LoadingZones: Leveraging Street Parking to Enable Vehicular Internet Access

Riccardo Crepaldi, Ryan Beavers, Braden Ehrat, Jason Sze, Matt Jaeger,
Steven Biersteker, Robin Kravets
Dept of Computer Science, University of Illinois
Urbana-Champaign, IL, US
{rcrepal2,beavers1,ehrat1,sze1,jaeger3,bierste1,rhk}@illinois.edu

Categories and Subject Descriptors

C.2.1 [**COMMUNICATION NETWORKS**]: Network Architecture and Design; C.2.4 [**COMMUNICATION NETWORKS**]: Distributed Systems

Keywords

VANET, Opportunistic, Internet Access, Ad-Hoc

Introduction

Wherever we are, we assume that we will be able to access the network to read our email, visit a website or download our favorite music, and to do this we often rely on the cellular infrastructure, especially when we are on the streets, walking or driving. With our mobile devices in our hands, we perceive network access as an ubiquitous service. However, we often overlook that the place where many people spend the majority of their mobile time, the car, is a lot more than a means of transportation. Not only are cars an extremely advanced mechanical tool, they are also equipped with cutting-edge communication and sensing technology. For the first time, users are not carrying their mobile computing device with them but are being carried by the device itself, and at ever increasing speeds. This change encompasses more than a simple size difference: it relaxes some of the constraints of mobile networks, and defines a set of new ones.

When we drive, our expectations about the connectivity we can achieve are very different compared to when we are inside a building. While ubiquitous connectivity has come to users in cars in the form of 3G and 4G networks, such connectivity is extremely unpredictable and has limited reliability and bandwidth. Essentially, although cellular networks at first glance offer an attractive backhaul medium, the reality is that the load that mobile users currently put on it already pushes the limits of the infrastructure. The sheer demand of today's 3G users, who put all of their data over the 3G networks, overloads the limited bandwidth and so adversely affects the original design goals. If more delay-tolerant communication, such as sensed data, sending and retrieving emails and other delay-tolerant personal communications, prefetching web content for off-line browsing, were supported by other means, the 3G network could operate far more effectively. We have lost the potential benefits of more efficient data transfer in a world where data availability has exploded, including hundreds of sensor readings per minute in each car, the need for constantly faster acquisition of traffic information that must be delivered to navigation devices within the cars, and targeted multimedia content, such as advertisements or information that users store in the cloud expecting to have ubiquitous access.

To enable the full potential of cars as mobile devices, we must be able to provide a network that insures predictability, reliability and high capacity. However, it is important to note that each of these requirements, despite all being necessary, may not always be needed at the same time. A predictable and reliable connection is required to enable time-sensitive operations such as credit card purchases, while a high bandwidth connection is necessary to transfer car-generated and other large chunks of data that can tolerate some delays and unpredictability. 3G networks were designed to satisfy part of this picture: predictability and reliability. When not overloaded, the cellular network is the best medium for timely responses. One step taken to relieve it of the unnecessary load is that of extending the in-building model of mobile handoffs to cars on the street [1, 2], allowing these cars to connect to any in-building access point (AP). While this seems to be a natural extension to mobile communications, a number of problems arise in a vehicular setting that have only partially been identified and addressed in the past. There are three major obstacles to such an approach:

Delay. Setting up an 802.11 link requires several steps: scanning available channels for AP beacons, associating to an AP, which often requires authentication, and finally running a DHCP request to configure the network parameters.

Quality. The connection between a vehicle and an indoor AP is impaired by several factors: the signal has to pass through thick walls, the AP location is not designed to cover the street, and contention happens with regular indoor users.

Duration. The coverage range of an AP depends on obstacles and transmit power. The former are inevitable, the latter limited by regulation. A car driving on the street passes through the area in which the indoor AP's signal is strong enough in only a few seconds.

Combined, these three factors have a deadly effect on the data transfer potential between cars and in-building APs. However, tackled individually, it is possible to provide the desired predictability, reliability and high capacity. While researchers have tackled the connection problem [1, 3], the only real solution to the other two problems requires bringing the APs closer to the cars. Researchers and companies have tried to deploy outdoor APs [4] and have failed simply due to the high cost of deployment. While we believe these outdoor AP approaches have merit, we take a novel approach to this problem that leverages inter-vehicle communications and enables individual cars to play the role of a relay to an AP.

Our solution to this problem of vehicle connectivity is LoadingZones, a communication system that uses parked cars as relay agents between moving vehicles and APs. The key to LoadingZones is a divide-and-conquer approach. By separating a single connection into two, moving vehicle to parked car and parked car

MobiCom'12, August 22–26, 2012, Istanbul, Turkey.
ACM 978-1-4503-1159-5/12/08.

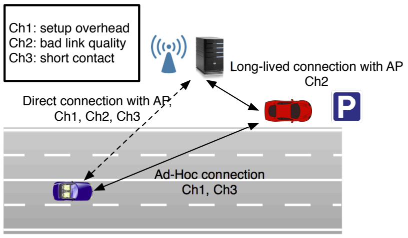

Figure 1: The two-hop approach of LoadingZones separates the challenges that make a direct link to an AP almost unusable.

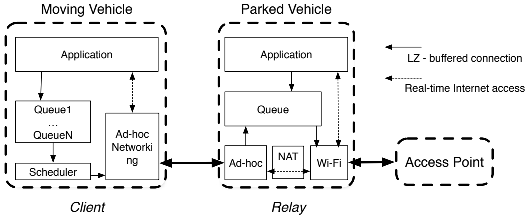

Figure 2: LoadingZones multi-layered architecture

to AP, as shown in Figure 1, LoadingZones isolates the detrimental effects of slow connection times and poor throughput to the car. The role of parked cars has been recognized as very useful in maintaining the DTN fabric of vehicular networks [5], but none of the existing solutions identified the primary role that parked vehicles can have in alleviating most of the issues that limit Internet connectivity for VANETs. Despite being stationary, they see both the vehicular network and the Internet from a vantage point, and can act as bridges between them. When parked in front of a building equipped with Wi-Fi, a vehicle can connect with it for a long time. The still long setup delay can be amortized over a much longer connection that may last several minutes if not many hours. The channel between the parked car and the AP still suffers from low quality, but the steady nature of the link helps alleviating this issue as well. The long life of this link, despite the poor quality and the setup overhead enables uploading and downloading large volumes of data to and from the Internet. A moving car consequently changes its search target. Not anymore an available AP, but parked car instead, that acts as a relay. Such a dedicated connection to the relay vehicle can be designed to significantly reduce setup delay. The channel quality between the moving car and the parked car is also better, and the contact duration is longer between two vehicles, due to a better position compared to that of an indoor AP, and lack of obstacles.

LoadingZones (LZ) is a system that can work as a companion to the cellular infrastructure with a dual benefit. It is a system that can support the increasing demand of an emerging scenario such as VANETS, and at the same time allows a cheap alternative to subscriber services for *delay-tolerant* data. This, in turn, releases the load on the current infrastructure, potentially providing a better service for everyone. Our work constitutes an initial effort to understand their potential and their limitations using a real implementation. In terms of throughput, our experimental results show that LoadingZones enables the transfer of several megabytes of data during a brief contact between a moving and a parked car. When the parked car is used as a relay, LoadingZones enables an increase in throughput of more than 100% compared to a direct connection to indoor APs even when the channel setup overhead is not considered. In the end, LoadingZones improves throughput, reduces user cost and results in a less loaded cellular infrastructure, ready to better serve users that require real-time access.

LoadingZones

LoadingZones (LZ) provides Internet connectivity to moving cars in short bursts as they pass by cooperating parked cars. Since applications may have diverse requirements, LZ's design enables two communication paradigms: *interactive* communication that supports short-lived web browsing or small data transfers such as sending and receiving email, and *bulk* communication for transferring larger data, although downloads are only available when prefetching content at a relay is supported. The application determines the specific style of communication, but LZ hides the complexity of supporting the distinct communication styles from the application.

System Design. The key design feature of LZ is a clear distinction between the role of a moving vehicle and that of a parked one. When a moving car has an *interactive* communication request, it initiates the search for a parked car to connect to. The parked vehicle acts as a relay to the Internet for the moving vehicle for as long as that connection lasts. Since *bulk* communication does not require any responses, LZ includes a local queuing system that resides locally on each vehicle to store the bulk data for uploading (or downloading in the case of prefetched data at a parked car). In the queuing system, a queue manager establishes a new queue for each application with bulk data. Applications can then continuously enqueue their data and continue their normal operations. This time when a connection between a moving and a parked vehicle happens, the moving car starts sending data to the parked car. The transmission order is determined by the *scheduler*, which can implement specific policies to guarantee fairness across applications. Since the channel between the vehicles generally has higher throughput than the channel between the parked car and the indoor AP, the *relay* does not forward the data immediately, but caches it in its own queue first. In this way, for the duration of the connection, the Internet connection is free for servicing *interactive* requests. A similar procedure happens in the opposite direction, with prefetched data such as map updates or local advertisements.

Implementation. We implemented a prototype of LZ using the Illinois Vehicular Project (IVP) [6] In IVP, devices are being installed in University of Illinois service vehicles that normally operate within the campus boundaries. The university campus, as with many other campuses worldwide, has good Wi-Fi coverage that includes all university-owned buildings, and extends to a large portion of the outdoor campus area.

Connectivity is automated by a number of scripts running in the background. First, the GPS receiver measures location and current speed to determine whether the vehicle is parked or moving. In a parked vehicle, a background process tries to locate the campus network and associate with it. Once associated, the parked vehicle broadcasts a beacon over the car-to-car interface. Upon establishing a connection with a moving car, the *interactive* connectivity is supported using NAT, IP forwarding and the `iptables` MASQUERADE policy. For *bulk* communication, the local queuing system simply binds to a local socket that applications can connect to and send messages to at any time. Applications must be specifically programmed to send messages to the queue. Beyond the local queue, however, applications do not have to implement any particular intelligence for sending packets through the network. They simply specify during enqueue the recipient's address and port.

The queue and message passing system is built to handle both small messages (e.g. sensor readings) and large messages (e.g., images and videos). For fairness, an independent queue is created for

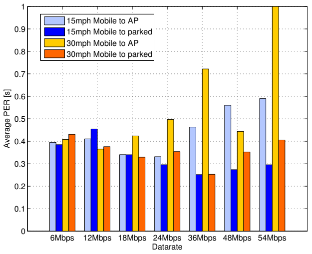

Figure 3: Average PER during a contact.

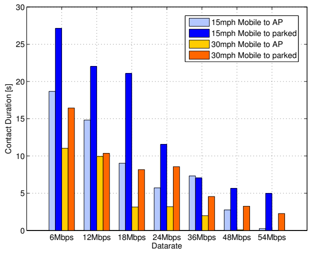

Figure 4: Link duration to an AP Vs a link between vehicles.

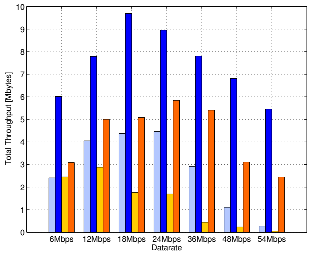

Figure 5: Total traffic delivered per contact.

data from each individual application. When a car-to-car connection is available, a scheduler is triggered and starts pulling packets from the queues according to its policy (the default is Round Robin for our prototype). The messages are then transferred and stored in the queue of the parked vehicle, where another scheduler dequeues and transmits them to the AP once the network is idle (i.e., the *interactive* applications do not require any bandwidth). Upon the messages' arrival at the destination, the data is simply passed to the specified port.

Evaluation

We compared LZ's performance to using a direct connection to an indoor AP. In both cases, we ran experiments with the moving car driving at 15 and 30mph to evaluate the impact of speed on system performance. In high mobility scenarios, automatic rate adaptation algorithms are known to perform poorly due to slow reactivity, which keeps them in a less than optimal state all of the time. For this reason, we turned off rate adaptation.

Delay. Connection setup time accounts for three operations: scanning time, association and authentication, and IP setup. The impact of security is substantial due to the additional overhead required. Our experiments show that the use of WPA2 accounts for more than 33% of the total connection time. Our experiments show that the average time of a successful connection setup, including association, authentication and DHCP, is not significantly influenced by mobility. Regardles of whether the car was parked or moving, the average time is $8.8s$, with a variance of less than $2s$. However, these values are computed only for the successful connections, ignoring timeouts. In our experiments, only 54% of the connection attempts worked on the first try. This is mainly due to the poor channel quality. 18% succeeded after the first timeout, 11% after 2 timeouts, and the remaining required even more retries. The key observation here is that parked vehicles have enough time to eventually recover from these timeouts. A moving vehicle, on the other hand, has a limited time to connect to to an AP before leaving its coverage range. As a result, for vehicles moving at $15mph$, the success rate was approximately 90% and decreased to 57.14% when for 30mph.

Quality. To measure the quality, and next duration, of a connection, we isolated the effect of association by establishing an ad hoc link between two IVP devices. We then placed one of them in a car, and drove it at different speeds. We measured the packet reception rate at the other device placing it inside a building first, to emulate the channel to an infrastructure AP, as in the single-hop approach. We then repeated our experiments with the second device outdoors, in a parking spot on the street shoulder, to emulate the channel with a parked vehicle. We measured the packet error rate (PER) for each link from the moment the first message was received to the moment the last one reached the destination. We define a contact as the interval between the moment the success rate is higher than 0.3 for the first time, until it reaches the same value for the last time before the connection fades out completely.

At low data rates, the robustness of the modulation scheme and coding dominates the effect of mobility or AP placement, and all configurations have a similar PER (see Figure 3). However, as data rate increases, more aggressive modulation schemes are required, and the robustness of coding is sacrificed to improve the bandwidth. In these configurations, speed has a severe impact on the channel quality. This impact is higher for connections to an indoor AP.

Duration. Channel quality considered alone is not representative of contact performance, since it does not indicate how long the contact was. As expected, faster speeds result in shorter contacts, and the line-of-sight channel with a parked vehicle enables contacts that can last more than twice as long as those with an indoor AP at the same data rate (see Figure 4).

Throughput. As always, throughput must be the deciding metric. To understand the total amount of data that can ideally be transferred using the LZ as compared to the single-hop approach, we evaluate the aggregate effects of quality and duration (see Figure 5). The higher throughput to the parked vehicles is a consequence of comparable or lower PER and longer contact durations. Our experimental results show that LZ enables opportunistic data delivery of large volumes of data over a single contact, and unleashes the potential role of VANETs as an ubiquitous opportunistic mobile network. LoadingZones is a first effort in understanding the potential of cooperation between vehicles to expand the range of cheap, high bandwidth links that are so far confined to indoor environments.

1. REFERENCES

[1] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden, "CarTel: a distributed mobile sensor computing system," *ACM SenSys*, 2006.

[2] Z. Zheng, P. Sinha, and S. Kumar, "Alpha Coverage: Bounding the Interconnection Gap for Vehicular Internet Access," *IEEE INFOCOM*, 2009.

[3] J. Eriksson, H. Balakrishnan, and S. Madden, "Cabernet: vehicular content delivery using wifi," 2008.

[4] R. Frenkiel, B. Badrinath, J. Borres, and R. Yates, "The infostations challenge: balancing cost and ubiquity in delivering wireless data," *IEEE PerCom*, 2000.

[5] N. Liu, M. Liu, W. Lou, G. Chen, and J. Cao, "PVA in VANETs: Stopped cars are not silent," *IEEE INFOCOM*, 2011.

[6] R. Crepaldi, R. Beavers, B. Ehrat, and R. Kravets, "Illinois vehicular project, live data sampling and opportunistic internet connectivity," *ACM MobiOpp*, 2012.