

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

1. PROBLEM AND MOTIVATION

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. 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.

When we drive, our expectations about the connectivity we can achieve are very different compared to when we are inside a building. On the road, Internet access is limited by high cost and low bandwidth. Personal content, such as email and short messages, can be accessed with long latency using mobile phones, enabling a somewhat limited but acceptable web browsing experience. However, we have lost the potential benefits of more efficient data transfer in a world where data availability has exploded and users expect to have ubiquitous access to it. Enabling all of this together would either bring down the current infrastructure, or require the deployment of an ever more expensive infrastructure.

Treated as a large smartphone, communication to and from our cars is dumped on the already overtaxed cellular infrastructure. But again, we are overlooking the capabilities of these new generation of cars and the role they should play in this extremely large scale network. When not overloaded, the cellular network is the best medium for timely responses. However, many applications may be able to handle more opportunistic services. One step in this direction is to extend the in-building model of mobile handoffs to cars on the street [1, 2, 3], allowing these cars to connect to any in-building access point (AP) they can reach. 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:

- Connecting to an AP is slow and even slower with secured networks, which are becoming increasingly prevalent.
- The signal that an AP leaks onto the street is generally of poor quality, resulting in a similarly low data rate.
- The link lifespan is very short, due to the high mobility of the vehicles and limited Wi-Fi footprint of the in-building APs on the street.

Combined, these three factors have a deadly effect on the data transfer potential between cars and in-building APs. While researchers have tackled the connection problem individually [1, 4], 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 [5, 6] 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 to AP, LoadingZones isolates the detrimental effects of slow connection times and poor throughput to the car. A parked car still suffers from slow connection times and a poor channel, but the contact duration with an AP is very long. A moving vehicle instead can only communicate with a parked vehicle for a short time, but the channel setup is quicker and channel quality higher.

When parked vehicles participate in the network, the whole system performance benefits from their unique properties. Our work constitutes an 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. However, these improvements do not come for free. Parked vehicles are actually energy-constrained, especially when parked for long periods of time, and particular care must be used in deciding when LoadingZones should be active, and when instead energy should be saved.

2. BACKGROUND AND RELATED WORK

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. For example, in Chicago there are on average 2500 vehicles per km^2 . If only 10% of the users in those vehicles are generating data, the network must be able to support 250 data streams per km^2 . With this density, a typical urban 3G basestation covers approximately 100 users. Given an upload capacity of around 500Kbps per basestation, that leaves only 5Kbps per device if all users access the network concurrently. Users of such systems must tolerate long delays even before a simple webpage is fully loaded and in the end they limit Internet access to the strictly necessary tasks.

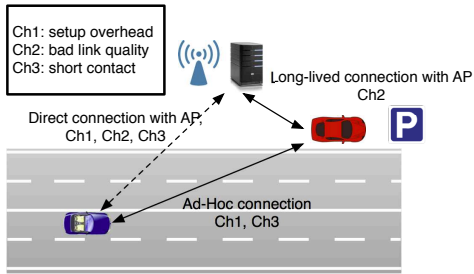


Figure 1: The two-hop approach of LoadingZones separates the challenges that make a direct link to an AP almost unusable.

To enable the full potential of cars as mobile devices, we must be able to provide a network that supports predictable, reliable and high capacity communication. 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. However, 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, delay-tolerant personal communications and prefetching web content for off-line browsing, were supported by other means, the 3G network could operate far more effectively.

Several systems have been trying to address the problem of limited bandwidth by using a traditional 802.11 approach instead of 3G or equivalent networks. Some solutions, such as [5, 6] rely on a specifically deployed infrastructure to service the vehicles on the road. These networks provide high capacity, but demand an unacceptable cost to deploy the required infrastructure [7]. On the other hand, solutions such as Cartel [1] used a completely opportunistic approach by locating and associating with open access points. The result is an extremely cheap solution, since it does not require any additional infrastructure. However, such solutions only tackle part of the problem. Essentially, Cartel and the other approaches [2, 3] that propose a single-hop link to available access points provide the means for obtaining a cheap connection, but fail to satisfy the high bandwidth requirement. This is due to the effect of three challenges:

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 a wireless 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.

Essentially, our experimental results show that tackled together, these problems can only be solved by deploying a more expensive infrastructure. However, tackled individually, it is possible to provide the desired predictability, reliability and high capacity.

Further reducing authentication delay and adding more fixed infrastructure are not options, since they would require changing server-

side configuration protocols, which network administrators are not likely to support since their networks are not primarily intended to service vehicular networks. However, it is interesting to consider the extended role of parked vehicles in this problem. Although these vehicles certainly act as data sources, they are also fully functional networking nodes that support vehicle-to-vehicle communication. And, incidentally, are typically closer to the moving cars on the street. While ad-hoc connectivity between cars has been proposed to avoid overloading the infrastructure [8, 9, 10], most solutions have focused ad-hoc or DTN systems for inter-vehicular communication. While some of these solutions do support communication between the moving vehicles and the Internet [1, 11, 12, 13], they try to achieve the best of a connection to an AP by either proposing proper AP deployment techniques [12], or optimizing the connectivity time and reducing the overhead at the client side [4]. To extend even further the coverage range, multi-hop communication between vehicles to reach the AP has been proposed in [14], which however does not address the connectivity issues from the last hop and the AP. The role of parked cars has been recognized as very useful in maintaining the DTN fabric of vehicular networks [15], 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 vehicular networks. 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 the building 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 alleviate 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.

The final piece of the puzzle is changing the search target of a moving car from an available AP to a parked car instead that acts as a relay (see Figure 1). 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 vehicular networks, 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.

3. APPROACH AND UNIQUENESS

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 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. Although the application determines the specific style of communication, LZ hides the complexity of supporting the distinct communication styles from the application.

The key design feature of LZ is a clear distinction between the role of a moving vehicle and that of a parked one, both of which have *active* and *inactive* roles. A moving car is generally in the

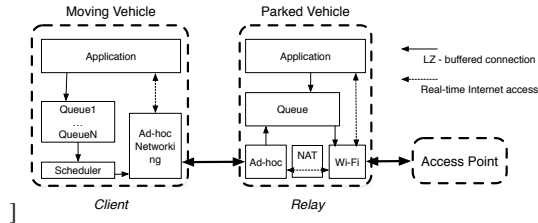


Figure 2: LoadingZones multi-layered architecture

inactive *idle* (*I*) state. It transitions to the active *client* (*C*) state when a connection to a parked vehicle is available and data can be transferred. Similarly, a parked car can be an active *relay* (*R*) or inactive and *off* (*O*) when the vehicle is parked in an area where no Wi-Fi connectivity or if there are any energy constraints. Energy consumption, which is a negligible component of the total car energy balance when the engine is on and the car moving, must be considered when the car is parked for long periods.

When a moving car has an *interactive* communication request, it enters the *client* state, and waits for a connection. As soon as a connection is established between a moving and a parked vehicle, 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 creates a new queue for each application with bulk data. Applications can then continuously enqueue their data and continue their normal operations. Now when a connection is established between a moving and a parked vehicle, the moving car enters the *client* state and 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.

We implemented a prototype of LZ using the Illinois Vehicular Project (IVP) [16], which is a vehicular network testbed based on customized hardware, built around a main board based on the AMD Geode processor. Our customized hardware reflects the capabilities of a vehicular communication system that is better than limited devices such as netbooks or mobile phones. With the space available, two radio interfaces, each with multiple antennas, is realistic for a vehicular device.

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. To verify the feasibility of LZ, we collected campus network RSSI beacons on campus streets by repeatedly scanning all frequencies while driving. In almost 60% of the streets we tested, the signal is higher than -90dBm. A signal as low as -90dBm is enough to establish a connection, although its quality would be poor.

The LZ software stack is based on the design principles presented in Figure 2. 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. The queue and message passing system is built to han-

dle both small messages (e.g. sensor readings) and large messages (e.g., images and videos). For fairness, an independent queue is created for 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. Upon the messages' arrival at the destination, the data is simply passed to the specified port. The amount of data that can be transferred during a contact is limited by the duration of the contact itself. Large messages are fragmented when they are enqueued, to enable delivery over multiple contacts. The message is reconstructed at the final recipient (the intended destination). The successful delivery of each fragment is guaranteed by TCP. Nonetheless, the probability that, after a successful delivery to the parked car, the fragments are not delivered to the AP is small, but not null (e.g., a parked car could move before it emptied its queue). LoadingZones does not attempt to reach 100% delivery guarantee for bulk data.

In our prototype, the test application sits on a smartphone connected to the device in the moving vehicle. The phone captures pictures and immediately enqueues them with no additional burden for the application. For our later analysis, the vehicle ID, along with the time and GPS location of both enqueue and dequeue, is prepended to each message at each queue. The entire journey of a message can be reconstructed, including where it was generated, when it passed between vehicles, when it was uploaded to the infrastructure, and the final AP's MAC address. A web front-end has been developed on the server to display all collected data and metadata on a map.

4. RESULTS AND CONTRIBUTIONS

Our evaluation of LZ investigates the benefit of the two-hop design of LZ compared to single-hop approaches. For measurements, we equipped a vehicle with the IVP hardware and instrumented the LZ code to measure the connection setup delay, the channel quality and the contact duration. 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 [10]. For this reason, we turned off rate adaptation and ran our experiments at fixed bit rates. We leave to future work the research of rate adaptation protocols specifically designed for vehicular scenarios.

Delay. Connection setup time is the composition of three operations: scanning, association and authentication, and IP setup. Scanning time depends on several factors, such as the channel switching period, the beacon interval, and the channel order. However, the majority of the setup time is spent to perform the last two operations. Timeouts during the scanning or association process have a significant impact on this time. Generally, timeouts happen when the signal quality is very low. Association and authentication time depends on the network configuration. Table 1 shows the results of experiments between the moving car and APs configured with different security policies. For these results, we did not consider the effect of timeouts.

The impact of security is substantial due to the additional overhead required to create the secure channel. Our experiments show that the use of WPA2 accounts for more than 33% of the total connection time. Even the use of optimizations such as those introduced in [1] would not help this overhead. In the IVP setup, the

	None	WEP	WPA	WPA2
t_a	1.97s	2.55s	3.41s	3.10s
σ	0.38s	3.02s	3.03s	0.74s

Table 1: Average association time (\bar{t}) and standard deviation (σ) using different security mechanisms.

use of WPA2-EAP on the campus network implies a remote handshake with the authentication server, which can be several hops from the AP. A comparison between traditional passphrase-based mechanisms and WPA2-EAP was not possible since we did not have access to the same APs as our infrastructure network does. Another possible source of delay, also ignored by Cartel, comes from open networks that may require extra actions such as accepting service terms on a web form.

DHCP servers can also be the source of significant delays for different reasons. First, in very complex structured networks, the DHCP server can be several hops away and then RTT dominates, especially in the case of timeouts. Our experiments show that the average time of a successful connection setup, including association, authentication and DHCP, is not significantly influenced by mobility. Regardless 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. We ran experiments at different locations and a successful connection was eventually established in 100% of tries. A moving vehicle, on the other hand, has a limited time to connect 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% for 30mph.

Quality. To measure the quality and contact 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, which also suffer from a non-optimal channel. In comparison, connections to a parked car, which in most cases is almost line-of-sight with the moving vehicle, mitigates the effect of velocity with a better channel. The average PER at different data rates for the same speed and placement does not vary significantly when the receiver is parked on the side of the street, while it increases to very high levels when the receiver is indoors.

Duration. Channel quality considered alone is not representa-

tive 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. Essentially, by using LZ, not only are contact durations longer, but the channel is also more predictable between a moving and a parked car. To understand the total amount of data that can ideally be transferred using 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 PER and longer contact durations. At higher data rates, where the channel quality also is affected by recipient location, the gap between the two approaches is even larger.

In our measurements, we do not account for connection setup time, which would take a substantial fraction of the overall contact time for AP-based solutions, even when all possible optimizations are applied. However, channel setup between a moving and a parked vehicle can be made more efficient. Authentication and IP configuration servers can sit on the car, and they are not required to be compliant to the networking standards. The design of secure authentication protocols is an orthogonal problem which must be solved, but it is out of the scope of this paper.

Energy. Compared to other forms of mobile networks, energy efficiency in vehicular networks is generally considered as a marginal problem, since the energy spent for communication is negligible compared to the total energy balance of a vehicle. However, LoadingZones establishes an active role for parked vehicles, which have a very different configuration. All energy available when the engine is off comes from a battery, which can have a large capacity, but is still a finite source of energy, which might not last for the whole duration of a vehicle’s stop. When LoadingZones can be active for only a fraction of the time, it is very important to understand when the energy is better spent, so that it provides the largest benefits to the whole system. To automate this decision process, we designed and implemented an energy governor for LoadingZones [17]. The main idea behind the decision process of our Governor is that a service is only as good as the benefits it provides to the system. For LoadingZones, we use traffic density to measure the benefits provided to the system, since when this metric is high, more vehicles can use LZ, while the energy cost to keep the service active is constant. Our governor implements the logic to automatically determine what traffic density should trigger the activation of LoadingZones.

To summarize, the benefits of LZ come for the complimentary characteristics of the moving car’s connectivity to the parked car with the parked car’s connectivity to indoor APs. Directly connecting to an indoor AP from a moving vehicle, as proposed in many existing solutions, is limited by a combination of three challenges: setup delay, channel quality and contact duration. In comparison, LZ’s two-hop solution, separates the effect of these challenges. LZ moves the low-quality link to the parked vehicle and the AP, on which setup time can be amortized over a longer contact lifetime, and creates a better quality link between the parked and moving vehicles, which also suffers less from setup time and short 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 these new networking paradigms, and reveals the need for future research on cooperation between vehicles to reduce the energy con-

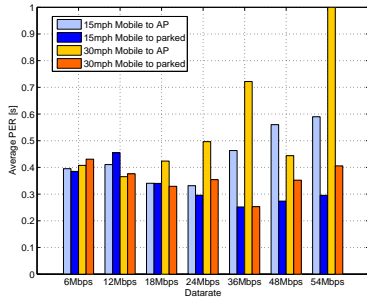


Figure 3: Average PER during a contact.

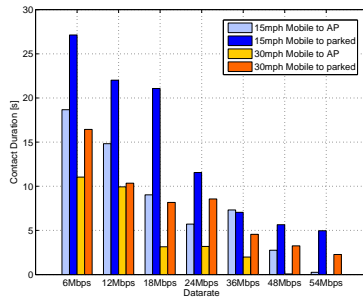


Figure 4: Link duration to an AP Vs a link between moving and parked vehicles.

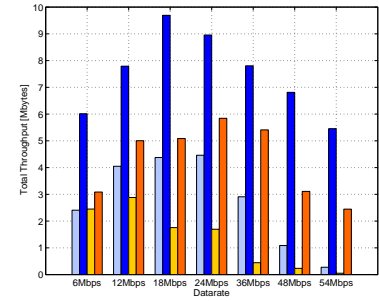


Figure 5: Total traffic delivered to an AP versus a parked vehicle per contact.

sumption of single vehicles without affecting the overall performance. We also intend to study efficient ways to discover and establish links between moving and parked vehicles to minimize the effect of connection setup delays. Finally, while deploying LoadingZones in a campus-wide scenario such as in IVP is simply a technical challenge, a more wide use of LoadingZones requires an appropriate incentive system for both the car owners and the network administrators that must share their AP signal.

5. 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] S. Yousefi, E. Altmaiv, and R. El-Azouzi, "Connectivity in vehicular ad hoc networks in presence of wireless mobile base-stations," *ITST*, 2007.
- [4] J. Eriksson, H. Balakrishnan, and S. Madden, "Cabernet: vehicular content delivery using wifi," 2008.
- [5] R. Frenkiel, B. Badrinath, J. Borres, and R. Yates, "The infostations challenge: balancing cost and ubiquity in delivering wireless data," *IEEE PerCom*, 2000.
- [6] M. H. Arbabi and M. Weigle, "Using vehicular networks to collect common traffic data," *ACM Mobicom*, 2009.
- [7] V. Kone, H. Zheng, A. Rowstron, and B. Zhao, "On infostation density of vehicular networks," in *WICON*, 2010.
- [8] K. Ibrahim and M. Weigle, "Optimizing cascade data aggregation for vanets," in *IEEE MASS*, 2008.
- [9] J. Ott, E. Hyytia, P. Lassila, T. Vaegs, and J. Kangasharju, "Floating content: Information sharing in urban areas," in *IEEE PerCom*, 2011.
- [10] Pralhad Deshpande and Samir Das, "BRAVE: Bit-Rate Adaptation in Vehicular Environments," in *Proc. of ACM VANET*, 2012.
- [11] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden, "A measurement study of vehicular internet access using in situ wi-fi networks," in *ACM Mobicom*, 2006.
- [12] Z. Zheng, Z. Lu, P. Sinha, and S. Kumar, "Maximizing the Contact Opportunity for Vehicular Internet Access," *IEEE INFOCOM*, 2010.
- [13] N. Frangiadakis, D. Kuklov, and N. Roussopoulos, "Pegasus: 802.11 connectivity at high speed," in *IEEE Globecom Workshops*, 2007.
- [14] J. Yoo, B. S. C. Choi, and M. Gerla, "An opportunistic relay protocol for vehicular road-side access with fading channels," in *The 18th IEEE ICNP*.
- [15] N. Liu, M. Liu, W. Lou, G. Chen, and J. Cao, "PVA in VANETs: Stopped cars are not silent," *IEEE INFOCOM*, 2011.
- [16] R. Crepaldi, R. Beavers, B. Ehrt, and R. Kravets, "Illinois vehicular project, live data sampling and opportunistic internet connectivity," *ACM MobiOpp*, 2012.
- [17] R. Crepaldi, R. Welsh, and R. Kravets, "Governing energy for parked cars," in *Proc. of WONS*, 2013.