

A Retrospective Look at the UMass DOME Mobile Testbed

Hamed Soroush^a
hamed@cs.umass.edu

Nilanjan Banerjee^b
nilanb@uark.edu

Mark Corner^b
mcorner@cs.umass.edu

Brian Levine^a
brian@cs.umass.edu

Brian Lynn^a
blynn@cs.umass.edu

^aComputer Science Department, University of Massachusetts, Amherst, MA, U.S.A

^bComputer Science and Computer Engineering Department, University of Arkansas, Fayetteville, U.S.A

In this paper we describe the evolution of DOME, a diverse outdoor testbed for mobile experimentation. In addition, while highlighting the challenges faced in construction of DOME, we describe a concrete set of scientific results derived from this experience in a retrospective study. First, we argue that a broad range of mobility experiments could be performed in a testbed which provides the properties of temporal, technological, and spatial diversity. We demonstrate these properties in our testbed through analysis of data collected from DOME over a period of four years. Second, we crystallize a set of design principles that others should use when constructing testbeds of their own, including those related to deploying and managing a diverse testbed, distributing experiments remotely, and fostering collaborations among testbed stakeholders. Finally, using traces collected by DOME, we provide insights into several important problems in mobile systems research.

I. Introduction

Mobile systems and networking researchers confront a myriad of challenges, including power consumption, channel and radio characteristics, mobility, and node density. Each of these factors alone is complex, and their combination can be difficult to model. Moreover, mobile systems span a wide spectrum of rapidly evolving radio technologies (WiFi, Bluetooth, UWB, 3G, GPRS, and 900MHz radios), mobile devices (laptops, tablets, smartphones and music players), and networking paradigms (mobile ad hoc, disruption tolerant, and infrastructure-based networks). Models that consider such complex interactions still may not cover indirect factors such as social trends and real-world distribution of resources. To account for these difficulties, most mobile systems researchers advocate driving evaluation with testbeds. To support the comparison of a wide array of systems, testbeds must provide spatial, technological, and temporal diversity, realistic mobility patterns, power consumption, latency, throughput, programmability, and end-user participation. For instance, measuring how the performance of cellular and organic WiFi have changed over time across urban and rural areas requires a testbed with a broad range of capabilities.

Unfortunately, building a sufficiently general testbed is typically infeasible due to the time and expense

required—the majority of existing testbeds are tuned to a particular area of research. For example, CarTel [18] and VanLan [24] study dense vehicle-to-access point (AP) communication over WiFi links. On the other hand, deployments like UCSB MeshNet [23], CitySense [26], GoogleWiFi [1], RoofNet [19], TFA-Mesh [13] study performance issues in WiFi based mesh networks. On the other hand, a primary goal of mobile computing research is the ability for systems to transparently move among any of these scenarios. Hence, testbeds for mobile computing research ideally possess both technological and spatial diversity, enabling the evaluation of different radio technologies and network architectures in varied densities.

Another major shortcoming of many existing testbeds is the relatively small time scale of data collection. Many trends in mobile computing take place over longitudinal time scales. For example, a great many projects rely on opportunistic connections to open WiFi APs, yet trends in open AP availability have not been measured. Similarly, the populations and geographic areas of most testbeds are small. Results based on only a few mobile or stationary nodes covering a relatively small area cannot, in general, be extrapolated to more extensive scenarios.

To address these shortcomings, we have deployed and evolved an architecture and implementation of a testbed, the Diverse Outdoor Mobile Environment

(DOME). To our knowledge, DOME is the longest-running large-scale, highly diverse mobile systems testbed. The testbed has been operational since 2004 and provides infrastructure for a wide range of mobile computing research. It includes 40 transit buses equipped with computers and a variety of wireless radios, 26 stationary WiFi mesh APs, thousands of organic APs, and half a dozen nomadic relay nodes. It provides support for diverse radio technologies, including WiFi, 900MHz, 3G, and GPRS. It covers an area of 150 square miles and provides spatial diversity; parts of the network form a sparse, disruption-tolerant network while others are more dense. With proper isolation the testbed can be used for research ranging from infrastructure-based networking to sparse and dense ad hoc networks.

While the construction of the testbed represents a significant engineering challenge, this paper describes a concrete set of scientific results derived from this experience in a retrospective study extending our previous work [36]. Based on the data collected from May 2004 to October 2008¹, we show that the properties of temporal, technological, and spatial diversity are essential in producing a generally useful mobile testbed. Second, we offer a set of design principles unique to mobile testbeds as compared to wired testbeds, such as PlanetLab [31] and EmuLab [38], or stationary wireless testbeds, such as Orbit [33]. These principles include notions of flexibility, consistency, experiment pre-staging, resource reservation, failsafe operation, and third-party collaboration. We demonstrate the ability of the testbed to measure temporal, technological, and spatial diversity by providing insight into several open questions in mobile systems research. The most significant of these results show: WiFi coverage became sufficient for some low-bandwidth vehicular applications, as many regions offer small amounts of throughput; users continued to operate open WiFi APs with no sign of abating; cellular coverage (such as 3G) provided greater coverage than WiFi, and far superior throughput; and interference did not play a significant role in degrading throughput for vehicular WiFi users.

II. The Case for DOME

After examining the motivation for building a great number of testbeds, both wireless and wired [1, 13, 17–19, 23, 24, 26, 31, 33, 38], several common themes emerge.

¹While the testbed has been operational since 2008, we stopped constant collection of the presented data in support of conducting other experiments.

- They are short lived or are no longer operating.
- They lack sufficient hardware diversity.
- They lack sufficient geographic diversity.
- They cannot compare systems on equal footing.
- They cannot be programmed remotely or lack sufficient control over certain hardware or software.
- They lack appropriate or realistic mobility.
- They lack real users and traffic.
- They have insufficient scale.
- They cannot be combined with other testbeds.

Most common are the first three limitations: mobile testbeds typically lack the temporal, technological, and spatial diversity to answer a number of questions about mobile systems and technology, so new testbeds are continuously created. However, as mobile systems mature—much as wired networks matured—centralized, remotely programmable testbeds must take over from short-lived or narrowly focused deployments. This evolution is needed to address a number of open and ongoing questions in mobile research:

Challenges related to temporal trends: A large body of recent work is based on the availability of open WiFi APs that provide free, ubiquitous connectivity to mobile users. At the same time, off-the-shelf APs increasingly help less tech-savvy users to restrict access. *Are research systems that use open APs for ubiquitous connectivity viable in the long term [1, 15, 37]?* The answer requires a study of the longitudinal trends in open versus encrypted APs. Similar questions can be asked of research relying on the popularity of peer-to-peer networking connections [17] or ubiquitous cellular deployment. Also, *how do AP selection algorithms [27] or other systems affected by changes in wireless environments perform over the long-term?*

Challenges related to technological diversity: The use of different radio technologies, such as WiFi, 3G, and proprietary 900 MHz radios among others, presents a fundamental cost-benefit trade-off. Opportunistic connections to open AP WiFi is free but can suffer from disrupted coverage or poor quality. On the other hand, cellular technology like 3G has better coverage but comes at a higher monetary cost. Several fundamental questions are relevant: *What are the performance characteristics of each type of network? Which applications can be supported by only free AP access? Can multiple radios support and complement one another [34]? How does the performance of open*

free WiFi infrastructure [18] compare to self-deployed mesh nodes [8, 13, 19, 24]?

Challenges related to spatial diversity: The performance of many network scenarios is dependent on spatial density of infrastructure or peers, spanning issues of coverage, mobility, and interference. While spatial density is easily parametrized, values observed in the field are due to a complex set of user needs and demographics and off-the-shelf availability. Observations about spatial diversity in the field can help address questions such as, *For what densities are MANETs or DTNs practical [9]? Similarly, at what density are infrastructure networks sufficient to support delay-intolerant applications [4]? And are organic WiFi deployments sufficiently ubiquitous to support mobile computing [18]?*

Some of these questions have been answered in isolation, on a small scale, or for short periods of time; however, there is a need to answer these questions on a continuous basis, confirming trends or discovering new ones, and evaluating systems over longer time scales, wider geographic regions, and through heterogeneous hardware living under a common testbed. We have constructed DOME to help address these challenges. However, it is important to note that DOME does not, and cannot answer, all of these questions at present. Rather the intention of DOME is to answer a large number of questions, evolve the testbed to answer more questions, and to create a set of principles that guide the development of future DOME-like testbeds. One sign that this maturation of testbeds is occurring is the Global Environment for Network Innovations (GENI) testbed. The goal of GENI has been to combine the largest, and longest running, wired and wireless testbeds, including DOME, PlanetLab [31], Emulab [38] and Orbit [33] under a common framework. DOME is now integrated with GENI.

III. Design Challenges and Principles

From our initial construction and subsequent expansion of DOME, a number of design principles have emerged that we enumerate here. Our focus in this section is largely on challenges that are specific to, or exacerbated by, the case of a wireless, mobile testbed; this focus is in contrast to wired testbeds, such as PlanetLab [31] and Emulab [38], and wireless stationary testbeds, such as Orbit [33] and CASA [32]. Like these previous successes, we seek to support experimenters through generality and programmability. To those

ends, we have followed several important tenets involving flexibility, consistency, experiment pre-staging, resource reservation, failsafe operation, and third-party collaboration.

III.A. Flexibility and Consistency

At the highest level, our testbed architecture is shaped by two competing principles. First, **Tenet 1: Testbeds must be flexible and evolve with changing standards, opportunities, and technology.** Otherwise, a great deal of effort can be put into a system that is available to run only a limited set of experiments, and then is quickly forgotten. Our own enhancements have included the addition of 900 MHz, GPRS, and 3G radios, the upgrade of core hardware to support virtual machines, and upgraded WiFi equipment to allow virtualization. Our most significant change has been modifying the software infrastructure to link into the GENI confederation of testbeds. GENI similarly embraces a core principal of *spiral development*, which seeks to leverage any newly available opportunity developed by the research community while not slowing progress.

Conversely, **Tenet 2: Testbeds benefit from stability that enables consistent measurements and observations of long-term trends.** Any change in node mobility, hardware, and power regime can change fundamentals of network connectivity and performance and subsequent experiments. For example, in Section V.B, we report on the percentage and number of APs that are open for third-party connections to the Internet. These trends are key to a great deal of research that leverages opportunistic WiFi connections [4, 6, 18, 37]. Updates to our WiFi hardware and mesh network have greatly increased the testbed connectivity. The increased connectivity had the unfortunate effect of changing our sample population, resulting in two phases of the same experiment that cannot be compared directly.

Balancing these two issues of flexibility and consistency is perhaps simpler for wired testbeds, such as PlanetLab, because the hardware is of secondary concern to the testbed and typically does not constitute the primary bottleneck. Further, the topology of a mobile network is heavily determined by how the testbed is built, whereas in PlanetLab, the testbed nodes form a small percentage of the overall Internet. In a mobile setting, evolving the testbed while still measuring long-term topological-based trends is thus more challenging.

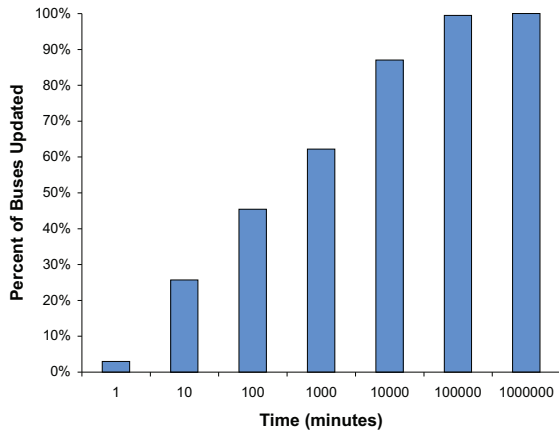


Figure 1: The empirical cumulative distribution of update propagation to mobile buses. The results are averaged over 11 updates.

III.B. Remote Programmability

Systems testbeds are subject to a variant of Metcalfe’s law. **Tenet 3: The value of a testbed is proportional to the number of experimenters that have access to the system.** Initially, DOME had been open to researchers at UMass and partners that visited for an extended time. More recently, DOME has become a Global Environment for Network Innovations (GENI) site, with the goal of following a community standard for cross-linking resources with other testbeds in the program. As part of that program, we have enabled more automatic programming of DOME, including virtualizing many of the resources.

However, the unique, primary challenge we have confronted in providing PlanetLab- or Emulab-like functionality to DOME is node mobility. Mobile nodes that are unreachable through the (wireless) network may still be functional and interacting with other nodes. In contrast, PlanetLab nodes that can’t be reached are assumed to be down since the network that connects nodes is the stable Internet. Mechanisms for staging experiments on PlanetLab, including PLUSH [2] assume that unreachable nodes cannot be part of an experiment; accordingly, PLUSH will not stage software on unreachable nodes. In contrast, mobile testbeds like DOME require the converse approach. **Tenet 4: Disconnected operation of mobile nodes requires that the administrative substrate support pre-staging of software needed for experiments, before reservations of resources can occur.** Once experiments are pre-staged, then the experiment can be launched, and then we take a count of the number of nodes used in the experiment.

III.C. Robust Operation

There are several practical difficulties involved in supporting experimentation on mobile testbeds. These can be generalized as **Tenet 5: Mobile testbeds must have an automated update mechanism, sufficient administrative bandwidth for updates, and a fail-safe mode to recover from corruption.**

Manually reprogramming all nodes in a mobile, disconnected testbed is impractical because it is too time consuming to track down all components. DOME uses a custom-built push-based mechanism to propagate updates. Originally, DOME used the opportunistic connections of buses to open WiFi APs to distribute updates. This worked sufficiently, but was slow and is impractical for supporting rapid reprogramming. Currently, each bus has a 3G data connection, which is highly available and has high throughput, but we have found that distributing updates to the full testbed takes an inordinate amount of time. Figure 1 is a CDF of the amount of time taken for an update to propagate through all the nodes in the mobile network. The results are based on 10 actual updates with an average size of 4.5KB (and a max of 18KB) all issued between June 11 and Nov 25, 2008. Note that these updates were sent while the buses had 3G access, and the update sizes are very small. Several factors affect the update times. For example, an update on a Friday or weekend should, on average, take longer because of reduced weekend schedules (where only about 5 buses are in use). Also, some buses can be in maintenance for extended periods of time.

High priority should be given to detecting and recovering from hardware and software failures. This is especially important in a mobile testbed where hardware failures are more prevalent and software failures take much longer to diagnose and fix. Nodes deployed in harsh environments, such as vibrating vehicles, outdoor building-tops, or attached to bicycles, suffer from an accelerated failure rate. Furthermore, nodes often consist of experimental and commodity hardware, none of which is necessarily hardened to the environment. For example, in our testbed, we suffer from frequent disk-related problems, including shock and vibration induced hardware failures, as well as corruption of the file system due to random power losses.

As our system has evolved, we have found the need to track failures at a fine granularity and filter our traces accordingly after the fact. In experimentation with software components, such as routing protocols, it is crucial to know how many nodes in the system are operating correctly. Since physical inspection of the mobile nodes is often not an option, the nodes must

be able to detect and isolate failures. Using diagnostic tests (see Section IV), checking peripherals and software components is extremely useful to track the quality of the data.

III.D. Resource Reservation

DOME was built with the intent that it be able to host a variety of experiments and collect measured data over a long period of time. It is designed to allow a varied complement of technologies in numerous, configurable topologies.

The PlanetLab testbed addresses this through *slivers* and *slices*, which are small segments of resources, such as CPU and network) on distributed nodes [31]. However, the crucial issues are that many events are ephemeral in nature, switching times can be long, and nodes must be coordinated beyond the ability of the testbed. **Tenet 6: Mobile testbeds are not easily amenable to fine-grained resource sharing.** For instance, two vehicular mobile nodes may pass one another at high speed, but both nodes must be running the same experiment at the same time, and must not switch to running another experiment in mid-stream. While some relatively fine-grained switching is available for some WiFi devices [21], such a system would be extremely difficult to synchronize between two mobile nodes at a sufficient granularity. Further, the routing in a disconnected network depends on these rare events for its performance and if multiple experiments are competing for each opportunity, it is difficult to scale the results accurately. Thus far, to support multiple experiments accurately, we have instead opted for more coarse-grained testbed-sharing, often with slices of a week or a month. Though we support finer granularity, we believe that the appropriate granularity of DOME is on the order of a day, and we have developed a straightforward mechanism to run different experiments on separate days.

III.E. Partnering with Third Parties

There are unavoidable political aspects to running a testbed that relies on the cooperation of third parties. In the case of DOME, we benefit greatly from the participation of the public transit authority (Pioneer Valley Transit Authority), the Town of Amherst, and IT staff of UMass Amherst. We have found these relationships to be a benefit for the project, but we know of other projects where analogous entities prevented a testbed from being deployed.

We believe our success benefited from the open minds of the people in our community, but more impor-

tantly, we have pursued many goals that were mutually beneficial: **Tenet 7: The deployment of an outdoor testbed must provide a reciprocal benefit to relevant third parties.** For example, each of the UMass buses offers Internet access through the 3G modem to riders of the bus. The benefit to the Town has been even greater, giving free Internet access to thousands of people in the downtown area. However, the Town has greater goals in mind, including moving municipal sensors, such as sewage meters and parking stations, from expensive leased lines to WiFi. Wired testbeds such as PlanetLab have followed a different model, where the researchers using the testbed have donated the resources, thus providing reciprocal benefit only to the research community. This often doesn't translate into a mobile testbed—covering large geographic regions requires more than researcher participation.

IV. Implementation of DOME

To give other testbed designers a starting point, and to place our traces and evaluation in context, here we provide an overview of the hardware and software that comprises DOME.

IV.A. Hardware Components

The DOME testbed consists of three major hardware components: the DieselNet vehicular network, a set of nomadic throwboxes, and an outdoor mesh network. At various times since DOME's inception in 2004, we have upgraded or improved virtually every hardware and software component. This has created unique challenges for extracting longitudinal data, as we discuss in Section V.

DieselNet Vehicular Nodes: Mobility in DOME is provided by a vehicular network called DieselNet [11]. It provides nodes that operate year-round, across a micro-urban and rural environment. DieselNet is comprised of 40 transit buses, each equipped with Hacom OpenBrick 1GHz Intel Celeron M systems with 1GB of memory running Ubuntu Linux, 60GB 2.5 IDE inch hard disk, USB GPS receiver based on the SiRF chipset, Compex WLM54AGP23 802.11abg mini PCI cards using the Atheros AR5413 chipset (upgraded from 802.11b Prism2-based USB WiFi dongles), 802.11g wireless AP used as a bridge to an Ethernet port on the OpenBrick, Sierra Wireless 881 or 885 3G USB Modems operating on the AT&T network (upgraded from a MultiTech GPRS modem attached to a serial port), Digi XTend 900MHz USB RF modem, and an inverter to convert 24VDC to 120VAC.



Figure 2: DieselNet equipment in a bus locker

We have installed the DieselNet node, which we generically refer to as a *brick*, in two different locations in the buses. In older buses, we have installed the equipment above the driver’s head, behind an electronic sign that displays the bus route. Newer buses contain a special locker to hold electronic equipment including video surveillance and radios, as shown in Figure 2. The locker also contains removable trays that gives us easier access to the equipment over the older buses. In both cases, we place antennas in the best available locations without drilling through the bus exterior.

A brick’s AP allows other buses, or bus riders, to establish 802.11 connections into the brick, giving them access to the Internet via the 3G modem, or to a locally cached web page when disconnected. The WiFi interface is used by a brick to connect to foreign APs, including the APs on other buses. The SSID broadcast by a brick’s AP allows it to be identified as belonging to a DieselNet bus by other buses, throwboxes, and any other DOME device.

Throwboxes: Throwboxes are wireless nodes that can act as relays or mesh nodes, creating additional contact opportunities among DieselNet buses [7, 8]. They are essentially nomadic nodes allowing for flexible placement in the DOME testbed. Unlike the vehicular nodes, the throwboxes use batteries recharged by solar cells. Also in contrast to the vehicular nodes, a throwbox will often remain stationary for several hours or days. We have deployed throwboxes within the testbed by attaching the boxes to the front of bicycles and then locking the bicycles to bicycle racks. This setup gives us the ability to easily reconfigure nodes in the testbed to support different placements and functions. While we can place the nodes virtually anywhere, we have typically used an algorithm that has provably good performance [40].

Each throwbox contains a Crossbow Stargate board with a 32-bit, 400MHz PXA255 XScale processor (having 64MB of RAM and 32MB of internal flash memory on board), TelosB Mote with an 8-bit 8MHz microcontroller with 10KB of RAM and 1MB of external flash memory, D-Link Air 802.11b CF WiFi card, modified to support a second external antenna, DiGi XTend 900MHz OEM module attached to the Mote, Three 1.2V 10Ah batteries (though we also have configurations with a single 1Ah battery), a custom board with a Maxim DS2770 fuel gauge chip to monitor the batteries, a Maxim 2751 current-sense amplifier to monitor power consumption, and two 5V PowerFilm solar panels as an additional power source.

Mesh Network: In cooperation with the Town of Amherst, we have installed 26 Cisco 1500-series WiFi APs. These are lightweight APs, managed by a central controller, and they support seamless hand-off for mobile nodes. The APs use two radios: an 802.11g radio for the public and mobile nodes to connect to, and an 802.11a radio to mesh APs together.

The nodes are mounted on a variety of town-owned buildings and light poles. While both locations provide power, only the buildings provide connectivity to the local fiber infrastructure. Consistent with research findings [13] and Cisco’s direction, the network is laid out such that there are never more than three hops without connectivity to the wired network. We have also installed a Procera PacketLogic packet logger and traffic shaping box for monitoring the network. This yields statistics about users, node mobility, and traffic patterns.

IV.B. Software Components

Virtual Machine: When we initially developed DOME the only users of the DieselNet Vehicular Nodes were people associated with the University. As DOME became part of GENI and was made available to researchers around the world, we needed to ensure that rogue experiments could not corrupt the bricks. We implemented Xen virtual machines on our nodes. An active experiment in DOME is an instance of a virtual machine. Experiments are distributed to the nodes as files formatted as disk images. When DOME starts a virtual machine, it mounts the appropriate disk image. We also support experiments providing disk images with their own customized kernels.

Experimenter Portal: The DOME experimenter portal allows researchers to upload experiments (disk images). Researchers use the portal to schedule when their experiments will run on the test bed. The DOME portal interacts with the GENI ORCA management

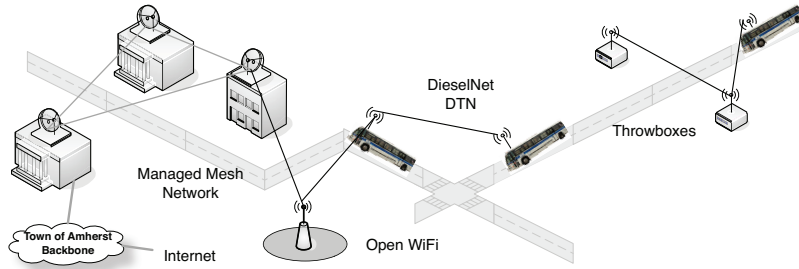


Figure 3: An overview of the DOME testbed.

framework [14] to allocate slices. The bricks interact with the portal to install experiments, and to get leases issued by ORCA. The leases authorize a specific experiment to execute on a defined number of buses.

DieselNet Control Plane: The computers on the buses have many services active to ensure the proper execution of experiments. These services include prestaging and garbage collecting experiments, and launching and terminating virtual machines based on the availability of leases. Another service manages the 3G modem, testing connectivity and reestablishing connections if necessary. The control plane also allocates resources to experiments. This includes exporting the WiFi card’s PCI address to the virtual machine, tunneling XTend radio data between the virtual machine domain and the control plan domain, configuring the network so that VM LAN traffic is routed over the 3G link, and providing read access to the GPS device.

Remote Update Mechanism: We provide a mechanism for updating the DieselNet Control Plane software. This is done by publishing an update on a DOME server, and having the nodes periodically use available 3G connectivity to check for any updates. A challenge to updating the software in our DieselNet testbed is that we have no control over when a system shuts down; a brick simply loses power when a bus’s engine is turned off. A concern is losing power during a critical section of an update, which can render the brick unusable until we are able to physically access it and make manual repairs. Our solution is to use a shadowing approach, isolating a copy of the software being updated in a directory and preserving atomicity by modifying a single symbolic link.

Logging: Logging of experimental results presents a challenge in a disconnected network. Results must be written to an intermediate but persistent location, and then offloaded from the bus when the network is available. The challenge is exacerbated since the user has no control over when the bus is turned off. DOME provides a logging service to experiments. DOME

ensures that data is eventually offloaded, even after an experiment’s lease has expired.

Maintenance Monitoring: A monitoring service allows us to track the health of the testbed and know how many nodes are operational and which peripherals are malfunctioning. Even if components fail, the DieselNet monitoring software will attempt to establish connectivity to the DOME servers to provide notification. We also correlate vehicular node activity with the bus schedules, allowing us to detect nodes that have no connectivity or that do not boot.

V. Characterization of DOME

One of our primary goals for DOME is ensuring that the testbed is capable of supporting a wide spectrum of research in mobile systems. In our previous work, we have demonstrated that DOME’s components are capable of supporting research efforts ranging from energy management [7, 20, 40] to routing [3, 4, 8, 11, 12, 39] to security [10] to application enhancement [5, 6]. In this section, we use traces from the project to answer a series of open questions specifically not addressed in our prior work.

- Has WiFi been trending towards ubiquitous deployment and open availability outdoors? Have WiFi APs been increasingly restricted, and what impact has this had for research and services that rely on opportunistic connections? How has coverage been affected by WiFi technological improvements?
- What is the impact of WiFi AP density on usable throughput from vehicles?
- What are the quantifiable and relative strengths of commercial 3G infrastructure for wireless Internet service over organically available WiFi connectivity?
- How does the spatial diversity of a network impact the available throughput from a vehicle?

- What are the relative strengths of a WiFi system specifically planned and deployed for outdoor coverage, versus organic connectivity deployed for home use?

V.A. Evaluation Methodology

The DOME testbed has collected a number of logs from May 2004 to October 2008 that we have used to derive the results in this section. As mentioned earlier, while the testbed has been operational since 2008, we stopped constant collection of the presented data in support of conducting other experiments. The most crucial of these logs is a list of contacts between vehicles and other vehicles, as well as between vehicles and the infrastructure. Logs include duration, GPS location, and speed at the beginning and end of every contact. Since September of 2007, the nodes also collected the number of APs seen in each scan, as well as what portion of those employ some form of access restriction (e.g., WEP or WPA). Also since September of 2007, the vehicles collected additional information about contact with APs, including successful associations and DHCP leases. During short term tests, we have deployed measurement apparatus to measure the fraction of time a node spends connected to a cellular network (GPRS and 3G) and connected to WiFi APs.

We have been able to answer most of the questions posed above using data originally collected for markedly different purposes. In other cases, we have deployed short-term experiments for additional clarification in this paper. To answer certain questions we have made do with incomplete data, demonstrating the principle importance of tracking the quality of data with the data itself. Over the lifetime of the testbed many changes have occurred, including: planned software maintenance and upgrades, replacement of hardware, long periods of neglected maintenance resulting in reduced data points, increases and decreases in log fidelity and measurements, and dedicated reservation of the testbed for individual researchers.

We divide the area that the vehicles and infrastructure inhabit into $100\text{m} \times 100\text{m}$ regions. This size is on the order of the range of a WiFi AP. We do not know the true locations of the majority of organic APs; thus, connecting to an AP at a mobile node's GPS location is only a rough measure of the AP's actual location.

In all cases, we have removed the effects of the varying number of vehicles operating, such as summers and vacations, which have a much reduced bus schedule, and aberrant vehicle behavior, such as temporary use of a bus for a field trip. For experiments that depend on regions, if there are less than 30 visits to a region

during a month, we discard all measurements from that region, and normalize the results based on the number of remaining regions. Given the scale of the testbed—there are often 30 or more buses operating 18 hours per day—the testbed has yielded an enormous amount of data. Our presented dataset includes 8,679,179 contact attempts between our 40 vehicles and 28,776 unique APs; of those attempts, 2,110,595 were successful. In addition, 1,091,307 successful contacts between vehicles on the road occurred.

V.B. Organic WiFi

A number of research projects have proposed the use of organic, open WiFi APs for opportunistic networking, particularly for vehicular networking [5, 16, 18]. While the number of deployed APs has certainly increased over recent years, it does not imply that coverage has dramatically improved, as many of the additional APs may have been deployed in the same region or are not open to the public. To quantify this trend in our area, we analyzed our traces from January of 2005 to October 2008 to find which $100\text{m} \times 100\text{m}$ regions had at least some connectivity, meaning that at least one successful ping was sent to our server during that month from a bus in that region. A plot of that analysis is shown in Figure 4. Measuring the number of regions that have some connectivity is somewhat complicated by several changes that have occurred in our testbed as shown in the timeline in Figure 5. (See our discussion of Tenets 1 and 2 in Section III.) The strength of this data is that it is a longitudinal study over a diverse geographic region (c.f., shorter tests over a more homogeneous set of regions [18]).

From January 2005 to May 2008 we used USB 802.11 interfaces with the Prism2 chipset—these interfaces exhibit range similar to what one might find in a laptop computer. Over the course of January 2005 to May 2006, the vehicles only found connectivity in 20% to 40% of the regions, with no significant trend over the course of that year and a half. This data demonstrates that building a mobile application on top of such a system would have exhibited significant disconnections and outages. During the summer of 2006, we added the first cellular modems but failed to disambiguate connections through the two interfaces, so that data is omitted from the graph. During the fall of 2006, we continued to use USB WiFi devices, and still found that more than 50% of the regions had no WiFi coverage whatsoever, although the increase in use of wireless APs brought coverage to almost double the number of regions that were covered in early 2005.

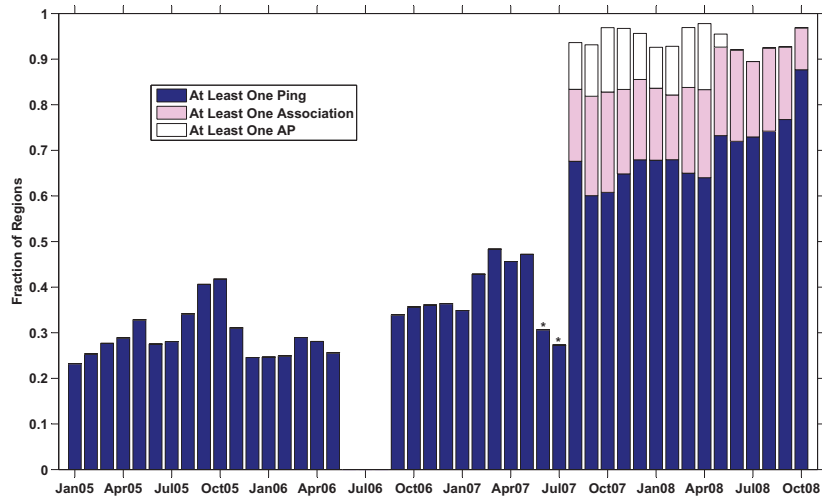


Figure 4: This graph shows the number of regions that the vehicles frequent that could support at least one successful ping through WiFi during a month. The summer of 2006 is omitted due to software problems, and the summer of 2007 began access to the UMass and Town WiFi networks. From September 2007 on, we show which regions supported at least one association or a scan with at least one AP (but not a successful ping).

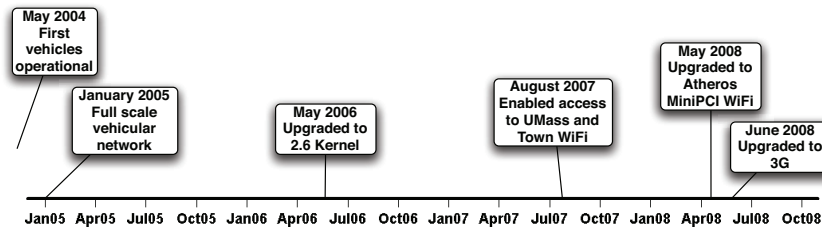


Figure 5: The timing of significant hardware and software changes made to the DOME testbed.

During the summer of 2007, we installed and enabled access to the Town Mesh (the network is open, but requires a click-through agreement that we worked around) and enabled access to the UMass wireless network. This required several months to fully take effect with full access starting in September of 2007. At that point, we started to track pings, associations and regions where scans revealed APs, but no connection was possible. In May of 2008, we upgraded from USB 802.11 interfaces to Atheros MiniPCI cards with external antennas. This yielded increased range and further improved the number of regions covered. The collection of improvements shows that given the proper hardware in an environment with a combination of managed and open APs connectivity can reach nearly 90% of regions. However, many environments do not have the benefit of a deployed infrastructure and will see much less coverage. Given applications that send relatively short messages, or are insensitive to throughput, organic WiFi will provide sufficient coverage in our environment.

There has been anecdotal speculation that while the use of WiFi is expanding, APs are increasingly protected by encryption as setup becomes easier and there is more attention paid to the importance of encryption. The previous experiment answers part of the question: encrypted APs do not seem to have significantly impacted coverage as far as our data-set is concerned. However, the next questions are what portion of APs have been encrypted, and if there has been any noticeable trends. We show the results of analyzing our traces for unencrypted versus encrypted APs in Figure 6.

The overall number of APs discovered per scan increases as we upgraded the WiFi interface on the vehicles. However, the increased range also discovered an increased proportion of encrypted APs, but did not show a noticeable trend of open APs disappearing. We speculate that the increased proportion of encrypted APs can be explained by the increased range of the radios capturing more residential APs, as many of the open APs are located in businesses with free WiFi.

We have also tracked the lifetime of organic APs in the network. In Figure 8, we show the amount of

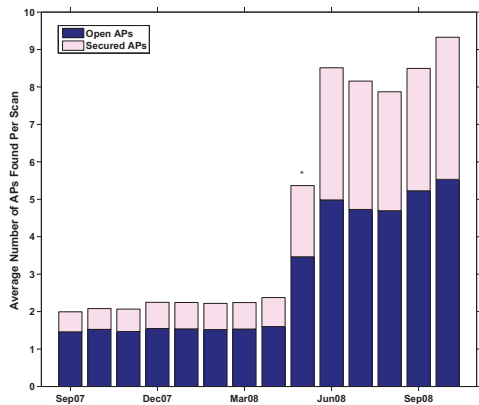


Figure 6: The number of APs found per scan over a 13 month period, broken down by encrypted APs versus non-encrypted APs. The increase in overall APs in the summer of 2008 is due to replacing the USB WiFi interfaces with MiniPCI cards with external antennas.

time between the first and last sighting of each organic AP. We have limited this dataset to useable APs that permitted ping messages on at least 10 different occasions to avoid spurious contacts with distant APs. The results show that open APs exhibit an extremely high amount of churn, with most APs only visible for less than three months. Some of this effect is certainly due to the transient nature of college students, but other effects may include users who were dissatisfied and returned their APs or users who quickly upgraded their hardware. However, there remain a number of APs that have been available for up to three years.

In Figure 7, we demonstrate the spatial diversity of DOME by plotting an estimate of the aggregate throughput available from a vehicle in different regions, sorted by the total amount of throughput available from that region. The data presented in this graph comes from an experiment from August 2008 to November 2008 in which the buses exclusively connected to APs and not to other buses. For each region, we have calculated the average per-connection throughput and then multiplied these values by the availability of bus-to-ap links which we calculated as 12.3% for the duration of this experiment. The results demonstrate the diversity of the DOME testbed for experimenters, providing regions that are well-connected to the infrastructure, as well as regions that are poorly connected.

V.C. Density Effects

As DOME has collected data from a wide variety of geographic regions, we can also examine the relationship between the density of APs found in a region to the

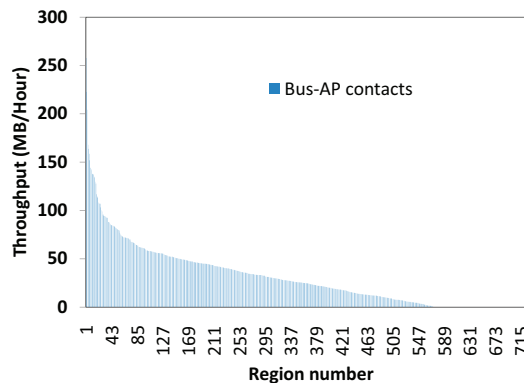


Figure 7: Vehicle-to-AP throughput for each geographical region during experiments from August 2008 to November 2008. The regions are sorted by the measurement of the throughput. The throughput values are normalized by 12.3% availability of bus-to-AP WiFi links.

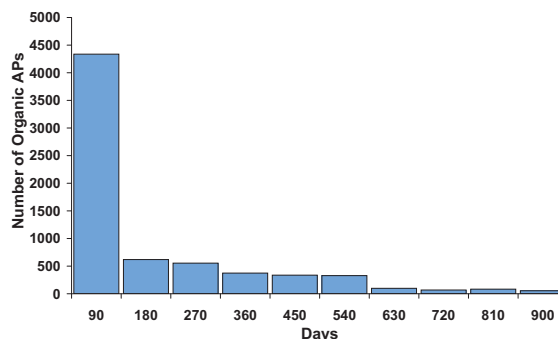


Figure 8: This shows the lifetime of open, organic APs in DOME that permitted pings to the Internet on at least 10 occasions. The results show that organic APs exhibit a large amount of churn, with most existing less than three months within the network.

usable throughput between a vehicle and a single AP. Due to the relatively small number of non-interfering channels to choose from, a primary concern in the dense deployment of APs has been the possibility of interference between nodes. Using five months of data collected from June 2008 to October 2008, we plot the throughput attainable to a single AP. We divide these results into bins by the aggregate number of unique APs discovered in that region during an entire month. The results are shown in Figure 9. Note that past 120 APs, there is considerably less data to form conclusions, even over a five month period.

The results show several effects. First, for smaller AP densities, the achievable throughput generally lies between 50 and 125 kB/sec; however, the large number of outliers show that in low density environments it is possible to achieve much higher throughput, but

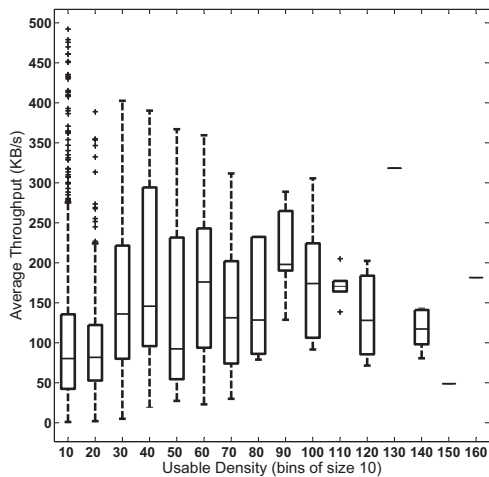


Figure 9: The correlation between WiFi throughput and the density of APs in a region for a period of five months from June 2008 to October 2008. The box plots shows the first and third quartiles, while the center line shows the median. The whiskers show the maximum and minimum values, while the stars demonstrate extreme outliers.

not predictably. As the density of APs grows there is generally more throughput available, as the vehicle can choose between a great number of APs and is more apt to select those with greater signal strengths. However, as the number of APs grows to as high as 120 or more, throughput is generally as good as lower numbers of APs. With the density of APs available in the DOME testbed, which we believe to be fairly high, we have been unable to demonstrate appreciable negative effects from interference in real-world settings.

V.D. A Comparison of WiFi, GPRS, and 3G

Opportunistic WiFi offers the opportunity for no-cost access, but it may provide less reliable access and lower aggregate throughput. Using an experiment deployed on the buses in November 2007 with USB Prism2 WiFi and GPRS, and a second experiment in November 2008 using Mini PCI Atheros WiFi and 3G, we compare the overall availability and upstream and downstream throughput of each of the interfaces. This experiment, shown in Figure 10, demonstrates several hardware trends.

The results show that the availability of GPRS and 3G in the DOME testbed is excellent; however, in 2007 GPRS and WiFi provided very comparable aggregate throughput over the course of a day. In 2008, the overall availability of WiFi connectivity from the vehicles

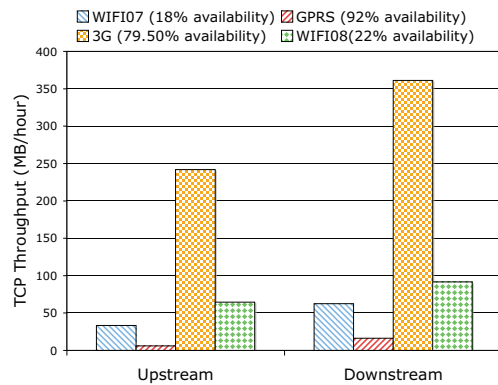


Figure 10: The amount of throughput (upstream and downstream) available from USB WiFi interfaces and GPRS during a test in 2007 and MiniPCI WiFi interfaces and 3G during a test in 2008. The legend shows the percentage of time that each interface was connected to the Internet. WiFi availability would have to increase to roughly 90% to surpass the throughput of 3G.

had substantially improved (due to the increased range of the Atheros radios), but the overall throughput during the course of a day lags behind 3G. To meet the overall throughput of 3G, WiFi would need greatly expanded coverage to give connectivity nearly 90% of the time—we believe this to be generally infeasible given the amount of time needed to search and associate with APs.

V.E. A Comparison of Organic WiFi and Planned WiFi

DOME incorporates three types of WiFi networks: the Town of Amherst managed mesh network planned and deployed specifically to support DOME, with APs mounted outside and directly over the roadway; the managed UMass WiFi network, which is a planned deployment, but was deployed primarily for indoor access with some outdoor coverage; and the organic WiFi APs which are unmanaged and were deployed in an ad hoc manner. To examine one aspect of these deployments, we measured the durations of the WiFi connections from vehicular nodes. It is clear that a planned network should have greater connectivity durations, as managed networks are capable of seamlessly roaming between APs, but quantifying the effects of planned versus organic networks remains an open question. The results of this experiment are plotted in Figure 11, which shows the durations of connectivity using the MiniPCI Atheros WiFi interface in 2008.

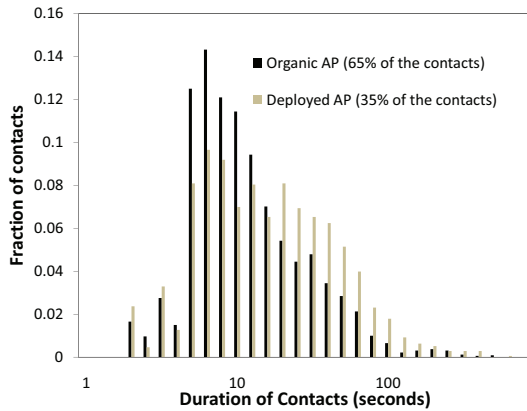


Figure 11: The duration of connections for the vehicles using the MiniPCI WiFi interfaces during a test in November 2008. The connections are separated by organic APs and managed APs (Town of Amherst mesh and UMass wireless network).

The results show that most connections to open APs cover 40 seconds or less, while connections to the managed infrastructure sustain higher durations typically 100 seconds or less. However, one concern is that there is some bias in the speed of the vehicles in relation to the APs: the buses move at a slower rate of speed while in downtown Amherst and near the campus where the planned networks were installed. To examine this effect, we combine the previous data with data about the speed of the vehicles while connecting to those APs. We use the average of the speed of the vehicle at the beginning and end of the connection. We then corrected for speed by normalizing the data to the overall average speed of the buses (14.4 km/hr), and plot the results in Figure 12. Note that all of the static connections during which the speed of the bus was zero have been omitted in that graph. The legend shows what portion of the connections fall into that category. While the overall distribution does not change very much, a larger portion of the organic contacts are made while the vehicle is stationary. We believe this is due to organic APs not supporting roaming, and thus the managed nodes are more frequently used while the vehicle is moving. However, most of the contact time is spent using organic APs, marking them as a vital contributor to WiFi connectivity.

V.F. Summary of Results

In this section we have presented a number of results from the DOME testbed that highlight the goals of enabling the measuring temporal trends, technological diversity, and spatial diversity.

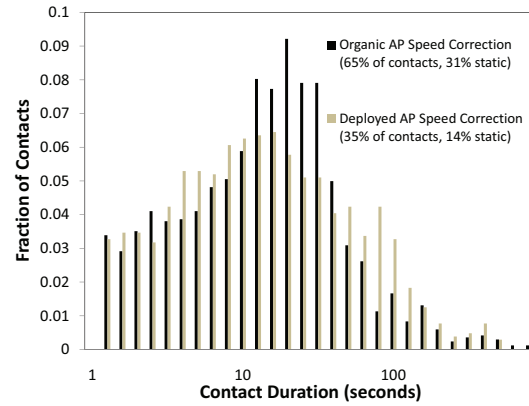


Figure 12: The duration of WiFi connections normalized by the average speed during each connection. The connections to the managed WiFi infrastructure yield longer connections.

We have shown that a network of organic and managed APs covers nearly 90% of the regions visited by the testbed vehicles. Although more than 40% of the APs were encrypted, the data showed no indication that this proportion increased over a year. The testbed covers a large geographic region, from micro-urban to rural, enabling a diversity of access profiles. Some regions have extremely high throughput—more than 150 MB/hr—and some near zero. Thus, DOME can support experiments that depend on high throughput and dense connectivity, or regions of sparse connectivity and low contention. We have also presented one such experiment, based on collected data, that demonstrates the correlation between AP density on per-connection throughput. At low densities throughput is impacted as the small number of APs limits finding APs with strong signals. At larger densities, interference limits throughput, but the throughput is comparable to regions with a small number of APs, showing the limited effect that interference has on vehicular throughput.

DOME supports a wide variety of radio technologies, including historical results from previously installed hardware. As recent as 2007, an older USB WiFi interface provided much greater aggregate throughput than cellular GPRS, even through the GPRS connection was four times more available. However, in 2008, after upgrading the WiFi interface to a more reliable and powerful MiniPCI card, and the cellular modem to 3G, the results were very different: 3G provides four times as much aggregate throughput as WiFi in our mobile setting. DOME also includes a WiFi network deployed specifically for vehicular use, as well as a managed University WiFi network, and a large set of organic APs. We have shown results that highlight the ability to compare organic APs and

managed APs, focusing on the managed APs' ability to maintain longer connections through WiFi roaming.

VI. Related Work

There is a growing number of academic mobile testbeds with similar goals to ours, and all are addressing a critical need within the mobility community [22, 29]. These testbeds include CarTel [18], the Drive-Thru Internet platform [28], CitySense [25], VanLan [24], Orbit [33], DakNet [30] and KioskNet [35]. While CarTel, Drive-Thru Internet, VanLan, and CitySense are focused on urban environments, DakNet and KioskNet are deployments in very sparse, developing regions. Orbit is a *wireless* testbed, in that it uses actual wireless nodes, but the nodes are fixed and mobility and interference are simulated effects. We see great value in both mobile and wireless testbeds; however, it is clear that many of the insights we have gained from building DOME could not have been discovered with a fixed deployment. There are many other short-term experiments that have been done whose traces have been collected at <http://traces.cs.umass.edu> and the Dartmouth CRAWDAD site.

VII. Future Work and Conclusions

We reviewed many challenges and lessons learned from our deployments and experiments with the DOME testbed. While this paper highlights DOME's varied capabilities, we are constantly expanding the technological diversity of DOME. In particular, we have deployed a WiMAX base station and installed WiMAX clients on several DieselNet nodes for studying cellular-like connections from both endpoints.

VIII. Acknowledgments

This work is supported in part by NSF awards CNS-0519881, CNS-0447877, DARPA award W15P7T-05-C-P213, and the NSF GENI program. We would like to acknowledge John Burgess for much of the early development of DieselNet and Aruna Balasubramanian for her contributions to the data that we have presented. We would also like to thank Kristopher Pacunas, David Soucie, and other officials of the Town of Amherst, Al Byam, Adam Sherson, Jeff Jason and others at the Pioneer Valley Transit Authority, and Randy Sailor, Chris Misra, Jake Cunningham, and Michael Dickson and others at the UMass Amherst Office of Information Technology for their partnership in the deployment of DOME.

References

- [1] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren. Analysis of a Mixed-Use Urban WiFi Network: When Metropolitan becomes Neapolitan. In *IMC*, pages 85–98, October 2008.
- [2] J. Albrecht, C. Tuttle, A. C. Snoeren, and A. Vahdat. Planetlab application management using plush. *SIGOPS Oper. Syst. Rev.*, 40(1):33–40, 2006.
- [3] A. Balasubramanian, B. N. Levine, and A. Venkataramani. DTN Routing as a Resource Allocation Problem. In *Proc. ACM Sigcomm*, August 2007.
- [4] A. Balasubramanian, B. N. Levine, and A. Venkataramani. Enabling Interactive Applications in Hybrid Networks. In *ACM Mobicom*, September 2008.
- [5] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. Levine, and J. Zahorjan. Interactive WiFi Connectivity for Moving Vehicles. In *ACM SIGCOMM*, August 2008.
- [6] A. Balasubramanian, Y. Zhou, W. B. Croft, B. N. Levine, and A. Venkataramani. Web Search From a Bus. In *Proc. ACM Workshop on Challenged Networks (CHANTS)*, September 2007.
- [7] N. Banerjee, M. D. Corner, and B. N. Levine. An Energy-Efficient Architecture for DTN Throwboxes. In *IEEE Infocom*, May 2007.
- [8] N. Banerjee, M. D. Corner, D. Towsley, and B. N. Levine. Relays, Base Stations, and Meshes: Enhancing Mobile Networks with Infrastructure. In *Proc. ACM MobiCom*, September 2008.
- [9] V. Borrel, M. H. Ammar, and E. W. Zegura. Understanding the Wireless and Mobile Network Space: A Routing-centered Classification. In *Proc. CHANTS*, pages 11–18, 2007.
- [10] J. Burgess, G. Bissias, M. D. Corner, and B. N. Levine. Surviving Attacks on Disruption-Tolerant Networks without Authentication. In *Proc. ACM Mobihoc*, Montreal, Quebec, Canada, September 2007.
- [11] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM*, April 2006.
- [12] B. Burns, O. Brock, and B. N. Levine. MORA Routing and Capacity Building in Disruption-Tolerant Networks. *Elsevier Ad hoc Networks Journal*, 2008.
- [13] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *ACM MobiSys*, pages 96–109, 2006.
- [14] J. Chase, L. Grit, D. Irwin, V. Marupadi, P. Shivam, and A. Yumerefendi. Beyond virtual data centers: Toward an open resource control architecture. In *Selected Papers from the International Conference on the Virtual Computing Initiative (ACM Digital Library)*, May 2007.

- [15] D. J. Goodman, J. Borras, N. B. Mandayam, and R. D. Yates. INFOSTATIONS: A new system model for data and messaging services. In *Proc. Vehicular Technology Conference*, pages 969–973, Phoenix, AZ, May 1997.
- [16] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal. Vehicular Opportunistic Communication Under the Microscope. In *MobiSys*, pages 206–219, 2007.
- [17] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket Switched Networks and Human Mobility in Conference Environments. In *Proc. ACM Workshop on Delay-Tolerant Networking*, pages 244–251, Aug. 2005.
- [18] 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. In *ACM SenSys*, October 2006.
- [19] S. B. John Bicket, Daniel Aguayo and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *Mobicom*, August 2005.
- [20] H. Jun, M. H. Ammar, M. D. Corner, and E. Zegura. Hierarchical Power Management in Disruption Tolerant Networks with Traffic-Aware Optimization. In *Proc. ACM SIGCOMM Workshop on Challenged Networks (CHANTS)*, September 2006.
- [21] S. Kandula, K. C.-J. Lin, T. Badirhanli, and D. Katabi. FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In *Proc. NSDI*, April 2008.
- [22] D. Kotz, C. Newport, R. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental Evaluation of Wireless Simulation Assumptions. In *IEEE MSWiM*, pages 78–82, Oct 2004.
- [23] H. Lundgren, K. Ramach, E. Belding-royer, K. Almeroth, M. Benny, A. Hewatt, E. Touma, and A. Jardosh. Experiences from the Design, Deployment, and Usage of the UCSB MeshNet Testbed. *IEEE Wireless Communications*, 13(2):18–29, April 2006.
- [24] R. Mahajan, J. Zahorjan, and B. Zill. Understanding WiFi-based Connectivity From Moving Vehicles. In *Proc. IMC*, 2007.
- [25] R. Murty, A. Gosain, M. Tierney, A. Brody, A. Fahad, J. Bers, and M. Welsh. CitySense: A vision for an urban-scale wireless networking testbed. Technical report, Harvard University, 2007.
- [26] R. Murty, G. Mainland, I. Rose, A. R. Chowdhury, A. Gosain, J. Bers, and M. Welsh. CitySense: A Vision for an Urban-Scale Wireless Networking Testbed. In *IEEE International Conference on Technologies for Homeland Security*, May 2008.
- [27] A. Nicholson, Y. Chawathe, M. Chen, B. Noble, and D. Wetherall. Improved access point selection. In *Proc. MobiSys*, 2006.
- [28] J. Ott and D. Kutscher. Drive-Thru Internet: IEEE 802.11b for “Automobile” Users. In *Proc. IEEE INFOCOM*, pages 362–373, March 2004.
- [29] K. Pawlikowski, H.-D. Jeong, and J.-S. Lee. On credibility of simulation studies of telecommunication networks. In *IEEE Communications Magazine*, volume 40, pages 132–139, Jan 2002.
- [30] A. Pentland, R. Fletcher, and A. Hasson. DakNet: Rethinking Connectivity in Developing Nations. *IEEE Computer*, 37(1):78–83, Jan 2004.
- [31] L. Peterson and T. Roscoe. The design principles of PlanetLab. *ACM SIGOPS Operating Systems Review*, 40(1):11–16, January 2006.
- [32] B. Plale, D. Gannon, J. Brotzge, K. Droegemeier, J. Kurose, D. McLaughlin, R. Wilhelmson, S. Graves, M. Ramamurthy, R. D. Clark, S. Yalda, D. A. Reed, E. Joseph, and V. Chandrasekar. CASA and LEAD: Adaptive Cyberinfrastructure for Real-Time Multi-scale Weather Forecasting. *Computer*, 39(11):56–64, 2006.
- [33] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-Generation Wireless Network Protocols. In *WCNC*, March 2005.
- [34] P. Rodriguez, I. Pratt, R. Chakravorty, and S. Banerjee. Mar: A commuter router infrastructure for the mobile internet. In *In Proc. of ACM Mobisys*, 2004.
- [35] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *Proc. ACM MobiCom*, pages 334–345, Sept. 2006.
- [36] H. Soroush, N. Banerjee, A. Balasubramanian, M. D. Corner, B. N. Levine, and B. Lynn. DOME: A Diverse Outdoor Mobile Testbed. In *Proc. ACM Intl. Workshop on Hot Topics of Planet-Scale Mobility Measurements (HotPlanet)*, June 2009.
- [37] N. A. Thompson, P. Zerfos, R. Sombrutzki, J.-P. Redlich, and H. Luo. 100% Certified Organic: Design and Implementation of Self-Sustaining Cellular Networks. In *Proceedings of ACM HotMobile*, Napa Valley, CA, United States, February 2008.
- [38] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An Integrated Experimental Environment for Distributed Systems and Networks. In *Proc. OSDI*, pages 255–270, Boston, MA, 2002.
- [39] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a Bus-Based Disruption Tolerant Network: Mobility Modeling and Impact on Routing. In *Proc. ACM Mobicom*, September 2007.
- [40] W. Zhao, Y. Chen, M. Ammar, M. D. Corner, B. N. Levine, and E. Zegura. Capacity Enhancement using Throwboxes in DTNs. In *Proc. IEEE Intl Conf on Mobile Ad hoc and Sensor Systems (MASS)*, pages 31–40, Oct 2006.