

Policy-Based Network Management For Generalized Vehicle-To-Internet Connectivity

Joshua Hare, Lance Hartung, Suman Banerjee
Dept. of Computer Sciences, University of Wisconsin-Madison
Madison, WI, USA
{hare, hartung, suman}@cs.wisc.edu

ABSTRACT

Spectrum is the scarce and limiting resource for all wide-area wireless communication activities. As user loads on this limited resource continue to grow rapidly, it is necessary for “providers” to have effective tools, frameworks, and functionalities to manage allocation of this resource to users. In this paper, we introduce a policy framework, called *Virtuoso*, that manages heterogeneous spectrum resources for large collections of users aggregated over multiple operators. The goal of Virtuoso is to enable providers to flexibly manage their spectrum resources while ensuring that users have suitable access for their communication needs. Instead of describing the problem in the abstract form, we draw all requirements of this system from our multi-year effort of designing, implementing, deploying, and operating an Internet service for multiple public transit systems that operate in the state of Wisconsin, USA, and in some neighboring states of the upper Midwest. In Virtuoso we propose a simple syntax through which users and administrators can express their network control preferences. We believe that centralized control of network resources, as defined in Virtuoso, will simplify the management of connectivity for large mobile populations, of which vehicular fleets are just one example.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communications; C.2.3 [Network Operations]: Public Networks, Network Management

General Terms

Experimentation, Measurement, Performance

Keywords

Network Policies, Vehicular Connectivity, Wireless, Virtuoso

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.

CellNet'12, August 13, 2012, Helsinki, Finland.

Copyright 2012 ACM 978-1-4503-1475-6/12/08 ...\$15.00.

1. INTRODUCTION

Demand for wide-area wireless access services have continued to grow at a rapid pace. With the emergence of highly capable and intuitive mobile devices, the demand for this resource is expected to continue growth in a significant manner. A central challenge plaguing most providers is to identify the most efficient ways to distribute this spectrum resource amongst the multitude of users, with their diverse demands, loads, and expectations. As with any constrained resource, one efficient way of managing spectrum is to define usage policies that maximize a specific objective function (e.g., a utility function). Utility functions, often, are hard to define very precisely, as utility of any data flow is a complex function of the task, users involved, location, time, and many other parameters. Typically it is one or a few network administrators who have an informal notion of this utility function of different data flows and they enumerate various policies guided by this informal notion. The goal of this paper is to provide administrators and users of these wide-area wireless network providers with a suitable policy framework (which we call Virtuoso) that allows them to manage these wide-area wireless network policies in an effective and intuitive manner.

1.1 Network Model and Defining the Entities

In today’s model of cellular systems, users typically acquire a contract (maybe a long term one, or a pay-as-you-go contract) with one specific cellular operator. Infrequently, users change their operators for various reasons, e.g., price, performance, and flexibility. We envision a not-too-distant future where there will be a level of indirection between users and cellular operators, which we refer to as aggregators or providers. In our model, providers buy spectrum resources from the base spectrum operators, and provide connectivity services to users. Thus, providers are essentially resellers of the spectrum access service operating over the physical infrastructure of multiple cellular operators. Thus, in our model, cellular operators are entities that manage the physical infrastructure, e.g., base stations, fiber to the base stations etc., and provide a certain bit-pipe or spectrum-pipe abstraction to providers. Providers aggregate multiple such bit-pipes to provide an abstraction of a seamless city-wide, nation-wide, or a world-wide network.

While this precise model of wireless networks is not in place today, very close approximations already exist. The Mobile Virtual Network Operators (MVNOs) are essentially examples of providers in our model that re-sell spectrum resources of physical infrastructure owners, who are the opera-

Entity	Definition
Operator	Builds the physical infrastructure, e.g., base stations, fiber, and basic software and protocol services to allow client-devices to communicate
Provider	Aggregates “bit-pipes” by acquiring resources from multiple network operators (or other providers) and makes them available to users
User	The end-consumer that utilizes a client-device and has a connectivity contract with a provider

Table 1: Definition of the various entities in the context of Virtuoso.

tors. Currently, the only issue limiting the true aggregation of resources from multiple operators is that in many cases different operators often employ different radio front-ends, protocols, technologies, and access methods, and a user’s mobile device usually does not operate over multiple operator’s networks. Two emerging trends are taking some initial steps to eliminating this constraint. First, there is an emerging class of “multi-sim” phones, which can operate on two different networks. Second, is the concept of “soft-sim,” the ability of a phone to manage its SIM card parameters through software. As such approaches take hold in the marketplace, we envision that the model of network providers as aggregators across multiple operators will provide greater specialization of functions — some providing basic access mechanisms (operators), and others managing the end-user connectivity through multiple pathways (providers).

The definition of who a provider is, in this model, is fairly generic. An operator can serve as the provider (as happens frequently today). Alternatively, a provider may be more like a single MVNO, or it can aggregate resources from multiple other operators as well as other providers. In a certain way, a single large enterprise may be viewed as a provider for all its employees, where it buys data plans in bulk from different operators (say, 500 GB per month from each of AT&T, Sprint, T-Mobile, and Verizon) and makes the entire 2 TB of data available to all its employees as a giant “family-style” data plan. Assuming employee devices can be flexibly migrated from one operator network to another, the core challenge in this setup is for the enterprise administrators to efficiently manage the use of the aggregate data to support all the needs of its employees. We envision that a policy framework that incorporates a number of complementary functions — traffic filtering, shaping, partitioning individual flows across multiple paths (with or without redundancies), flow migration, etc. — all implemented in a centralized control framework that allows administrators to express time-varying policies would be necessary and useful to manage this complex problem.

While the general problem of any provider managing aggregated (spectrum, bandwidth, or data) resources is interesting in its own right, in this paper we focus on one specific instance of this problem in the domain of vehicular communication systems, that we have encountered through deployment and operationalization experience for more than two years now. More specifically, over this period we have implemented a vehicle-to-Internet connectivity solution, called WiRover [3] and deployed it for two bus companies: one a city transit system in Madison, WI, and the other operates a fleet of coach buses that travel various long distance

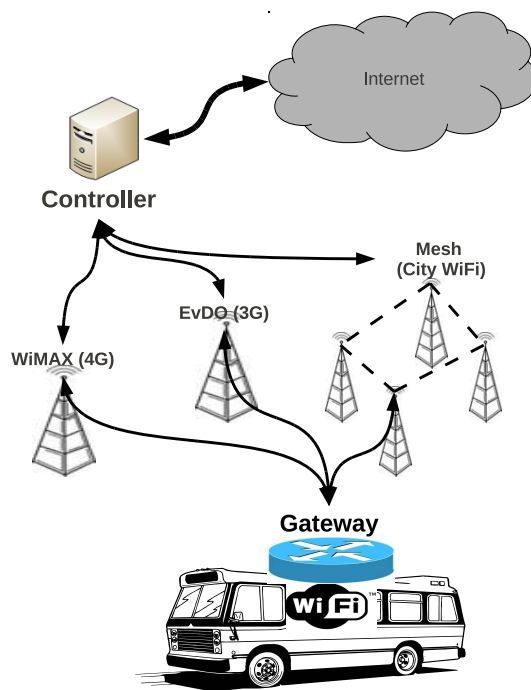


Figure 1: Overview of the WiRover system deployed on multiple bus networks that provides Internet connectivity to passengers through multiple cellular and WiFi networks simultaneously.

routes in the upper Midwest of the US. The WiRover system connects through multiple cellular networks simultaneously, currently among AT&T, Sprint, Verizon, a city-wide WiFi mesh network, and a custom WiMAX network (Figure 1).

To manage the spectrum, bandwidth, and data resources in this vehicular context, we describe a flexible policy framework, we call *Virtuoso*. This paper provides a detailed motivation, background, and design approach we have taken in Virtuoso, and provides specific details on various next steps in bandwidth management for large vehicular fleets.

2. REQUIREMENTS IN THE VEHICULAR CONTEXT

The automobile industry is moving towards greater connectivity in the vehicles. The ability to provide various Internet-based services, e.g., real-time vehicle telematics, traffic, road and weather updates, navigational information, and various other forms of infotainment, increase the efficiency and safety of vehicles. In many installations today, connectivity is tightly integrated with each individual application. For example, a navigational unit can come pre-configured with a cellular access service, while a built-in emergency calling and location service may be pre-configured with a different cellular service.

However, as the needs of connectivity pervade across different application silos, a single access gateway that serves as the connectivity hub for the vehicle becomes a natural choice. In such a setting, there are at least two obvious ways of offering Internet connectivity through these access gateways. We call them: (i) *Subscribe To Your Own Net-*




Figure 2: Multiple vehicles (black squares) on a road partitioned into three arbitrary segments. A provider can manage bandwidth resources across a fleet of vehicles customized to application bandwidth needs.

work: In this model, the access gateway in each vehicle is already equipped with a cellular radio modem, and the vehicle owner can contract with any one specific cellular operator for monthly service; and (ii) *Bring Your Own Network:* In this approach, the access gateway in the vehicle is configured to tether with the owner’s personal cellphone as the radio modem, and all Internet traffic piggybacks on the cellular access plan of that cellphone. However, there is an intriguing third option that we propose and explore in our framework that can put the vehicle manufacturer (or a designated third-party provider) in charge of managing cellular services to vehicles. We call it *Synthesize The Network*.

In this synthesis approach, we assume that the auto manufacturer (or a designated third-party serves as a cellular provider, by contracting bulk data plans with various cellular operators. When a vehicle owner wants to configure data services to the vehicle, it obtains a suitable contract from this provider. We assume that access gateway can operate through the diverse cellular networks, *simultaneously*, if so desired. This provides the provider with a great flexibility in optimizing connectivity performance to different vehicles. When a vehicle is passing through different road segments, it might be efficient to connect to different networks based on historical information of performance of these networks. As Figure 2 illustrates, the middle road segment may be a poor connectivity zone for Net A, and hence when a vehicle transits that region the vehicle’s accesses the Internet through Net B only. Occurrences of such performance diversity has been demonstrated in prior measurement studies [6]. At other times, the provider may prefer to use Net A for vehicles (maybe because of lower data rates of Net A in general). Finally, the provider may also increase the throughputs for a certain vehicle which is imposing a higher traffic demand (e.g., kids watching a streaming movie in the back seat). Network diversity and aggregation can, thus, be used in many interesting ways by the provider to meet various performance needs.

Our WiRover system and requirements for the flexible policy framework: The WiRover system consists of an access gateway mounted in the vehicle that is equipped with multiple wireless radios to connect to different cellular and WiFi networks in range, simultaneously, when possible (Figure 1). Passengers can access Internet services by connecting to this gateway over WiFi. In addition, various vehicular services (e.g., telematics, electronic displays, etc.)

can also connect to the gateway over Bluetooth, WiFi, or Ethernet.

The deployment experience brought various bandwidth and data management challenges and requirements to the fore. The following are some examples.

- *Keeping to data limits:* Our WiFi service to bus passengers is freely available to all. Under such scenarios, many passengers may impose a high traffic demand. If this spirals out of control, then the monthly data limits would be easily breached leading to high, and variable, costs for operation of these systems [4]. So a key requirement for our system is to ensure that the Internet traffic from bus passengers do not lead to breach in data limits. Of course, one simple way to implement this is to impose hard filtering rules when certain thresholds are crossed. However, a more benign and continuous traffic shaping and management approach is required for a more graceful experience. For instance, we would like to identify heavy flows and temporarily shape or limit them. The aggressiveness in such shaping decisions could depend on the likelihood of breach in the data limits. Given that each gateway is equipped to connect to multiple networks simultaneously, this process could easily be coupled with decisions of migrating flows between networks when needed. This requirement implies that all such traffic shaping and filtering decisions would need to be adapted over time.

- *Shared data plans:* A more interesting version of this problem arises when a number of vehicles can share a common data plan. For example, if a fleet of 100 buses is provided a monthly limit of 500 GB in aggregate for each network in use, we need to design efficient ways to manage the data aggregate across all vehicles. This brings an entire new dimension to data allocation and management.

- *Value of data and traffic priorities:* Clearly not all flows and their data are equal. Flows carrying vehicle diagnostic and telematics data to a central monitoring application for the entire fleet are more important than the Facebook updates being posted by passengers. In certain cases, real-time data from security cameras mounted on the buses might be even more important. Therefore, when the system is approaching data limits, certain flows may need to be shaped or filtered aggressively while others should not be. Further, to do this efficiently, the system would benefit from the ability to predict the traffic demands of various flows. Overall, administrators need to be able to flexibly design and update these policies that determine how different flows are handled.

- *Content Filtering:* A basic request from our bus operators was to ensure that passengers could not download offensive content. Since the bus is a public space it would be possible for one passenger to subject other unwilling passengers to content that would be deemed offensive. Bus operators were also not interested in supporting the unauthorized download of copyrighted material. Further, bus operators did not feel that it was necessary to support high bandwidth video streams, if it caused poor performance to other users. So under periods of overload, filtering streaming video applications was also considered a desirable feature by our operators.

- *Network prioritization and camera video offload:* Fleet operators desired that they be able to specify priorities across different networks available for specific types of traffic. For example, there are multiple security cameras mounted on buses that record and store events both inside and outside

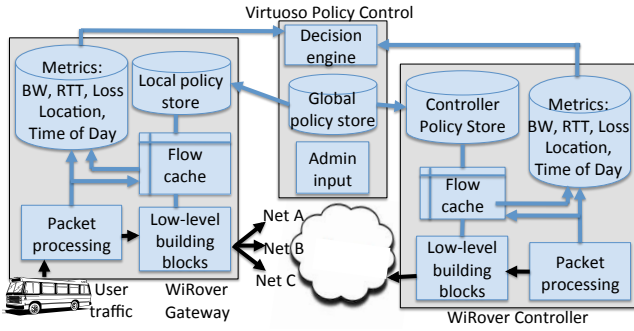


Figure 3: Overview of the Virtuoso policy framework and how it interacts with the WiRover gateway and controller subsystems.

the bus. In normal circumstances, no real-time access to these camera feeds were necessary and they could be offloaded from the buses once a day. However, shipping this data over cellular networks would be expensive. Therefore, it is best to ensure such offload happened over WiFi networks (available in the bus depot) [2]. However, if some emergency situation occurs, then public safety agencies may want immediate access to these feeds requiring an immediate change in this policy. Our policy framework needs to suitably support such flow requirements.

To address all of these requirements in a flexible manner and to empower the fleet administrators to manage aggregated bandwidth and data resources in an efficient manner, we propose the Virtuoso policy framework. This is a centralized management structure through which the fleet administrators can input all policies at multiple levels — gateways (in vehicles), networks, applications, and the various flows to and from each gateway. Through this framework, global policies are specified by the administrator which translates to various local policies and are downloaded, implemented, and enforced at distributed points in a dynamic manner. In certain cases, a global policy requires coordination across multiple points, e.g., both at one or more WiRover gateways and their controllers.

Some example scenarios: Different applications and scenarios require different ways to provide the best performance end-to-end. Our proposed policy framework handles a large range of application requirements in its design. We present a few simple examples:

- *NAT vs encapsulation and network failovers:* WiRover’s design offers a simple but useful design choice for each flow — it can be directly NAT’ed at the gateway and sent into the Internet, or it can be encapsulated through the WiRover tunnel and sent to the controller (which acts as a proxy) for further forwarding. While NAT’ing is more efficient in general, encapsulation offers the opportunity to protect long-lived flows better by transparently migrating them from one network path to another, when there are failures in its original network path. In fact, the tunnel structure allows stripping of packets from a single flow across multiple paths, if so desired. This can achieve higher aggregated throughputs for flows and new robustness strategies. However, the cost of encapsulation is some centralization in the data plane at the controller leading to potential bottlenecks. Hence, admin-

istrators are empowered to identify individual flows which should be encapsulated.

In our example, we prefer that all user (web) traffic be directly NAT’ed, while some of the critical and long-lived flows carrying important telematics and navigation data be encapsulated. Exceptions can be made on-demand for users or applications that require greater robustness.

- *Shaping a flow, a user, or a network:* Administrators have the ability to perform traffic shaping at various granularities — an individual high-bandwidth flow, an individual high-bandwidth user, or across all traffic transiting through a specific network. Different circumstances may lead to each of these scenarios, e.g., if a single video-streaming flow is causing congestion for all other flows in the same vehicle, it may be worthwhile to just shape that single flow. Shaping different users may be relevant to provide certain classes of paid services in the vehicle, while shaping all traffic for a single network can be useful if certain data limits for that network are approaching [5].

- *Location-specific flow duplication:* If the vehicle is passing through a region of poor performance of the network paths, and if capacity permits, it might be useful to duplicate the traffic of certain demanding flows across multiple network paths. We observed this for certain latency-critical applications, e.g., audio- or video-conferencing, when the performance of any individual network is not adequate to maintain a suitable quality of service.

3. IMPLEMENTATION AND INTERFACES

The Virtuoso policy framework empowers the provider (in this case, the fleet administrator) to centrally manage data and bandwidth resources across all its users (vehicles). It is implemented as a cloud service. All policies specified by the provider are entered through a web interface. A decision engine translates them to specific gateway and controller level policies. For instance, the administrator may want to implement a policy as follows: “If the aggregate data usage on Net A across all vehicles exceeds 200 GB (when the limit is, say, 250 GB), then limit all users to a maximum of 50 Kbps, while imposing no limits on telematics and navigation related traffic.” This policy will translate to various local policies (at the gateways) that trigger at different time instants. At the beginning of the month, the gateway imposes no limits on flows from the users. However, once the 200 GB limit is reached, each gateway is instructed to shape traffic aggregated per user to the 50 Kbps limit. This policy gets removed at the beginning of the following month. The overall logical structure of Virtuoso is shown in Figure 3, and Table 2 lists a number of example policies relevant to our system. The syntax of the policies are available as a command line interface.

To implement these policies, we have created a set of low-level building blocks — (i) token bucket based rate limiter for individual flows or users, (ii) rate limiters for a specific network interface (aggregated for all traffic flowing through it), (iii) filtering capability, (iv) ability to prefer one or more network paths over others, (v) ability to identify different types of applications, and (vi) a decision component that can choose to directly NAT traffic to the Internet or encapsulate traffic via the proxy (controller). For encapsulated traffic, we also allow flow migrations from one network path to another, traffic duplication across different paths, additional robustness through multi-path network coding,

Type	Command	Notes
Limit user	policy -src 10.0.0.1 -limit 200K	Limit user at 10.0.0.1 (for a given gateway) to a 200Kbps limit
Limit dest.	policy -dst youtube.com -limit 200K	Limit all YouTube traffic from this gateway to 200 Kbps
Limit network	policy -net Verizon -limit 200K	Limit all traffic using the Verizon network to 200 Kbps
Block an app	policy -drop -app torrent	Block all BitTorrent traffic on this gateway
Link Selection	policy -link WiFi -app video	All video traffic should only be sent over WiFi links
Set priority	policy -priority high -app music	All flows carrying music content has the highest scheduling priority
Packet striping	policy -bwagg Verizon,Sprint -proto TCP -dport 20	All TCP flows using port 20 should be striped across the two specific cellular networks

Table 2: Example policies relevant to the WiRover deployments.

striping an individual flow across such multiple paths, traffic compression and de-duplication, and encryption and decryption across the tunnel. Based on the specified policies, different building blocks are invoked on the traffic on the end-to-end data path for each flow.

For deciding when and where to instantiate various policies, the gateways and the controllers gather various traffic metrics which are communicated in aggregate to the decision engine. (Some of these metrics are also communicated to the local policy stores.) The metrics of direct interest include various estimates of available bandwidth, link latency, and loss rates. In addition, it is also useful to convey the location and the time of these measurements to the decision engine. This is necessary because historical performance could be used to predict network performance based on a vehicle’s current location and the time of the day.

Most of the policies in local stores (in the gateways and controllers) may not be active at any given instant. Hence, we use a simply flow cache that maps an abstract policy to a usual 5-tuple flow and various actions on packets of that flow. A single policy may lead to multiple entries in this flow cache and is used in our system for efficient packet processing on the data path.

4. PRELIMINARY ANALYSIS

To demonstrate the need for particular policies we analyzed some of the usage data we collected during our tests of a vehicular Internet system over the past 20 months. Our system has primarily been tested on passenger buses which provide Internet connectivity to bus passengers. The goal of such a system is to ensure all passengers have an acceptable experience while ensuring that the system remains within operational budgets. Currently, our system uses multiple USB cellular modems over which the passengers traffic is forwarded to the Internet.

4.1 Avoiding Cellular Coverage Fees

To remain within operational budgets it is important to ensure the network usage does not exceed the transfer limits of any one cellular network. There are multiple methods to reduce data consumption such as caching, compression, and traffic shaping. Traffic shaping can target large network flows, heavy users, or all traffic over a particular network.

The difficulty with avoiding cellular overages is that actual passenger usage is unpredictable. As can be seen from Figure 4, the actual cellular usage without any usage restrictions on passengers’ traffic frequently exceed the limits imposed by the cellular providers. To solve this issue we suggest using a combination of traffic control mechanisms.

At a basic level caching techniques can be used to re-

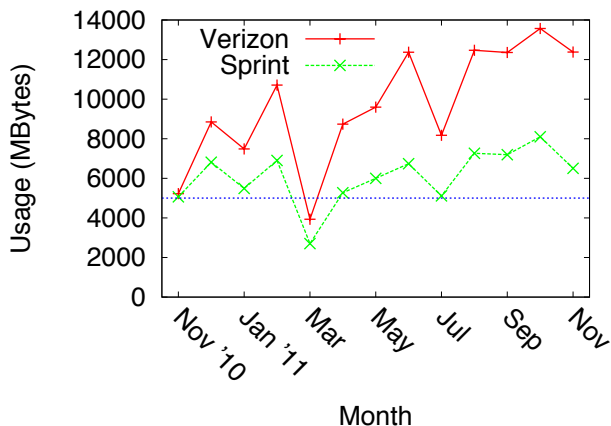


Figure 4: Data usage by network as observed for a single coach bus. The horizontal line represents the 5 GByte cellular transfer limit.

duce the number of bytes sent and received. Further, since most observed traffic is HTTP further optimizations can be employed such as HTML/CSS minification, image optimizations, and file compression. Such features can easily be realized through open source proxies like Ziproxy [1]. To further reduce the data consumption various methods of filtering could be implemented to remove unwanted traffic. For example, offensive content could be filtered out as well as unwanted protocols or applications like BitTorrent.

As passengers consume more and more data, caching and compression can only reduce the amount of bytes transferred over each network to a certain point. In such a situation it might still be necessary to implement traffic shaping to encourage users to consume less data.

4.2 Limiting Heavy Usage

Traffic shaping is sometimes viewed as a poor solution to managing network resources. However, we have found situations in which traffic shaping would provide many benefits. For example, we have found that a minority of users consume the majority of available bandwidth and also consume a large portion of bytes. These users hurt the performance that other concurrent passengers receive but also transfer a large amount of data that can cause a particular network to exceed the imposed transfer limits.

As can be seen in Figure 5, passengers are connected to the Internet for varying amounts of time. Despite the pas-




Figure 5: User connection durations during a representative coach bus route. Users are sorted by decreasing connection time.




Figure 6: The number of bytes received by each user during a representative coach bus route. User indexes are correlated to Figure 5

sengers' connection durations some users exhibit behaviors of heavy usage, Figure 6.

To improve fairness between concurrent users the system could implement various fair queuing methods. However, to solve the issue of overall data consumption it may be beneficial to further shape the traffic from heavy users. Traffic shaping is further motivated by the views of the bus operators in which they desire to provide basic Internet connectivity but do not necessarily want to support video stream or torrent downloads.

To provide a basic level of network performance bus operators must balance the amount of bandwidth available to each system with the cost of providing that bandwidth. As usage grows it may be necessary to increase the system's bandwidth, possibly by adding another cellular modem or upgrading to faster wireless technologies.

5. THE "ROAD" AHEAD

Generally, there are many critical applications for bus operators which require a certain amount of bandwidth. With standard cellular data plans a customer often is allocated a certain amount of data for each month which must be greater than or equal to the amount of data required by their critical applications. The remaining resources should then be used by non-critical applications such as passengers' Internet connectivity. To accomplish such a task the policy framework can predict the amount of resources required by critical applications to ensure these resources are available when they are needed. The policy framework could then evenly distribute the remaining resources over time to the non-critical applications.

Traffic shaping can further be beneficial to the operation of some vehicle-to-Internet solutions; however, more work is needed to better understand how and when traffic shaping is acceptable. Using traffic shaping to avoid overage fees when a global pool of resources is shared across a fleet of vehicles is a difficult problem as a certain level of resources are needed to provide applications with a basic level of performance. With pooled resources it may be necessary

to distribute resources such that application performance is relatively consistent across all vehicles of the fleet.

In some environments traffic shaping may require the support of network proxies. For example, providing Internet connectivity to bus passengers can generate a large amount of network traffic in the downlink direction. Shaping or policing this downlink traffic from the vehicle may be an extremely difficult method of avoiding cellular transfer limits.

6. ACKNOWLEDGMENTS

We thank the Madison Metro Transit and the Van Galder Bus Company for partnering with our research efforts and for their feedback and suggestions.

7. REFERENCES

- [1] Ziproxy. <http://ziproxy.sourceforge.net/>.
- [2] Aruna Balasubramanian, Ratul Mahajan, and Arun Venkataramani. Augmenting mobile 3G using WiFi. In *MobiSys '10: Proceedings of the 8th international conference on Mobile systems, applications, and services*, New York, NY, USA, 2010. ACM.
- [3] J. Hare, L. Hartung, and S. Banerjee. Beyond deployments and testbeds: Experiences with public usage on vehicular WiFi hotspots. In *ACM MobiSys*, 2012.
- [4] Andrew Odlyzko. Internet pricing and the history of communications. *Computer Networks*, 36:493 – 517, 2001.
- [5] Barath Raghavan, Kashi Vishwanath, Sriram Ramabhadran, Kenneth Yocum, and Alex C. Snoeren. Cloud control with distributed rate limiting. In *ACM Sigcomm*, 2007.
- [6] S. Sen, J. Yoon, J. Hare, J. Ormont, and S. Banerjee. Can they hear me now?: A case for a client-assisted approach to monitoring wide-area wireless networks. In *Internet Measurement Conference*, 2011.