest along popular highways, which travelers might use to query for advice in selecting a local restaurant or hotel. There might also be VSNs established to coordinate carpooling, to communicate roadway events (e.g., accidents, congestion, etc). Finally, VSNs can become particularly useful in the case of an emergency. Such emergency-based VSNs might serve as a way for people to ask for and provide assistance to each other in critical situations.

This paper presents a framework for building Vehicular Social Networks (VSNs) and, as a proof of concept, a VSN-based system, called RoadSpeak. RoadSpeak allows drivers

^{*}This work has been supported in part by the NSF grant CNS-0520123.

[†]Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. SocialNets'08, April 1, 2008, Glasgow, Scotland, UK
Copyright 2008 ACM ISBN 978-1-60558-124-8/08/04...\$5.00

to automatically join VSNs called *Voice Chat Groups* (VCGs) along popular highways and roadways. The primary goal of RoadSpeak is to facilitate communication between people who are already physically present in the same location, but who are limited in their ability to communicate with each other. In most social scenarios, for example, a cafe, nightclub, etc., people can easily communicate with each other. They are not restricted by their activity, as they are when they are driving on highways or other roadways. Today, if people wish to speak with each other while occupying a vehicle, they must either be collocated in the same vehicle, join a radio talk show, or initiate a direct connection to each other (e.g., make a cell phone call.) RoadSpeak takes advantage of the common situation drivers on the same road segments share, to allow opportunities for these people to socialize and communicate in an easy and natural manner.

RoadSpeak provides a simple asynchronous communication protocol for drivers to participate in the discussion in a given VCG. Users can both listen and speak to others in their VCG. The length of a voice message intervention is limited by the RoadSpeak client to a certain quantum of time. Once the intervention is completed, the corresponding voice message is sent to the RoadSpeak server. The server acts as the VCG coordinator, which admits participants, enforces the discussion protocol, collects and timestamps the drivers' interventions before delivering them in order to the entire VCG.

In our framework, a VSN such as a Voice Chat Group (VCG) is defined by a profile consisting of the triplet $\{time, location, interests\}$. Aside from interests, on which traditional social networks are formed, VSNs also include time and location in their definition, which may limit membership to specified roads and intervals of time. The VCG profile is defined by the group owner (which can be the manager, an organization, or simply a user), when the group is created. Users are admitted to groups based on matching their user profiles to the group's profile. Admission can be free or it can be restricted by certain requirements, which the user must prove, such as a location frequency threshold for a specified roadway segment, during a specified time slot. By limiting group membership based on the users' roadway recurrence frequency, VSN-based groups may be limited to commuters, the best target for this system, rather than including transient travelers.

To the best of our knowledge, this paper is the first to introduce the idea of Vehicular Social Networking. We present the design of a centralized system for automatically forming VSNs based on time, location, and interests. We describe a particular VSN system called RoadSpeak, a voice chat system for drivers on the road, which we preliminarily prototyped using 3G-based cellular communication.

2. ROADSPEAK SCENARIO

To illustrate our VSN vision, we describe a usage scenario for our RoadSpeak project. Joe would like to start participating in RoadSpeak during his daily commute. From his home computer, he goes to the RoadSpeak portal to create an account and to download the RoadSpeak client to his smart phone. While he is logged into the system from home, he browses through the various groups available along the route

he takes between home and work. He chooses to join a sports talk group active on the NJ Turnpike between 6 AM and 8 AM, since he is usually on the highway from 6:30 AM - 7:15 AM.

The next morning, Joe gets into his car, sets his smart phone in its cradle and puts on his hands-free kit. As he travels to work, the RoadSpeak client running on his smart phone tracks his location via the built-in GPS receiver, and when he enters the NJ Turnpike at 6:32 AM, the client contacts a RoadSpeak server and adds Joe to the group. Joe receives an audible alert from his phone, notifying him that he is joining the chat group. Joe introduces himself to the group and begins listening to the current sports talk voice chat messages that his RoadSpeak client is downloading from the server. These messages are similar to SMSs, or text-based chat, in the way they are delivered, but they are audible and are transmitted to everyone in the group.

After Joe listens, for a while, he intervenes and speaks to the group. The RoadSpeak client captures and transmits his voice message to the server. As the server continues to send out messages in the sequence it received them, it reaches Joe's message and distributes it to the other users currently participating in the VCG. Two people, Fran and Harry, hear Joe's question to the group and respond, about the same time. Their responses are ordered by the server based on their arrival time, and scheduled to be delivered in this order to everyone in the group. As Joe leaves the NJ Turnpike, the RoadSpeak client running on his smart phone, detects his departure and notifies him that he will leave the sports talk chat group, shortly. Joe transmits a farewell message to the group and continues along his route to work.

3. DESIGN

3.1 Vehicular Social Network Model

A Vehicular Social Network (VSN) is a social network of users who traverse popular roadways. This section defines our model of VSNs and describes the issues related to VSN creation, enforcement, privacy, and infrastructure requirements.

VSN Definition: A Vehicular Social Network is a social network that is enabled by: (i) spatial locality and (ii) temporal locality. Roadways naturally corral drivers into very regular and predictable patterns. Aside from very exceptional cases, vehicles are confined to travel on roadways, and there is a natural spatial locality enforced by the fact that drivers must do so.

There is also temporal locality in the way many drivers utilize roadways. For most commuters, this locality is brought about by various recurring daily deadlines. Commuters must arrive at work by some fixed time, and they are usually required to remain at work until some later fixed time.

Although spatial and temporal locality are enablers of VSNs, they are not the motivation for their formation. As with any social network, it is the social aspects that define the utility of the network and VSNs share the same goals as any other social network. They exist to allow people to form new connections, expand their social connectedness, and to support various social applications, based mostly upon their common

interests.

VSN Deployment: In order to realize a VSN, we must first answer some basic questions about their formulation and deployment. These questions can be grouped into three categories: (i) state maintenance, (ii) VSN centrality, and (iii) communication.

State related to VSN specifications and user profiles must be maintained in order to support user-defined criteria for VSN creation. This state can either be stored at central trusted servers or locally with users. Additionally, state can be made persistent, to support stable VSNs, or it can be considered transient, in the case of ad-hoc VSN formation.

VSNs may be formed with differing centrality. A location-centric VSN is one that is formed from the perspective of some fixed location. Membership in location-centric VSNs is restricted to those drivers matching some predefined social criteria and passing through predefined locations with some recurrence frequency. At the other end of the spectrum, are VSNs that are formulated from the driver’s perspective. A driver-centric VSN fixes each driver at its center, regardless of location, and connects others to the driver as he encounters them on the roadways. In a driver-centric VSN, the strength of connections between drivers is determined by their level of commonality in social interests as well as the frequency of their recurring encounters.

Finally, we must choose a model to support inter-vehicular communication. There are a number of models in existence today. Among them, are car-to-car approaches including Vehicular Ad-Hoc Networks, centralized communication infrastructure approaches such as cellular communication, and various hybrid approaches.

VSN Creation: A VSN is created by specifying its profile: (i) the time interval, (ii) location/road, (iii) and group’s interests. For example, a user might create a VSN of users for the time slot starting at 4 PM and ending at 6 PM. Similarly, a user could define a VSN for a specific section of a roadway, and this VSN would include any user who is currently located on that section of the roadway. Location properties could be specified using either coordinates (i.e., GPS values) or some logical description of the location (e.g., miles 501-576 on Route 1). Finally, VSNs can be specified using group’s interests. These are user-defined keywords that provide a meaningful definition of the purpose or collective interests of the users included in the VSN.

Location Validation: Although part of the definition, location may or may not be enforced by the VSN. When it is not enforced, a user that provides false location information would be able to participate in a VSN, contrary to the VSN specification properties. When location is enforced, the VSN must use a method to validate the location information provided, which could be either (i) direct validation, for example, using road-side infrastructure to periodically sample user positions, or (ii) cross-validation performed between users.

Privacy: Since a VSN may use per-driver time-location information, there is a privacy issue beyond what traditional social networks must handle. The privacy requirements for VSNs will differ based on their specification properties. A VSN model that includes cross-validation will directly expose user

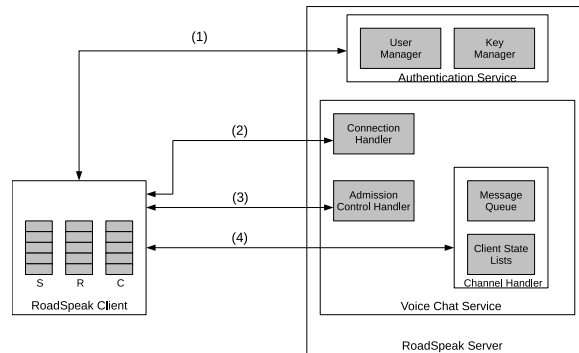


Figure 1: RoadSpeak vehicular voice chat application overview.

location information to users around them. Therefore, it may be important to include mechanisms to anonymize the data with respect to the individual users, while only exposing location information to a central authority [17]. Alternatively, a model that does not include cross-validation would only require users to expose location information to a centralized, trusted, authority or to an anonymizer. In either case, this data must be handled as highly sensitive to its users and the system must not directly leak this information between users or allow indirect inferences based on the information that is exposed.

Infrastructure Requirements: There is a broad range of infrastructure that can be assumed to support Vehicular Social Networks. First, we assume that vehicles are equipped with some form of wireless communication. This may include wireless-enabled smart phones (e.g., 3G data access, WiFi, WiMAX, etc.), or vehicles may be equipped with embedded vehicular computing systems that support car-to-car communication (C2C). External to the vehicles, we can either assume a fully deployed road-side infrastructure, which supports car-to-infrastructure (C2I) communication, no infrastructure, or some deployment model in between, such as a CarTel-like [10] in-situ organic WiFi network approach. The infrastructure assumptions we make will have a substantial impact on the design of any VSN-based system.

3.2 RoadSpeak Voice Chat

Figure 1 shows an overview of the RoadSpeak vehicular voice chat application. Central to RoadSpeak is the concept of a Voice Chat Group (VCG). We define a VCG to be a VSN composed of users who wish to communicate with each other in near real time over a voice communication system. This can be done as a combination of voice over IP (VOIP) and Internet chat rooms. The closest applications similar to this are Internet voice chat systems, such as Ventrilo [21], TeamSpeak [19], or Asterisk [1]. These systems are designed to handle one-to-one or one-to-many voice communication between users over more stable Internet connections, and were not intended to operate as VSNs overlaid on top of vehicular networks.

3.2.1 RoadSpeak Overview

User Creation and Authentication: Users download a client to their smart phones (or embedded vehicular PCs) and use the client to create an account in the system (see (1) in Figure 1). To access groups, users login through the client. Both user

API Function	Function Description
<i>createUser</i> (< username >, < cprofile >)	Create user named < username > in the VSN and set its initial profile to < cprofile >.
<i>getUserProfile</i> (< username >)	Fetches the user profile data for < username >.
<i>setUserProfile</i> (< username >, < cprofile >)	Updates the user profile data for < username > to < cprofile >.
<i>createGroup</i> (< groupname >, < gprofile >)	Create a VSN group named < groupname > in the VSN and set its initial profile to < gprofile >.
<i>getGroupProfile</i> (< groupname >)	Fetches the group profile data for < groupname >.
<i>setGroupProfile</i> (< groupname >, < gprofile >)	Updates the group profile data for < groupname > to < gprofile >.
<i>enumGroups</i> ()	Enumerates the current list of existing groups.
<i>enumUsers</i> (< groupname >)	Enumerates the current list of users connected to the group < groupname >.
<i>lookupGroupKey</i> (< groupname >)	Fetches the current public group key for < groupname > group.
<i>generateGroupKey</i> (< groupname >)	Generates a new group key pair for group < groupname >.
<i>getModList</i> (< groupname >)	Fetches the current list of moderated users connected to the group < groupname >.
<i>setModList</i> (< groupname >, < modlist >)	Updates the current list of moderated users for group < groupname >.

Figure 2: RoadSpeak API.

account creation and login authentication are handled over a secure web protocol (i.e., HTTPS) by the User Manager. At the time of account creation, the client generates an asymmetric cryptographic key pair (using PKI) and registers its public key (K_{PUB}) with the User Manager.

Once a user has been authenticated to the system, they are free to create new groups, to invite other users to the groups they own, or to request membership to groups owned by other users. Group invitations and join requests are forwarded to other users (potential participants and group owners, respectively) to either accept or decline. In the case where a group is owned by multiple users, the owners vote to either allow or deny new users into the group.

Group Membership: Voice chat groups can be defined as public or private. Anyone is free to join a public group so long as they satisfy the group properties, while private groups require the explicit permission of the group owner. Groups can also be defined as open or moderated. In an open group, anyone can freely communicate with members of the group, whereas in a moderated group, permission to communicate is explicitly granted by the group owner and can be revoked at any time. Finally, a group can be hidden or advertised. Hidden groups are only advertised to those users who are allowed to participate in the groups.

Admission to a group is handled by the Group Admission Manager. To be admitted to a group, a user submits his profile. This data allows the Group Admission Manager to verify that the user attempting to join the group, actually satisfies the group's properties. If the Group Admission Manager validates the user, and allows her to join the group, the group key (K_G) is transmitted back to the client to be used by the client in the future for that group.

Connection Handler and Admission Control: A user chooses which group he wishes to participate in, once he has logged into the RoadSpeak system (through the local client on his smart phone). A client may only enter one RoadSpeak chat group at a time. When a RoadSpeak client connects to a RoadSpeak chat application server (see (2) in Figure 1), the Connection Handler spawns a new Admission Control Handler (ACH) and hands the connection to the ACH to perform admission control (see (3) in Figure 1).

To initiate Admission Control, a RoadSpeak chat client transmits the the group name, the user's name, profile, and the SHA1 hash of the group key ($sha1(K_G)$) to a vehicular chat server. This information is encrypted by the client, using its

private key (K_{PRIV}), prior to being sent to the server. The server validates the user based on the user's information and the hash of the group key submitted by the user. The ACH then hands the client connection off to the Channel Handler for the requested VSN group (see (4) in Figure 1).

Channel Handler: The Channel Handler spawns one thread per client connection. All of the Channel Handler threads share a global data structure, called the Message Queue. The Message Queue maintains a list of message chunk rows (< message #, chunk #, chunk data, username, status flags >). Each row represents a chunk of a voice chat message that has been submitted to the server by a user participating in the group chat channel. Messages in this list are maintained in the order in which they were received, such that they are sent out to each client, one message at a time in the same order.

Each Channel Handler thread also maintains a local data structure called the Client State List, which maintains a set of client state rows (< message #, chunk # >). There is one row per message being sent to the client. Each row tracks the last message chunk sent to and acknowledged by the client.

Message Handling: Each RoadSpeak client maintains a set of three message queues: (i) send queue, (ii) receive queue, and (iii) control queue. These queues are used to buffer outgoing and incoming voice chat messages, and for flow control.

Send handling. Voice messages are captured by the client's Voice Capture thread (e.g., using the smart phone microphone). Captured voice messages are then broken into fixed sized message chunks, encrypted with the group key (K_G), and placed into the client send queue. The client's Send Handler thread streams message chunks to the server, and the server responds with acknowledgments.

Receive handling. Once the server receives a message chunk, it places the message chunk on the global Message Queue. Each Channel Handler thread accesses this queue and streams chunks to its associated client. The client's Receive Handler thread receives the message chunk, places it into the client's receive queue, and sends an acknowledgment back to the server.

Message sequencing and flow control. The client's Voice Playback thread is notified by the receive thread whenever all chunks for the lowest numbered message have been received and the message is complete. The Voice Playback thread then removes chunks, one at a time, from the receive queue, decrypts them using the group key (K_G), and streams them in order to the user (e.g., using the smart phone's speaker). In this way, messages are played back to all users in the same

sequence that they were received by the server. Flow control is handled in a similar fashion to the TCP protocol.

The server Channel Handler threads can also send *pause* and *resume* messages to their associated clients. These higher-priority messages are placed on the clients' control queue when received and are handled by the receive thread. When a *pause* is received from the server, the client receive thread notifies the client send thread, to pause all chunk streaming from the client to the server. Upon receiving a *resume*, the client receive thread notifies the client send thread to begin chunk streaming, once again.

User Event Notification: There are a number of conditions that must be reported to the user as feedback, during a voice chat session. For each of these, the client provides audible and visual feedback to the user.

Application Programming Interface: Figure 2 lists the API that the RoadSpeak server exports for use by clients. This API can also be used to extend RoadSpeak clients to provide enhanced functionality.

3.2.2 Time-Location Traces

Location traces can be collected by having clients perform discovery of each other over WiFi, and submit the traces as proofs to the server for cross-validation. This method assumes both the existence of WiFi at each client, and a large enough concentration of clients such that enough discovery can occur to perform the cross-validation.

Time-Location Trace Collection: Time-location traces are collected by individual clients for drivers as they travel along roadway segments included in the system. Location is determined through the use of GPS. For any such roadway segment, a RoadSpeak client polls for the presence of other participating drivers in the same location, utilizing inter-vehicular wireless communication. In addition to the information gathered about others, each client keeps the same information for itself. Whenever a client leaves a known road segment, it transmits updated traces for the roadway segment it has just left, to the server.

Trace Submission and Cross-Validation: Once a client has joined a RoadSpeak group, it may be required to submit time-location traces in order to be admitted to that group. For example, a group may require that its members travel on the relevant roadway segment(s) for the group 4 out of 5 weekdays, each week, in order to be admitted to the chat group.

Once a trace has been submitted to a RoadSpeak server by a client, the server performs three checks on the trace. First, the server decrypts the trace using the client's registered public key (K_{PUB}). Next, the server check that the trace satisfies the matching criteria for the group. Finally, the server performs cross-validation on the traces using historical data stored from previous trace submissions from other clients to the server.

Cross-validation is utilized to verify that the traces provided by any client are authentic. To do so, a server searches its historical trace data set to verify that other clients who were physically located on the same roadway as the submitting client, discovered the submitting client. In this way, each client provides proof for the clients it discovers.

There are two issues in cross-validation that must be han-

dled. First, it is possible that a client submits traces that cannot be completely cross-validated. This may occur due to issues in wireless communication, where Client A discovers Client B, but Client B does not discover Client A. Therefore, we only require a client to meet a minimum threshold in terms of the number of clients that validates the submitting client's presence. Second, colluding clients could cross-validate each other, thereby circumventing the admission control policy. At present, we consider this to be outside the scope of this work, but acknowledge the issue, nonetheless.

4. IMPLEMENTATION AND PRELIMINARY RESULTS

We have implemented the RoadSpeak system using Java and tested it using laptops with Verizon EVDO PC cards, which provide 3G connectivity. We deployed two clients that periodically send their location and neighborhood traces to the server over 3G. One of the clients forms a RoadSpeak group, and the other client shared its route, and hence becomes a part of the group, and the two clients can send messages to each other.

Our preliminary results show that the 3G network bandwidth is able to successfully transfer voice messages at a sampling rate of 32 Kbps. We chose 32 Kbps, as it represents the worst case sampling rate for the system, since it is common for voice sampling rates to be substantially slower (typically 8 Kbps). However, owing to the unstable 3G connection, at times, we noticed large spikes in propagation delay of these voice messages. Since RoadSpeak is less sensitive to packet loss, but typically very sensitive to delays, we are working on an RTP/RTCP based implementation of RoadSpeak for exchanging streaming voice. We are also working on more detailed scalability experiments to study the effect of a large number of clients on the 3G network, as well as, the server and the client queuing delay.

5. RELATED WORK

Traditional online social networks such as Facebook [8] and LinkedIn [12] are essentially a virtual view of physical communities, however they are agnostic to location and proximity. Recently, there has been an increasing trend towards mobile social networks that can help people connect with their physical communities and surroundings. Social Net [20] infers shared interests between people by storing the IDs of the nearby devices and analyzing this co-location data collected over a long period of time. Social Serendipity [7] is another similar mobile phone-based application using Bluetooth and a database of user profiles to recommend face to face interactions between nearby users who share common preferences. MobiSoc [3] is a middleware that enables formation of mobile social networks and provides a platform for managing the social state of physical communities.

Vehicular networks have become increasingly popular in the recent past due to the interest from government as well as auto manufacturers. Various types of applications that make use of vehicular networks have been proposed, ranging from improving safety [14], traffic management [23], emergency

management [18], urban sensing [10], multimedia sharing [16], gaming [15], to disseminating advertisements to passengers [11]. All existing work in the field of vehicular networks consider links between vehicles to be spontaneously formed as a result of the vehicles being in wireless range of each other.

Vehicular networks are a special case of mobile ad-hoc networks (MANETs), as well as, delay tolerant networks (DTNs). In the field of MANETs and DTNs, various works have proposed applying knowledge about social networks in order to improve networking protocol performance. In [5], the authors propose a routing algorithm for MANETs that, based on the small world phenomenon seen in social networks, identifies bridge nodes for routing based on their centrality characteristics. [13] shows how knowledge of nodes' social interactions can help improve performance of mobile systems. [2] studies the impact of users' social relationships and movement patterns on the performance of routing protocols in opportunistic networks. Similarly, a social network based overlay for publish-subscribe communication in DTNs has been studied in [22]. None of the existing works, however, consider vehicular networks as the underlying network that can make use of a social network overlay.

A number of projects have described ways of spontaneously creating mobile social networks in different network settings. [4] describes a middleware for managing location-dependent relations in mobile social networks and for propagating the social networks' visibility up to the application layer. The Huggle project has performed [9] extensive studies of human mobility traces in opportunistic networks and extracted information related to clustering, and community structure.

Our work is distinct from all the earlier work in that RoadSpeak is the first application built on a vehicular network with an underlying social network overlay. We believe that our work will inspire existing vehicular applications to leverage the underlying social connections that exist in vehicular networks and form vehicular social networks.

6. CONCLUSION AND FUTURE WORK

We have proposed the concept of Vehicular Social Networks, and a framework for them. Using this framework different people traveling along the same roadways at the same time can form virtual mobile communities. We presented the design of RoadSpeak, an inter-vehicle voice chat system that allows users to automatically join Voice Chat Groups along these roadways. We presented mechanisms to discover and form these groups, which utilize both time and locality, along with user interests. As future work, we plan to evaluate our RoadSpeak prototype on the roads with a larger number of cars.

7. REFERENCES

- [1] Asterisk. <http://www.asterisk.org/>.
- [2] C. Boldrini, M. Conti, and A. Passarella. Impact of Social Mobility on Routing Protocols for Opportunistic Networks. In *Proc. of AOC Workshop*, 2007.
- [3] C. Borcea, A. Gupta, A. Kalra, Q. Jones, and L. Iftode. The MobiSoC Middleware for Mobile Social Computing: Challenges, Design, and Early Experiences. In *Proc. of Mobilware*, 2008.
- [4] D. Bottazzi, R. Montanari, and A. Toninelli. Context-Aware Middleware for Anytime, Anywhere Social Networks. In *IEEE Intelligent Systems Volume 22, Issue 5, Pages 23-32*, 2007.
- [5] E. Daly and M. Haahr. Social Network Analysis for Routing in Disconnected Delay-Tolerant MANETs. In *Proc. of MobiHoc*, 2007.
- [6] Dodgeball. <http://www.dodgeball.com/>.
- [7] N. Eagle and A. Pentland. Social Serendipity: Mobilizing Social Software. In *Pervasive Computing, IEEE, volume 4(2), pages 28-34*, 2005.
- [8] Facebook. <http://www.facebook.com/>.
- [9] P. Hui, E. Yoneki, S. Chan, and J. Crowcroft. Distributed Community Detection in Delay Tolerant Networks. In *Proc. of MOBIARCH Workshop*, 2007.
- [10] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *Proc. of SenSys*, 2006.
- [11] S.-B. Lee, G. Pan, J.-S. Park, and M. Gerla. Secure Incentives for Commercial Ad Dissemination in Vehicular Networks. In *Proc. of MobiHoc*, 2007.
- [12] LinkedIn. <http://www.linkedin.com/>.
- [13] A. G. Miklas, K. K. Gollu, K. K. W. Chan, S. Saroiu, K. P. Gummadi, and E. de Lara. Exploiting Social Interactions in Mobile Systems. In *Proc. of Ubicomp*, 2007.
- [14] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. TrafficView: A Scalable Traffic Monitoring System. In *Proc. of MDM*, 2004.
- [15] C. E. Palazzi, M. Roccetti, S. Ferretti, G. Pau, and M. Gerla. Online Games on Wheels: Fast Game Event Delivery in Vehicular Ad-hoc Networks. In *Proc. of V2VCOM*, 2007.
- [16] J.-S. Park, J.-S. Park, J. Yeh, G. Pau, and M. Gerla. CodeTorrent: Content Distribution using Network Coding in VANETs. In *Proc. of MobiShare*, 2006.
- [17] N. Ravi, M. Gruteser, and L. Iftode. Non-Inference: An Information Flow Control Model for Location-Based Services. In *Proc. of MOBIQUITOUS*, 2006.
- [18] M. Roccetti, M. Gerla, C. E. Palazzi, S. Ferretti, and G. Pau. First Responders' Crystal Ball: How to Scry the Emergency from a Remote Vehicle. In *Proc. of IPCCC 2007/NetCri07*, 2007.
- [19] Teamspeak. <http://www.goteamspeak.com/>.
- [20] M. Terry, E. D. Mynatt, K. Ryall, and D. Leigh. Social Net: Using Patterns of Physical Proximity Over Time to Infer Shared Interests. In *Proc. of CHI*, 2002.
- [21] Ventrilo. <http://www.ventrilo.com/>.
- [22] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft. A Socio-Aware Overlay for Publish/Subscribe Communication in Delay Tolerant Networks. In *Proc. of ACM/IEEE MSWiM*, 2007.
- [23] J. Yoon, B. Noble, and M. Liu. Surface Street Traffic Estimation. In *Proc. of MobiSys*, 2007.