

Data Muling with Mobile Phones for Sensornets*

ISI Technical Report ISI-TR-673

Unkyu Park and John Heidemann

{ukpark,johnh}@isi.edu

Information Sciences Institute

University of Southern California

Abstract

Sensors are all around us, in buildings, vehicles and public places, from commodity thermostats to custom sensor-nets. Yet today these sensors are often disconnected from the world, either because they are distant from infrastructure, and wide-area networking (by 3G cellular, satellite, or other approaches) is too expensive to justify. Data muling makes communication cost-effective by leveraging short-range wireless and mobility, perhaps by zebras, buses or farmworkers. In this paper we propose that *human-carried mobile phones* can serve as data mules for sensornet deployments, exploiting ubiquity of mobile phones and human mobility to bring low-cost communication to sensors. We use two mobile phone datasets to show that Bluetooth can serve as a viable muling network, and humans already see many potential sensors regularly. We have implemented a mobile-phone-based data muling system, and used it in four sensornet deployments totaling ten months operation. We find that muling can be the only cost-effective option for rural deployments, where it is critical to monitoring remote sensor networks. We also show opportunistic mobility can collect data without any extra effort in residential and office environments. Finally, we systematically evaluate our deployments to understand how contact duration and data size interact, and to evaluate the effect of muling on phone batteries.

1 Introduction

Yet our world is increasingly instrumented. Electronics generate useful information: a modern office building has a thermostat at every room, a motion detection sensor at the door; automobiles have dozens of sensors monitoring the engine, the passengers, and the outside environment. Today these devices serve their original, dedicated purposes—yet what if their data were made available through low-cost communication?

Sensor networks begin with the premise that sensors can communicate. SCADA systems have been used in industry for decades, and industrial SCADA systems and sensornets are increasingly sophisticated, but often price limits their use

to high value applications. Similarly, scientific uses of sensornets are growing, but connectivity to remote locations is expensive, limiting less fortunate scientists to datalogging without experiment supervision. How can we make wide-area communication more accessible to sensornets?

Wifi, cellular telephone data networks, and satellite communication make wide-area communication commonplace. But these approaches do not solve these problems because they are often too expensive or unavailable where needed. Satellite data is often prohibitively expensive. Cellular seems more affordable, but monthly fees for data service are often too high for many applications. (Even services such as the Amazon Kindle that are free to the user embed data fees in the cost of content, discouraging high-bandwidth content [12].) And if cost is not considered, *coverage* remains a problem. Wifi is free, but coverage can be spotty even in urban areas. Satellite coverage requires line-of-sight to the sky. Cellular coverage is generally good, but all providers have dead spots. Humans work around coverage problems by moving, but that option is not available to fixed sensors. Some sensornets use a local mesh to get to wide-area connectivity, but that approach greatly adds to complexity and failure cases for simple, small deployments.

In this paper we describe our system where *human-carried mobile phones* serve as data mules for sensornet deployments, exploiting ubiquity of mobile phones and human mobility to bring low-cost communication to sensors. Many other groups have explored the idea of data muling for sensor networks [21, 8, 28, 25, 3, 29, 20, 35], and some have proposed human mobility for communication in remote areas [18], in disasters [11] or from cars [19]. Our work makes three contributions beyond the prior work. First, we bring together sensing and the mobile phone as a data mule in our implementation of muling in an off-the-shelf mobile phone (Section 3). We also describe several applications where muling's cost reduction makes new applications viable or current applications easier to justify (Section 2.2).

Second and more importantly, we have used our phone-based mule in four sensornet deployments, a field deployment for subsidence detection in an oilfield, an experimental analog of that application in an urban area, and an office-based person-detector (Section 4). As far as we know, we are the first to use human-carried mobile phones in gathering data from real sensor deployments in remote infrastructure-less area for an extended time of period. Together these ten

* This research is partially supported by the National Science Foundation (NSF) under grant number CNS-0626702, "Sensor-Internet Sharing and Search" and by CiSoft (Center for Interactive Smart Oilfield Technologies), a Center of Research Excellence and Academic Training and a joint venture between the University of Southern California and Chevron Corporation.

months of deployment experience help us understand what makes data muling work in practice. We show that muling provides essential feedback for experimental deployments in remote areas (Section 6.1.2), halving the time outages of experimental hardware were unknown, from 60 sensor-days to 27. We also evaluate when energy consumption is a limiting factor (Section 6.4).

Our last contribution is to explore how human mobility patterns affect the potential of data muling. We examine two datasets of mobile phone contacts to show that humans see many potential sensors (Section 5.1), and some of these regularly (Section 5.2). Yet our deployments show that *intentional* mobility is often required when coverage of specific sensors is required, at least with ranges typical for Bluetooth (Section 6.1). We show the importance of the human's loiter time to effective muling (Section 6.2), and that that opportunistic mobility works best in our office deployment where sensors are dense and loiter times are long (Section 6.1.4). We also consider an alternative communication choice with long range radio and fast data rate, showing that it benefits the data muling and make opportunistic muling even more practical (Section 6.3).

2 Motivation, Applications, and Challenges

In this section we describe why mobile phones can make good data mules, then show potential applications that motivate phone-based data muling.

2.1 Why Data Muling using Mobile Phones?

Both sensors and people with mobile phones are all around us. However, the cost of bringing sensors on to the wide-area network is often prohibitive. We suggest that mobile phones can bridge this connectivity gap through data muling.

Mobile phones are attractive as data mules for three reasons. First, mobile phones are truly ubiquitous. In fact there are approximately 4.6 billion mobile phone users worldwide estimated by the International Telecommunication Union. That means 68% of the world's population already carry mobile phones all the time. And mobile phones are widely used in the developing world where the need for data muling is greatest since other forms of wide-area communication are often limited. Although currently most phones in the developing world are feature phones with limited extensibility, in principle even these telephones could support muling, and we expect phone capabilities to grow.

Second, current generation smart-phones are powerful, general purpose computing platforms. They are already equipped with energy-conserving, short-range radio networks like Bluetooth, and with wide-area Internet connectivity through 3G and now 4G telephony. We show later (Section 6.4) that we can use these networks for muling with minimal additional energy cost.

Third, mobile phones are already carried by humans as part of their daily life, so muling can easily piggy back on this mobility for free.

Finally, the large display and physical or virtual keyboard of modern smartphones provide a friendly interface to sensors. Many embedded sensors lack a sophisticated interface or on-site control, and as sensors become smaller, lower-

power, and cheaper, user interfaces become impossible to provide. We anticipate the mobile phone can be used a unified interface to various sensors.

2.2 Motivating Data Mule Applications

We next describe five applications that are good matches for data muling. Their common characteristic is the need for sensing and delay-tolerant communication with the outside world, coupled with a lack of cheap communication, and the presence of human mobility. With near-ubiquity of mobile phone coverage, it may seem that wireless coverage should be always available. However, we show that in our Subsidence/Oilfield application, cellular data coverage was so poor as to be unusable (Section 4.2), and even when cellular coverage is good, its price or energy draw may be too high to justify its use.

Assisted-reporting Garbage Bins: Today garbage is removed from bins in national parks and urban streets with a fixed, periodic schedule. A fixed schedule works poorly when bursts of use fill bins unexpectedly, or underuse results in needless trips for servicing. We expect that maintenance costs could be reduced and citizen satisfaction improved if garbage was removed on-demand rather than at fixed intervals.

Garbage bin monitoring is ideal for data muling because they are often sparsely deployed in remote areas, yet they serve humans carrying mobile phones [32].

Habitat Monitoring: Habitat monitoring has been studied by many sensor-net researchers [6, 34, 30]. Several deployments to-date have used long-range wireless or satellite connections to relay observations to researchers' institutions, but expensive and custom networking may be challenging to justify for smaller habitat monitoring projects. We suggest that data muling can lower the cost and technical requirements for habitat monitoring by exploiting the mobility of humans as they travel to the target habitat then back to urban areas with inexpensive networking. Even if some habitats lack regular hikers, opportunities for data muling may be sufficient with park rangers or scientists.

Car Blackboxes: Several projects have brought sensing into vehicles [4, 23, 19]. Vehicles have much service information to provide, including gas mileage, emissions, cooler temperature, and engine performance. Applications may mine these data archives to suggest needed car maintenance or to feedback to the car manufacturers. While projects such as CarTel have shown one can exploit opportunistic wifi connections [19], increased security concerns mean that availability of open wifi networks can be inconsistent, and owners may not be motivated to extend their home wifi network to parking areas. Our data muling system can replace wifi connections with Bluetooth connections to the driver's mobile phone (CarTel suggested, but did not explore, this possibility [19]).

Personal Energy Monitoring: Personal energy conservation is of growing interest, partly because simple knowledge of energy consumption allows individuals to reduce consumption by 5–10% [17]. Several projects have begun




Figure 1. An overview of our data muling system.

instrumenting home power usage, including Google Power-Meter [15] and MS Hohm [24]; others such as LoCal [22] are exploring smart grids that negotiate electricity usage between suppliers and consumers.

While these approaches provide a region- or home-centric view of energy use, we suggest that data muling can provide a *personal* model of energy use—capturing use at work and in public spaces as well as at home. A personal view of energy consumption requires harvesting data from sensors in our environment. We suggest that data muling can provide this data: imagine outfitting devices that consume energy with Bluetooth-based energy sensors, and allowing people’s mobile phones to gather this information directly from their environment. Data muling in this case is less about connectivity than about discovery and recording the correct information. While energy data may be too sensitive to post publicly on the Internet, the physical proximity provided by muling may reduce those concerns, since proximity implies some relationship with the provider of the data being collected.

Hiking Water Quality: Hikers often face questions about water quality, either at remote fountains or natural springs or streams. Governments or informed hikers may inspect water quality, but today there is little way to retrieve water data from remote locations, or to share that data with other hikers. Remote locations have little data infrastructure, and although important to health, the economic value of monitoring hiking water cannot support satellite data. As shown by CenWits [18], data muling can exploit hikers’ mobility to deliver the information of water quality cost-effectively to those who care about it. We suggest that the mobile phones likely already carried by the hikers are an ideal data mule today.

2.3 Data Muling Challenges

Prior work has shown the principles of data muling: relaying data between nodes upon rendezvous, exploiting random mobility [21] or expected mobility patterns [8, 25]. We build on this prior work and answer several new challenges:

How effective are current mobile phones as data mules? We examine mobile phone hardware, evaluating Bluetooth, wifi, and 3G cellular technologies for networking. We consider evaluate energy consumption and the unique constraints of mobile telephones, with typical daily charging and an important requirement to never run out of power.

Is data muling feasible for traditional sensornet applications? We explore four sensornet deployments retrieving

different kinds of data to understand how well data muling works in practice. Important in our study is understanding how mobile phones work in these environments, how real-world data sizes and latency requirements affect muling, and if muling is effective for these applications. We also compare loiter time, from human mobility, and muling time, driven by application requirements.

How well does human mobility support data muling? The success of muling depends on the mobility pattern. Ideally we would like muling to be “free”, leveraging human mobility. We study new and existing datasets of human mobility, and then evaluate intentional and opportunistic mobility in our deployments.

3 Design of Our Data Muling System

Our data muling system employs three components shown in Figure 1: sensors, mules, and gateways. *Sensors* generate data, store it locally, and periodically listen for a mule to come by. Humans carry *data mules* around sensors; mules periodically scan for sensors and automatically fetch any new data they see. Humans may carry the data mules with *intentional*, planned meetings with sensors, or they may simply go about their daily business with rendezvous happening *opportunisticly*. Mules and sensors communicate over a short-range, low-power wireless network. In our implementation this short-range network is Bluetooth, although we have considered 802.11 and 802.15.4 as third option.

Data mules only store data temporarily, they relay their data to Internet-based *data servers* where it is analyzed by users. Our implementation uses either 3G cellular networks or 802.11 for mule-to-server communication, based on the preference of the mule carrier and convenience of 802.11. Muling allows two important advantages relative to traditional sensor networks: first, the sensornet need not have a dedicated Internet gateway with corresponding need for extra power and cost. Bluetooth is inexpensive enough we use it on all sensors, and the mule amortizes the cost of WAN communication over all sensors. Second, when sensors originate data in an area with poor WAN coverage, a human can take a mule from this data-rich/network-poor area to another area with good or free WAN coverage. Both of these advantages motivated our adoption of data muling in our four deployments.

Next we describe each component of our data muling system in more detail.

3.1 Sensors

Details of sensor hardware depends on the specific application. The only additional requirement for muling is support of some low-power, short-range wireless networking protocol.

In our deployments our sensors use embedded PCs with Bluetooth as the short-range wireless protocol. Our choice of Bluetooth is driven primarily by what is available on the mule, but many devices today include built-in support Bluetooth, including weather stations and automobile accessories. In our case, our custom-built sensors did not include built-in Bluetooth hardware, but we easily added a Bluetooth via USB. In principle, other short-range, low-power wireless networking protocols could replace Bluetooth. One promising protocol is 802.15.4 because of its low power; we also considered 802.11 because of its higher speed, as discussed in Section 6.2. Ultimately we look for Bluetooth 3.0 to provide both low-power rendezvous and high speed data transmission when required.

Our muling scheme is agnostic to the type of sensor. We use two sensors in our deployments. Two deployments use GPS sensors that capture files that are 1–2 MB in size (after compression) every two hours, generating 12–24 MB/day. The third deployment tracks people via their carried Bluetooth devices; this sensor generates very small datasets, typically less than 800 kB per day.

3.2 Data Mules

Our data mules are mobile phones. We use mobile phones to exploit these intelligent devices carried by almost everyone today. We next describe our hardware and software choices for our mules.

We currently use four different Android-based smartphones as our mules: the HTC Hero, HTC Touch, Samsung Galaxy S, and HTC EVO. These platforms were chosen because of their availability and suitability for an individual's personal mobile phone. These platforms also drive our choice of Bluetooth as the sensor-to-mule wireless protocol, so that no hardware modifications to the mule is required. We also considered but have not yet implemented a mule based on an embedded PC with 802.15.4 support so we could mule from motes and other embedded sensors. We currently require support for Android 2.0 for Bluetooth scanning supported only since that release. In principle our muling software should port to other smartphones such as the Apple iPhone or phones based on Windows or SymbianOS.

On the mobile phone we run our custom muling service. The muling application runs as a background service, periodically scanning for neighboring Bluetooth devices to determine if they are sensors. We use a default scan interval of 2 minutes, although configured from 1 to 10 minutes to trade-off detection speed against energy consumption (we look at energy consumption in Section 6.4).

The mule has a list of known sensors, and when one is within range it connects to the sensor and retrieve and fresh data. In addition to gathering data from the sensor, the mule pushes a delivery report of what data has been delivered to the Internet, allowing the sensor to garbage collect data.

Although it is not essential for successful muling, our application includes a user interface that reports what sensors

are in range and what data has been collected on the mule. The interface can be used for other purposes well, if a sensor needs servicing (perhaps battery replacement or sensor cleaning), the mule could request assistance from the carrier.

3.3 Data Servers on the Internet

Finally, we expect that all data is ultimately hosted on servers on the Internet (as is the case in nearly all operational sensor networks). Our current implementation uses two different storage servers. We use a webserver with an off-the-shelf uploading extension to support our two data-intensive deployments. Subsidence/Oilfield and Subsidence/Urban (Section 4) each generates 84–168 MB per week. For our third application (People/ISI), we store data in Sensorbase.org [9], a sensor data sharing platform built on Apache and MySQL. Sensorbase also includes support for managing and sharing sensor data, and allowing users to query and interact with stored data.

3.4 Design Alternatives

We prototyped sensors and mules based on Mica2 motes, and considered multi-hop communication, although neither is active in our current system. We briefly discuss these design alternatives here.

Mote-based Mules: We prototyped a data muling system using Mica2 motes, and this prototype is representative of muling with an 802.15.4 network for mote-to-mule. Mica2s are very attractive sensors because of their proven success at long-term, energy conscious operation, combined with their easy customizability.

The chief barrier to 802.15.4-based sensors is with the mule, since few commodity devices include 802.15.4 network hardware. We considered having users carry a mote as the mule, but use of a second device just for muling is a significant burden. While we believe many users would run muling if it had no impact on their use of existing devices, we expect that few would add a new device to their daily lives for this purpose. As a secondary concern, standard motes have less than 1 MB flash storage, too little for our Subsidence applications. A more attractive alternative may be to add an 802.15.4 radio to a mobile phone or other already-carried device; we have not yet explored this option.

Multi-hop communication between sensors: Our current system assumes all sensors can directly communicate with the mule. With many multi-hop communication (mesh) protocols for sensor networks, we could easily employ multi-hop communication between sensors.

A mesh network between sensors is of interest only when sensors are clustered, physically close to each other. In that case, the sensors could preemptively push data to a designated collection point, or the appearance of a mule could prompt the sensors to gather their data on-the-fly. Either way, mesh communication increases opportunities to mule data. We believe the greatest advantage of multi-hop networking among sensors is that it can extend the energy-efficient network to locations that are difficult for the mule to reach. A second advantage may be coordinated sleeping among sensors to conserve energy (as explored for other purposes [36, 10, 27]). However, the cost of a sensor mesh is

much greater complexity, to insure a connected mesh and to manage resource usage at a designated collection point. We did not employ sensor-to-sensor communication because our deployments have one or two clusters of sensors, so multi-hop was not necessary at each site, where all sensors could be reached with one visit, nor possible between sites.

Multi-hop communication between mules: We also considered multi-hop communication between mules, as in prior sensornet muling deployments [21, 18] or between mobile phones during disasters [11]. Given the wide-spread availability of cellular data connectivity our goal is to mule data out of dead spots with poor network coverage but sensing interest, and to amortize the cost of the data connection among many sensors. Mule-to-mule communication does not help either of those goals. We therefore do not consider mule-to-mule connectivity.

4 Case Studies: Human Mobility and Sensor-net Deployments

We employ six datasets in this paper to understand data muling potential and practice. As shown in Table 1, the first two, Mobility/MIT and Mobility/ISI, are passive observations taken from mobile phones, while the others are four different sensornet deployments we undertook to study data muling. These cases show both opportunistic and intentional mobility patterns with different numbers of mules and sensors; we describe each case next.

4.1 Observations of Human Mobility

We began our work with a public dataset about mobile phone mobility, then conducted additional experiments to improve precision.

The *Mobility/MIT* trace is from the Reality Mining project, where they collected mobile phone activities from 100 mobile phones for 9 months [13]. (Their full dataset runs 18 months, but we use the 9 months starting in Sept. 2004 that they identify as their active data collection period.) Their public dataset includes rich information including calls, location, and Bluetooth contacts. We study their Bluetooth contact information to estimate regularity in human mobility to show the potential for data muling in Section 5.

The *Mobility/MIT* dataset scans for neighbors at 5-minute intervals. To observe brief connectivity, we carried out a smaller Bluetooth survey with mobile phones with the *Mobility/ISI* dataset. We scan every two minutes, a period chosen to balance battery life and frequency of detection. Each scan takes about 15 seconds. We use this additional dataset to update the prior dataset and to better understand opportunities for muling in Section 6.

4.2 Sensornet Deployments

We have employed our data muling system with four sensornet deployments.

The *Subsidence/Oilfield* deployment involves two pairs of GPS units observing subsidence in a production oilfield. The project carried out multiple deployments over several years and adopted data muling for the most recent 4-month deployment from Oct. 2010 to Jan. 2011 out of necessity. The experimental hardware required close monitoring to insure

correct operation. Unfortunately, the industrial field wireless network was not ubiquitous, nor were we allowed to access. Early deployments used 2G and 3G cellular modems for data, but we were unable to get consistent cellular coverage for more than a few days. The system’s primary method of data retrieval was with manually swapped flash memory cards, but month-long intervals and difficultly swapping cards (travel to site, open locked box, halt machine, etc.) made “sneakernet” untenable. We therefore turned to data muling.

Unlike the mobility datasets, muling for subsidence/oilfield is *intentional*: with only one mule and a large oilfield, field personnel would explicitly drive to each approximately weekly. Although in principle one could have swapped memory cards, data muling greatly simplified data retrieval, since it requires only wireless connection, personnel need only drive nearby, park, and push a button on the smartphone.

The *Subsidence/Urban* dataset uses the same equipment as Subsidence/Oilfield. However in this case, subsidence is part of a controlled experiment and the site is at a residence in an urban area. As a residence, we were able to frequently mule (except for travel and operator error). Because the sensors are out of Bluetooth range of the residence, most muling was again intentional, however in Section 6.1.3 we show that in many cases, normal movement was near enough to the sensors provide opportunistic muling as well.

Finally, the *People/ISI* dataset is designed to provide pure opportunistic muling in an office environment. We deployed sensors in four locations: two offices and a break room at ISI, and the home of a researcher. Each sensor tracked nearby people, as determined by scanning for Bluetooth contacts; each location had a number of visitors. Muling happened only opportunistically, as the operator carried his mobile phone as part of daily use. In Section 6.1 we use this dataset to evaluate the effectiveness of casual muling in daily life.

5 Evaluating the Potential for Data Muling

We next consider opportunities for data muling. We first look at how pervasive short-range wireless sensors could be, and how regularly humans visit them. For both of these questions we consider our two observational datasets (*Mobility/MIT* and *Mobility/ISI*) and use Bluetooth devices as a stand-in for sensors, as described below. Then in the next section we revisit these questions with our system in practice in four deployments.

5.1 How Many Potential Sensors Around Us?

Data muling presumes that short-range wireless sensors are pervasive and available for muling. Today wireless sensors clearly are *not* everywhere, although *each of the pieces exist*: sensors, wireless communication, and sensor networks. Sensors are deployed and operating in almost everywhere we go: thermostats, motion detector at door, smoke detectors, power meters and water meters, several sensors in each mobile phone, hundreds of sensors in automobiles, and cameras and pressure sensors on streets. Yet today these sensors often stand-alone or are used only in specific application “stovepipe”. Wireless communication is everywhere as well, with wifi, Bluetooth, and 3G and now 4G mobile phone data.

Dataset	Goal	Description	Mules	Mobility	Sensors	Start	Duration
Mobility/MIT [13]	observation	Bluetooth scanning log (5-minute interval)	100	opportunistic	815*	Sept. 2004	9 months
Mobility/ISI	observation	Bluetooth scanning log (2-minute interval)	3	opportunistic	226*	May. 2010	12
Subsidence/Oilfield	deployment	Oilfield subsidence monitoring	1	intentional	4	Oct. 2010	3
Subsidence/Urban-BT	deployment	Subsidence monitoring in urban area (BT)	1	mostly intentional	2	Nov. 2010	4
Subsidence/Urban-Wifi	deployment	Subsidence monitoring in urban area (802.11)	1	opportunistic	2	Jun. 2011	1
People/ISI	deployment	Person monitoring in office area	1	opportunistic	4	Feb. 2011	2

Table 1. Datasets considered in this paper: observations and sensornet deployments. (* indicates sensor stand-ins)

	Mobility/MIT	Mobility/ISI
devices encountered	25,687	12,954
mobile devices	23,814	12,699
stationary devices	815	226
unknown	1029	29
start	Sept. 2004	May. 2010
duration	9 months	12

Table 2. Devices seen in observation studies.

And there have been a number of long-term sensor network deployments as well. Yet we conjecture that there is a bootstrapping problem: there is no muling today because there are few public, wireless sensors, and yet there are few such sensors because there is no muling.

To break this deadlock we first wish to estimate characterize how wireless sensors *might* operate. Here we consider the effectiveness of short-range wireless communication to answer how many sensors might we see; in the next section we evaluate how regularly humans would see those sensors. In both cases, we use *stationary Bluetooth devices* as a stand-in for wireless sensors. We list the specific classes we consider stationary below. We target only *stationary* devices because we expect most environmental sensors to be stationary. We consider Bluetooth because it is cheap enough (current Bluetooth chipsets add only pennies to the cost of a device), and low-enough power that it is a plausible technology for a public wireless sensor. Therefore we leverage the installed base of stationary Bluetooth devices to approximate the connectivity future wireless sensors may have. (In Section 6.1 we compare these results with muling in our real sensornet deployments.)

We evaluate our two observation datasets in Table 2. The Mobility/MIT dataset dataset has 25,687 unique Bluetooth devices observed by 100 mobile phones over 9 months. In Mobility/ISI dataset, we observed 12,954 unique devices from a single phone. We attribute the count in Mobility/ISI to its collection 7 years after Mobility/MIT and increasing use of Bluetooth.

Potential sensors: Both traces show that humans naturally encounter many different devices, but most of devices are due to one-time encounters with other mobile devices. (Not surprising since nearly every mobile phone as Bluetooth, while its use outside of phones is growing but lags.) Since we expect most environmental sensors to be stationary, we wish to consider encounters with stationary devices. Each Bluetooth devices includes class information in its public announcements, in addition to its unique MAC address. The class tells what kind of Bluetooth device it is, telephone, headset, and stationary classes such as desktop computer, server computer, modem/gateway, ISDN, loud speaker, set

top box, and VCR as stationary. Table 2 identifies these stationary devices and we treat them as *sensor stand-ins*. In addition, 1058 devices omit class information. We find that some devices (both stationary and mobile) send class information occasionally; we classify these as mobile or stationary when possible, or as unknown when they never report.

We conclude that even today, humans encounter many short-range, wireless devices every day. If these devices made sensor data available we would be swimming in data.

Frequency: Muling depends on frequent encounters with sensors, and humans are only likely to be interested in sensor data they regularly encounter. We therefore next consider how often humans encounter each sensor stand-in.

Figure 2 shows how often each device was seen over the 9 months of Mobility/MIT. Figure 2(a) shows the number of visits for each sensor stand-in (running along the y-axis) on each day (the x-axis). While some sensor stand-ins are seen only a few times, many appear repeatedly. We define a *potential muling day* as a day when some mule sees a given sensor-stand-in at least once. If we assume data can tolerate up to 24 hours of latency, this metric represents muling “timely coverage”. Figure 2(b) shows the number of potential muling days for the top 100 most frequently seen sensor stand-ins. We see that 26 devices have more than 100 potential muling days; these are seen on about 35% or more of the days of the dataset. These devices are excellent candidates for opportunistic muling, without their own wide-area network infrastructure.

We repeated this study with our second observation dataset (Mobility/ISI). The results in Figure 3 are qualitatively similar to our findings from Mobility/MIT, even through the population of mobile observers was much smaller. Part of the reason for similar coverage with fewer mobile devices is much greater use of Bluetooth today.

We have shown that there are hundreds of sensor stand-ins that are frequently seen with casual movement, and some of these are seen quite frequently. We next look more carefully at the regularity of sensor encounters with opportunistic mobility.

5.2 Regularity in Human Mobility

We have shown that there are many potential sensors for which Bluetooth can provide reasonable coverage, and that some are seen many times. But how *regular* is communication? If data muling is to replace wide-area networking, we need guarantees that sensors are seen not only frequently, but regularly. We next re-examine our observation datasets to judge regularity in potential data muling.

To understand how often a specific sensor is visited, Figure 4 shows contact patterns (left) and inter-meeting times (right) for a specific sensor stand-in (sensor-stand-in 270

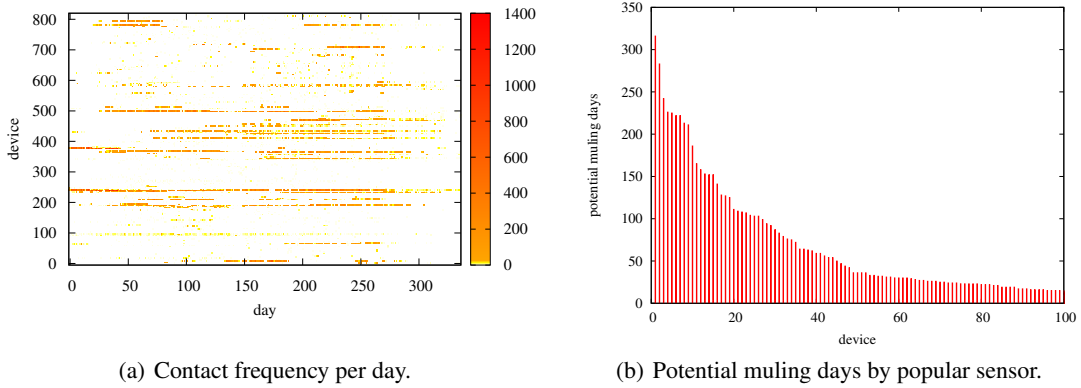


Figure 2. Contact with sensor stand-ins (dataset: Mobility/MIT).

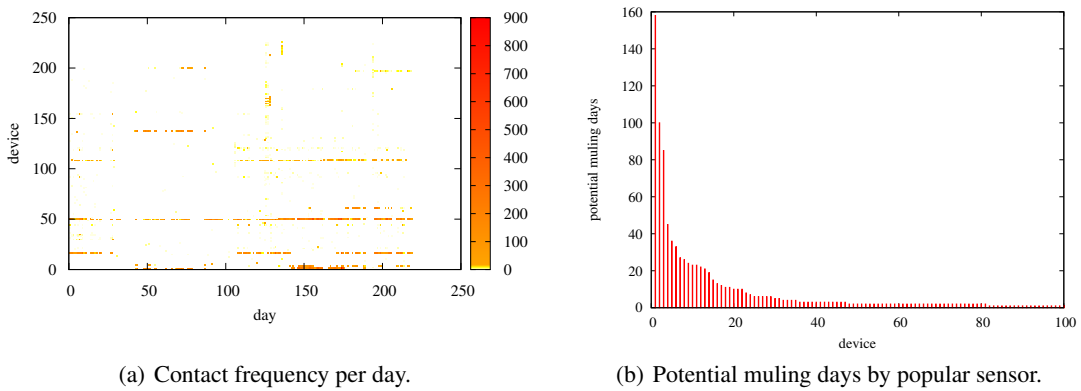


Figure 3. Contact with sensor stand-ins (dataset: Mobility/ISI).

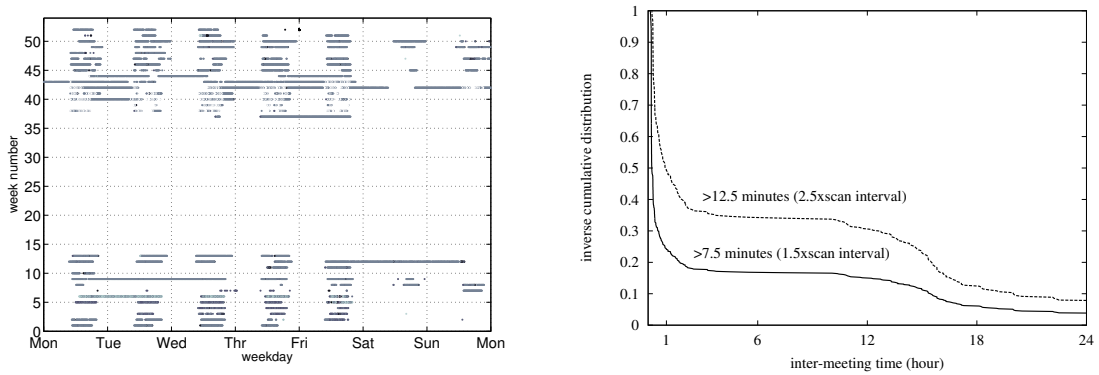


Figure 4. Contact with a specific sensor stand-in over time (dataset: Mobility/MIT, mule: 35, sensor-stand-in: 270).

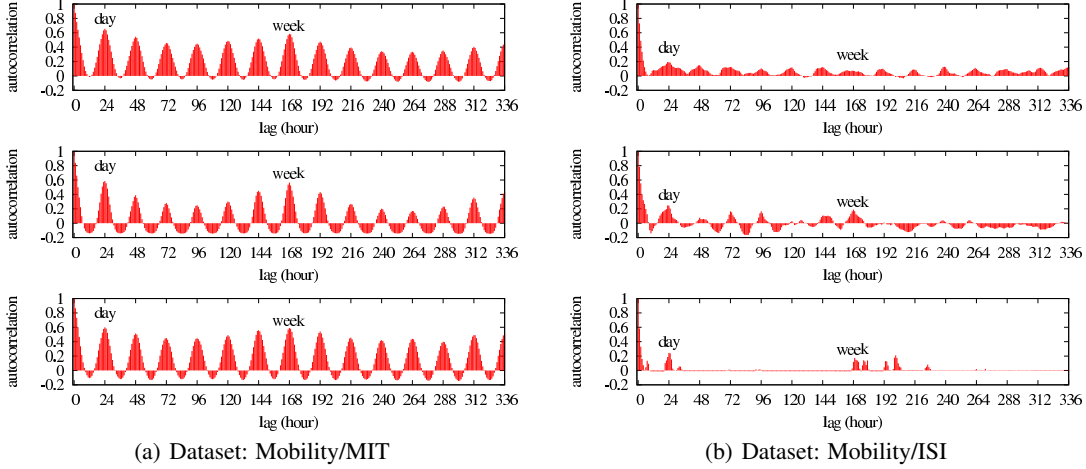


Figure 5. Autocorrelation of contacts between any mule the first, fourth, and sixth most seen sensors from two datasets.

from Mobility/MIT). We examined other sensors, this kind of contact is typical for sensors that are seen relatively frequently (this sensor was the fourth most common sensor).

Figure 4(a) shows the contact pattern over the dataset. This sensor shows strong regular daily contact on most weekdays (corresponding during a 9am to 5pm workday). It also shows occasional contact on weekends (Saturday and Sunday) and extended periods of contact (for example, at the beginning of the ninth week). We can infer that this target device is located in the person’s work area. For other sensors we see patterns of contact during non-work times (9pm through 7am). This data shows the potential for regular contact using mobile-phone-based mules.

To quantify the degree of contact, Figure 4(b) shows inter-meeting times. Since contact is determined only by scans every 5 minutes (for Mobility/MIT), we compute time between contacts by assuming contacts within some window show continuous connectivity. We use two different time-out windows, $1.5\times$ and $2.5\times$ the scan interval (7.5 or 12.5 minutes), to detect gaps, optionally bridging over a single missed scan. Both windows show similar behavior: the most frequent inter-meeting time is 5 or 10 minutes, corresponding to the window size, because it is the minimum timeout we can detect. They also a relatively sharp drop-off around 16-hours, corresponding to the non-work part of a weekday.

To better understand how often sensors are visited, Figure 5 examines autocorrelation values between for six different sensor stand-ins taken from the two observation datasets. All sensors show very strong daily periodicity (the peak at lag of 24 hours, this peak is second highest to the lag at the scan interval of 5 or 2 minutes). In addition, we often see the next highest peak at one week (168 hours), showing a strong weekly periodicity. In fact, for the sixth most seen sensor stand-in in Mobility/ISI, only immediate, daily and weekly periodicities are strong. This more complete analysis of mobility data confirms regular periodicity in human mobility and strongly suggests the potential for human-based data muling.

6 Evaluation of Our Data Muling System

Having established the potential for data muling, we next explore our data muling implementation. We begin (Section 6.1) by reviewing the potential we studied with mobile phones and sensor stand-ins in Section 5 translates into muling at our four sensornet deployments. Our deployments allow us to study loiter time in Section 6.2, to understand how data size and human movement interact. Finally, we evaluate the energy requirements of muling to understand trade-offs in network type and data size 6.4.

6.1 Does Data Muling Work in Real Deployments?

Our analysis of the observation datasets suggest the potential for regular data muling with mobile phones and opportunistic mobility. We next turn to our four deployments: Subsidence/Oilfield, Subsidence/Urban-BT, Subsidence/Urban-Wifi, and People/ISI, to evaluate how muling performs in practice.

6.1.1 Muling in the Oilfield

First we consider the Subsidence/Oilfield deployment. We turned to muling here to monitor experiment hardware, after being denied access to the industrial wireless network and having extremely poor coverage with the 3G mobile data network. Some form of regular communication with sensors was a requirement for this application because the hardware required regular evaluation and maintenance at unexpected times. Our muling system used a mobile phone carried by a field engineer who would gather data using the phone’s Bluetooth connection, then carry the phone into town at night where 3G coverage is quite good. The alternative to muling was a 6-hour round-trip drive into the field at periodic intervals; our goal instead was at least weekly updates and data-to-date.

This experimental deployment brought out three things: it required intentional movement, muling met our latency expectations, and it succeeded in reducing downtime. We consider each next.

First, this deployment required *intentional*, not opportunistic movement. Our studies of the observational datasets

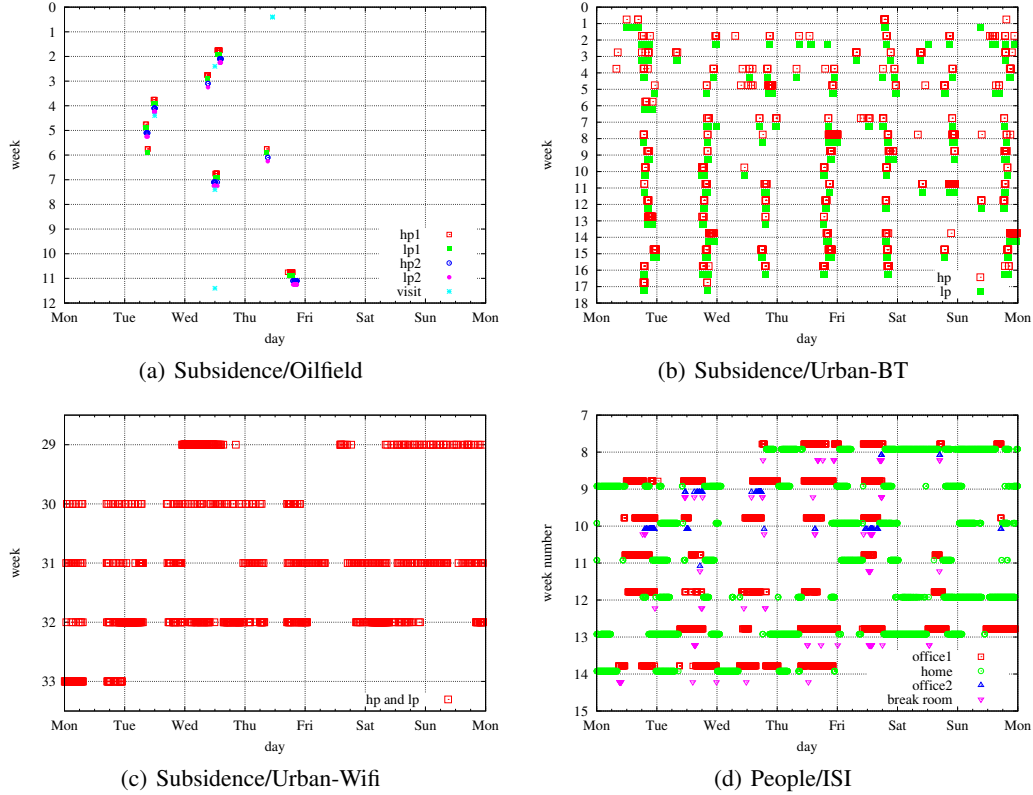


Figure 6. Weekly pattern of four muling deployments.

show that many sensor stand-ins are seen often. While true, in the Subsidence/Oilfield deployment we have only *one* mule and four *specific* sensors to visit, not 3 or 100 mules scanning for any sensors. The oilfield is a very large area (more than 50 km²), and Bluetooth radios typically operate with range limits of 10 m or less, so our single mule would nearly *never* meet our sensors accidentally. Instead, we asked the field engineer to *intentionally* travel to the sensors and wait next to them while data transfer takes place. As Figure 6(a) shows, regular weekly muling with intentional mobility was successful for the first 5 weeks of operation. However, the figure also shows a weakness in relying on human mobility: the gap in weeks 8 through 10 is due to December vacation by our mule carrier.

In a large outdoor area, muling with opportunistic mobility requires many more mules and longer-range radios. The site has many workers moving about, if we could equip each engineer’s phone with muling, or instrument company trucks with wifi-based mules, we expect we would reduce the need for intentional mobility.

Second, muling met our latency expectations, as Figure 6(a) shows, regular weekly muling with intentional mobility was successful. With intentional movement this result is expected, but it confirms that our muling system was successful at meeting its goal, and that our system was usable by a non-expert.

6.1.2 Oilfield Muling and Sensor Coverage

We have shown that muling works in the oilfield, but not that it *helps*. For us, muling helps by providing data and informing us of deployment problems more rapidly than periodic visits would.

Our alternatives were muling every week, or scheduled visits every three weeks that swap memory cards. On average, muling retrieves data 1.5 weeks before it would have been acquired with memory card swapping. (Although both muling and memory card swaps require a visit to the site, memory card swaps required much more time and a level of technical involvement that field personnel were unwilling to undertake.)

More important than getting data more rapidly, muling alerted us to problems with our deployment. As an experimental deployment, we encountered several hardware and software problems. The early detection of muling gave us an opportunity to understand when outages occur. Table 3 shows outages dates in our deployment, and when these outages were detected with muling, compared to when they would have been detected with a scheduled visit. As can be seen, muling halved the time outages were unknown, from 60 sensor-days to 27, allowing us to make an informed decision about the need for early field visits for maintenance.

6.1.3 Muling for Urban Sensing

Our second deployment is Subsidence/Urban-BT deployment. While the application is the same as Subsidence/Oilfield, in this case the site is an urban area with the

sensor	outage dates	detection date		dur. unknown	
		mule / scheduled		mule / scheduled	
hp1	Nov. 12 to Nov. 23	Nov. 17	Nov. 23	5	11
	Nov. 26 to Dec. 15	Nov. 15	Dec. 15	4	19
hp2	Nov. 12 to Nov. 23	Nov. 17	Nov. 23	5	11
	Dec. 09 to Dec. 15	Dec. 15	Dec. 15	6	6
lp1	none	—	—	0	0
lp2	Nov. 10 to Nov. 23	Nov. 17	Nov. 23	7	13
total				27	60

Table 3. Outages in Subsidence/Oilfield.

goal of collecting data for a controlled subsidence experiment.

This experiment again required intentional mobility. Although located at a residence, the sensors are behind the garage and so usually of Bluetooth range from the primary living areas. As Figure 6(b) shows, intentional mobility addresses this problem. However, in examining the data we note about 7% of the contacts appear to be opportunistic. Opportunistic mobility is therefore possible and may provide a benefit even when not planned. We expect use of wifi over Bluetooth would have allowed all-opportunistic muling in this scenario.

With *Subsidence/Urban-Wifi* deployment, we increase the radio range by replacing Bluetooth with an 802.11 radio. The longer radio range of 802.11 makes possible data muling with all-opportunistic mobility. Figure 6(c) shows that data mule covers sensors most of the time during four weeks of experiment period. Sensors are contacted with data mule opportunistically, because we do not employ intentional mobility used in the Subsidence/Urban-BT deployment. We conclude pure opportunistic mobility is sufficient for data muling in the Subsidence/Urban-Wifi deployment. In addition, we analyze how a faster data rate of 802.11 affects in muling data in Section 6.3.

Both sensors in Figure 6(c) make contact exactly same pattern. This is because sensors in the Subsidence/Urban-Wifi deployment share a wireless access point that data mule makes contact to. Whenever data mule connects to the wireless access point, both sensors are contacted by the mule at the same time. Therefore, sensors have the same contact pattern.

6.1.4 Muling for Office Sensing

Our final deployment was designed to test purely opportunistic data muling. We placed several sensors in an office environment better suited to Bluetooth’s short range. Here the sensors sense people (in our implementation, by looking for human-carried Bluetooth devices using a second Bluetooth adapter); we place sensors in four locations: two offices, each visited by a few people; a break room visited by many people. We also place one sensor at an apartment shared by several people. One individual carried the muling device, visiting the apartment and one office and the break room daily, the other office weekly.

Figure 6(d) shows muling opportunities at each sensor over the course of eight weeks. We see many opportunities to mule at office1 and home (square and circle), and regular opportunities at the break room and office2. There is actually some correlation between breakroom and office2

because they are at the edge of Bluetooth range. We conclude that opportunistic muling works very well when radio range and mobility patterns are well matched, as in an office environment.

6.2 Loiter Time Effects on Muling

For data muling to be successful, the mule must stay within radio range of the sensors long enough to transfer any pending data: the loiter time must be longer than the muling time. Muling time is a function of the size of each data item and the number queued up to send, which in turn depends on contact frequency. We next evaluate muling time and estimate loiter times to see how often successful muling is likely to occur.

Observing: We first estimate required muling time as a function of data size and number of queued data items in Figure 7(a). In this graph, each diagonal line represents a single data size, from 1 byte to 1 MB, and each point on that line a different number of data items, from 1 to 100 (for small sizes), or to 20 or 2 for the largest sizes. (Note that the observations in our subsidence applications are 1–2 MB in size, the largest size listed.) Each point is taken experimentally and is the mean of 10 measurements, with error bars showing standard deviation. We include all delays in muling: There is roughly fixed-duration overhead for a mule to discover sensors, and setup a Bluetooth connection, and determine if there is data to mule; together this setup requires about 17 seconds. Then the time to transfer data items is roughly linear with the quantity of data transferred.

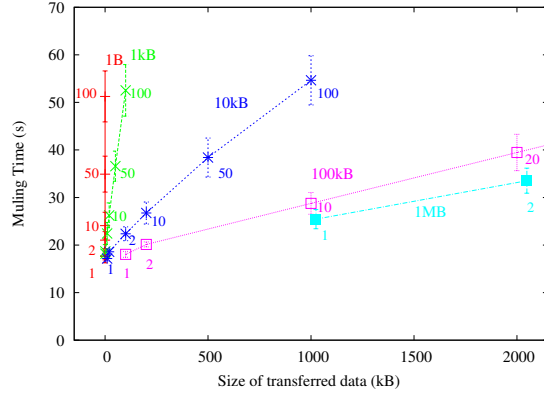
Modeling: To understand muling time, we fit a simple linear model to our observations. Muling time consists of three components: communication overhead (discovery and connection), data transfer time, and muling overhead. Communication overhead is almost constant; we measure 13 s for Bluetooth discovery and 1 s to open a connection, consistent with what is stated in the Bluetooth specification [16]. Transfer time changes according to the total size of data, and the muling overhead increases according to the number of files. If we define k_{xfr} is the transfer rate, k_F as the per-file overhead, and k_D as monitoring and disconnect overhead, we can solve for these constants using multiple linear regression.

$$\begin{aligned}
 T_{muling} &= t_{discovery} + t_{conn} + t_{transfer} + t_{muling_overhead} \\
 t_{transfer} + t_{muling_overhead} &= (1/k_{xfr})SN + k_F N + k_D \\
 &\text{where } S \text{ is the data item size} \\
 &\text{and } N \text{ is the number of files}
 \end{aligned}
 \tag{1}$$


Predicting: This analytic model helps us evaluate how muling can work with different kinds of mobility patterns.

Table 4. Parameters for the data muling time model (802.11 model is discussed in Section 6.3).

variable	Bluetooth	802.11
$t_{discovery}$	$\mu = 13.01 \text{ s}, \sigma = 0.448 \text{ s}$	$\mu = 3.70 \text{ s}, \sigma = 0.192 \text{ s}$
t_{conn}	$\mu = 1.616 \text{ s}, \sigma = 0.661 \text{ s}$	$\mu = 0.277 \text{ s}, \sigma = 0.179 \text{ s}$
k_{xfr}	1.143 Mb/s	14.435 Mb/s
k_F	0.3280 s/file	0.3642 s/file
k_D	3.402 s	1.670 s



(a) The transfer time between mule and sensor with different data sizes



(b) Loiter time of mules with 10 most frequently contacted nodes

Figure 7. Data muling time and loiter time (dataset: Mobility/ISI)

The large rendezvous time is a critical factor to muling: with Bluetooth, muling *any* amount of data takes at least 20 s. This limit places a bound on user movement: with a 10 m radio range, a user can move at most 1 m/s if a scan begins immediately on entering radio range. Typical human walking pace depends on age, but ranges from 1.35 to 1.5 m/s for an already walking individual, depending on age [5]. Therefore Bluetooth-based muling will *not work well* for constantly moving humans, even if mules constantly scan for sensors. This result suggests that muling must involve either accidental loitering, longer range radios (perhaps higher power Bluetooth with ranges to 100 m, or 802.11 with ranges of 35–70 m), or great improvements to device discovery protocols. This result is consistent with our observations in People/ISI, since we place each sensor at a location where the data mule is likely to loiter—offices, a break room, or home.

The model also shows that transfer time is irrelevant for almost all small data sizes (data items than 100 kB), since connect time dominates. However, with the large, 1–2 MB data items in our Subsidence applications, transfer time is very noticeable and transfer time dominates muling time when more than a few of items are queued for transfer.

Finally, we can compare this model to our observation datasets to evaluate how often muling would likely succeed. This comparison is difficult, because those datasets record only sensor-stand-in *detection* times, not contact times, so we know when a device was seen but not for how long. However, if we assume multiple consecutive detections correspond to continuous contact, then we can infer contact times in those cases. If consider the Mobility/ISI dataset since it has greater scan frequency than Mobility/MIT (2 rather than 5 minutes between scans), we can then classify more than two scans (loiter time more than 4 minutes) as enough time to mule almost anything, two scans (2 to 4 minutes) as enough time to mule all cases on Figure 7(a), and one scan (up to 2 minutes) as possibly enough to mule something

Figure 7(b) shows loiter times in these three categories for the ten most frequently contacted sensor stand-ins from Mobility/ISI. For these sensors, about 20%–60% of contacts are 2 or more than 2 scans; long enough to transfer at least

15 MB of data. We cannot judge loiter times for single scans, but there still seems some chance to transfer smaller data items (1–100 kB) in 17–20 s.

Large, multi-item transfers: Finally, our focus here has been understanding how brief loiter times interact with minimal, opportunistic data transfer. In our Subsidence applications, data items are each 1–2 MB in size, and muling is done once a day (Subsidence/Urban) or once a week (Subsidence/Oilfield). Each sensor generates an observation every 2 hours, so these applications require moving 12–24 MB or 84–168 MB of data per muling session. Our model predicts muling times of 2–20 minutes, and our experiences bear this out: muling takes a long time for large-data sensing with infrequent rendezvous.

The best way to reduce muling time for large data items is to increase transfer speed. We considered 802.11 as a muling network, and its 54 Mb/s transfer rate would slash a 2 MB transfer from 16 s to 0.3 s. Bluetooth 3.0 has recognized this need, and it will employ current Bluetooth for a low-power signaling and discovery protocol, then shift to wifi for bulk data transfer.

6.3 Data muling with a faster data communication: 802.11

In the previous section, we evaluated the muling time with Bluetooth. Although Bluetooth was successful at supporting these experiments, the data transfer rate was slow, with some scenarios requiring loiter times of 70–80 minutes. The short radio range of Bluetooth is also problematic in the Subsidence/Urban-BT deployment, since it requires intentional mobility of data mules to make contact with sensors.

Here we propose 802.11 as an alternative to Bluetooth for data muling communication. We re-evaluate 802.11-based data muling in the Subsidence/Urban-Wifi deployment to see how much it improves the muling data rate and muling opportunity. We see that the use of 802.11 permits shorter loiter times and provides longer communications range, allowing Subsidence/Urban-Wifi to work with only opportunistic muling; intentional mobility is no longer required.

First, 802.11 has much faster data rate and shorter discovery time than Bluetooth. These improvements significantly




Figure 8. Data muling time via Wifi: The transfer time between mule and sensor with different data sizes

reduce the required loiter time to get data from sensor. We conduct the same analysis with 802.11 (shown in Figure 8) and fit it to our analytic model of muling time (shown in Equation 1) to evaluate the muling performance improved by the faster data rate and shorter discovery time.

Table 4 shows a comparison of muling time of Bluetooth and 802.11. With 802.11, discovery time is reduced to 3.7 s, less than one-third of the time with Bluetooth discovery (13.0 s). It also takes only 0.3 s to connect to a sensor via 802.11, where Bluetooth takes 1.6 s. These short discovery and connection times allow data mules not only scan neighboring sensors rapidly, but also check the availability of new data quickly.

The transfer rate with 802.11 is 12 times faster than Bluetooth (shown in Table 4). This faster data rate enables data muling that were previously impossible when mules use Bluetooth. For example, a Bluetooth mule takes 23 s to mule 1MB of data. Hence a successful Bluetooth-based muling requires the carrier loiter around the sensor more than 23 s. However more than half of contacts last about 13–120 s, shown in Figure 7(b). In those short contacts, Bluetooth discovers sensors successfully (longer than 13 s), but may not be able to complete the data transfer (when loiter time is less than 23 s). However, a 802.11 data mule takes only 6 s to rendezvous and transfer a 1MB. Thus, data muling with 802.11 is quick enough to complete the data transferring process even in a very short period of loiter time.

The improvement can be more substantial when muling large and multiple data items. In the Subsidence/Oilfield deployment, it takes 20 minutes to transfer a week’s worth of data from a sensor with Bluetooth (84–168MB). A field engineer who carries a data mule with Bluetooth will have to spend 70–80 minutes to mule data from all four sensors. With 802.11 muling, data transferring time will take 1–2 minutes per sensor and less than 6 minutes for all, so the field engineer no longer need to wait hours to mule large and multiple data items. Instead just a few minutes of short contact will work. Data muling becomes more attractive with fast data transfer rate and scanning time.

As importantly, the long ranges of 802.11 (35–70 m) allow more frequent opportunistic muling. In the

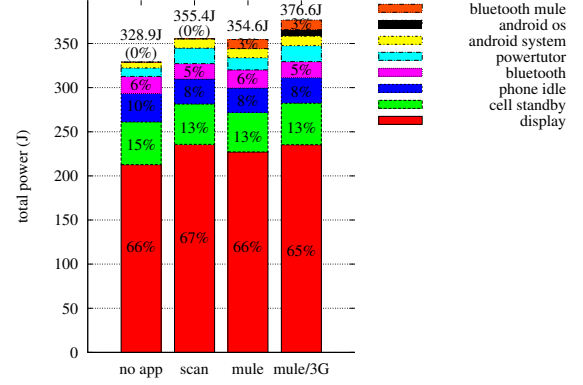


Figure 9. Battery energy consumption using Data Mule App.

Subsidence/Urban-BT deployment, sensors are located 10–15 m away from the main building where data mule loiters most of time. However Bluetooth has only 10 m radio range, so sensors are not able to contact mule unless data mule come quite close. Although we observe the data mule contacts sensors opportunistically, opportunistic contacts are only 7% of all contacts (Section 6.1.3). When we replace Bluetooth with 802.11 (in the Subsidence/Urban-Wifi deployment), the data mule contacts sensors opportunistically and gather data successfully, without intentional mobility that is required in the Subsidence/Urban-BT deployment. Thus a longer range radio increases the chance of data muling and allows opportunistic data muling to replace intentional muling.

6.4 Energy Consumption

Energy consumption of mobile devices can be critical for use. Although many mobile phones are charged every day, running out of power mid-day quickly draws user complaints. We next evaluate the energy consumption of our data muling system.

To measure energy consumption, we observe energy consumption over one hour of operation. Over that time we carry out regular scans for sensors every two minutes, and we relay 12 datafiles, each 1 MB in size. This amount of data is roughly equivalent to the amount of data sent in one daily rendezvous for the Subsidence/Urban-BT deployment. While doing muling we run the PowerTutor application [31] to record total energy consumption, and at the end of experiment we use Android’s Settings: About Phone: Battery Use application to evaluate the percentage energy consumed by each component. The data we report is the result of one set of experiments on a Samsung Galaxy S phone running Android 2.1 and BluetoothMule 2.2.3. We see similar proportions of energy use on a HTC Magic (branded as T-Mobile myTouch 3G) running Android 2.2.1. Finally, we check phone status for the first and last five minutes of the period, activating the display.

Figure 9 shows both total energy consumption (the heights of the bars, and indicated numerically above each bar), and component costs (shown as percentages for key components in each bar). We consider four scenarios: the standard phone without muling (“no app”), running muling

doing scanning with one sensor in range (“scan”), muling the amount of data (“mule”), and muling the data and sending it over the 3G network to the Internet (“mule/3G”).

Our first observation from this analysis was surprising to us: the display is by *far* the largest energy consumer. An early version of our mule intentionally left the display active to inform the user of progress; we quickly removed this energy-wasting choice.

Second, we see that frequent scanning consumes a noticeable amount of energy: about 30 J, comparing the scan vs. no-app bars. Over the course of 24 hours, scanning consumes about 5% of total battery energy.

We see that scanning takes noticeable energy. Whether or not energy consumption from scanning is a problem depends on if it runs the phone out of battery before the phone is recharged. In our use we found that scanning did not frequently exhaust phones batteries, however, on occasions when the phone was taxed for other reasons (say, long voice calls), muling contributed to forcing an early recharge.

By contrast, muling the data from the sensor does *not* consume much additional energy (compare the scan and mule bars in Figure 9). Bluetooth is optimized for energy-efficient short-range data movement, while the cost of listening many seconds to scan for potential devices is much more expensive. (This trade-off is the same one that prompted low-power listening [14, 26] and scheduling [37, 33, 38] optimizations in MAC protocols.)

Finally, we see that 3G wide-area communication approximately doubles energy consumption, consuming another 22 J (compare mule/3G vs. mule). In fact, this experiment was conducted in an urban area, and the energy costs of 3G can be greater still in areas where cellular coverage is poorer.

We conclude that energy costs of muling are noticeable, however they are relatively small in absolute terms compared to the primary function of most mobile phones: taking calls and communicating information, as reflected in the cellular standby and display energy costs. Energy consumption of muling must be considered, but in our experience it is usually not a primary factor in phone usability. In addition, when muling is the primary use for the phone, as it was in our Subsidence/Oilfield deployment, energy consumption is well within the capacity of today’s typical phones.

Based on our experience running muling we implemented several features to manage power usage. First, our muling program disables relay to the Internet when operating on battery is less than 20% (however, muling from sensor-to-phone is still done). This addition implement the policy of “phone first, Internet second”, on the assumption that the phone will likely be connected to grid power shortly and can complete data relay to the Internet at that time.

7 Related Work

Our work builds on prior work in data sharing applications, data muling in sensor networks, and understanding human mobility.

7.1 Data Muling in Sensor Networks

The concept of using data mules to support sparse sensor networks is an old one [21, 8, 28, 25, 3, 29, 20, 35].

The key idea is that a mobile mule can provide energy efficient data relay with a short range radio, while mobility can bridge long distances. Muling schemes can be categorized by the type of mobility they expect: *random*, such as with animals [21], humans [18], or simulated [28]; *controlled*, with robots [29], airplanes [35], or boat; an *predictable* mobility with trains, shuttles, or buses [8, 25], or semi-predictable farmworkers [3].

Different expectations about mobility usually are reflected by different research methodologies. Some prior studies of humans as data mules have modeled primarily random walk or random waypoint mobility patterns [28], yet we know human mobility is far from random. Other work has considered semi-random models [21, 29, 35]. We instead study human mobility using traces from mobile phones, sampling the mobility of real humans. We also and evaluate the feasibility data muling with real human movement in our four mule-based sensornet deployments.

Closest to our work, Burrell et al. [3] and CenWits [18] recognize the potential of human mobility in data muling from sensors. We compares these work with ours next.

Burrell et al. study use of sensor networks in vineyards [3]. Based on ethnographic studies, they identify farmworker mobility as capable of supporting data muling, similar to our identification of field engineer movement. Their work focuses on motes dedicated for muling, so they do not explore mobile phone traces, and their field system required sensor densities that eliminated the need for data muling [1].

CenWits is a search-and-rescue system for hikers using hiker-carried motes [18]. They recognized the need for communication in sparse areas, but as with Burrell, they propose a dedicated system.

Several recent applications have explored participatory sensing using mobile-phones [2]. These applications applications often focus on the mobile phone as the sensor, while the applications we identify in Section 2.2 assume mobile phones relay data from an in-situ sensor. Garbage-Watch is one proposed application: participants take photos of garbage bins on a campus to improve recycling [32]. We instead focus on data muling to permit reports on remote garbage bins.

7.2 Understanding Human Mobility

Many groups have studied human mobility for data communications, we summarize three very relevant studies here. As described above, Burrell et al. explored farmworker mobility with ethnographic studies [3].

Chaintreau et al. study transfer opportunities using wireless devices carried by humans as we do [7]. They find heavy-tailed inter-meeting times, and so recommend new opportunistic forwarding algorithms between mobile nodes. We discuss opportunistic data forwarding as possible future work for our system (Section 3.4), and instead focus on the data transfer between mobile phones and stationary nodes with a real working system.

Eagle et al. study human mobility patterns using mobile phones as proxies [13]. Their goal is to understand human and group relationships. They find that individual behavior over a specific day can be approximated by a weighted sum

of repeated behavioral patterns they call eigenbehaviors, corresponding to behaviors such as a normal day vs. traveling, weekends vs. weekday, etc. We build on their studies of human mobility to evaluate the feasibility of data muling in daily life. While they focus on extracting and analyzing the underlying structure of human behaviors, we exploit the routine of human mobility to cover sparsely deployed sensors efficiently.

8 Future Work

Although our system has been operational for some time, there remain several areas of future work.

As proposed in CarTel [19], a system should make the most of short periods of opportunistic muling. We currently implement a simple policy of relaying newest-data first, but summaries or sampling may be better policies for specific applications.

We have several ideas to make muling rendezvous more efficient. First, there are opportunities to improve protocol-level rendezvous, using techniques such as low-power listening [14, 26] or scheduling [37, 33, 38]. Also, we believe that study of prior mobility patterns can improve predictions about future rendezvous, allowing us to dynamically alter protocol-level behavior when conditions for rendezvous are favorable. In addition to speeding rendezvous, mules that visit sensors very frequently waste energy confirming there is nothing new to relay. We therefore may explore less frequent scanning for often-visited sensors.

Finally, our system focuses on human mobility, so human motivation plays a role. While in some cases (like Subsidence/Oilfield), muling may be company policy, in more general cases we need to consider reasons for users participate in muling. More importantly, our current system makes no attempt to influence human movement. We would like to explore giving users incentives to approach sensors to assist muling.

9 Conclusion

We have shown that data muling with human-carried mobile phones is both possible and practical. We have demonstrated the potential with analysis of two datasets of mobile phone movement, showing that individuals see many potential sensors, and see some regularly. Inspired by this potential, we implemented a data muling system and used it to share data in four deployed sensor networks. We showed that short radio ranges of Bluetooth require intentional mobility to make muling practical for industrial and even some urban applications, but that opportunistic muling is suitable for our office-based deployment. We investigated trade-offs in data size, visitation frequency, and how they interact with muling and loiter times, and we examined energy consumption. While work remains, we believe data muling has a role in bringing communication to sparsely connected sensors.

Acknowledgments

This work would not have been possible without the strong support of several Chevron personnel in cooperation with CiSoft. We thank Greg LaFramboise for project guidance and Tai Nguyen for support at Lost Hills. We also thank John Houghton and Roy Kutlik for involvement with Subsidence/Oilfield, and Iraj Ershagi (USC) and Mike Hauser

for their guidance as CiSoft directors. Finally, at USC/ISI, Fabio Silva lead Subsidence/Oilfield development, and Zhihan (Sam) Lu and Hsing-hua Chen did early development and experiments.

10 References

- [1] Richard Beckwith. Personal communication, April 2011.
- [2] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava. Participatory sensing. In *Proceedings of the 1st ACM World Sensor Web Workshop*, pages 117–134, Boulder, Colorado, USA, November 2006. ACM.
- [3] Jenna Burrell, Tim Brooke, and Richard Beckwith. Vineyard computing: Sensor networks in agricultural production. *IEEE Pervasive Computing Magazine*, 3(1):38–45, January 2004.
- [4] Vladimir Bychkovsky, Bret Hull, Allen K. Miu, Hari Balakrishnan, and Samuel Madden. A measurement study of vehicular internet access using in situ wi-fi networks. In *12th ACM MOBICOM Conf.*, Los Angeles, CA, September 2006.
- [5] Nick Carey. Establishing pedestrian walking speeds (draft results). Technical report, Portland State University, 2005. http://www.westernite.org/datacollectionfund/2005/psu_ped_summary.pdf.
- [6] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San Jose, Costa Rica, April 3-5 2001.
- [7] Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proceedings of the IEEE Infocom*, 2006.
- [8] Arnab Chakrabarti, Ashutosh Sabharwal, and Behnaam Aazhang. Using predictable observer mobility for power efficient design of sensor networks. In *Proceedings of the The second International Workshop on Information Processing in Sensor Networks (IPSN)*, pages 129–145, 2003.
- [9] Kevin Chang, Nathan Yau, Mark Hansen, and Deborah Estrin. Sensorbase.org - a centralized repository to slog sensor network data, May 2006.
- [10] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, pages 85–96, Rome, Italy, July 2001. ACM.
- [11] Xiangchuan Chen. Enabling disconnected transitive communication in mobile ad hoc networks. In *Proceed-*

ings of the Workshop on Principles of Mobile Computing, colocated with PODC01, pages 21–27, 2001.

- [12] Barry Collins. The true cost of publishing on the Amazon Kindle. <http://www.pcpro.co.uk/blogs/2011/02/15/the-true-cost-of-publishing-on-the-amazon-kindle/>.
- [13] Nathan Eagle and Alex Pentland. Eigenbehaviors: Identifying structure in routine. In *Behavioral Ecology and Sociobiology*, pages 1057–1066, 2009.
- [14] A. El-Hoiydi, J.-D. Decotignie, and J. Hernandez. Low power MAC protocols for infrastructure wireless sensor networks. In *Proceedings of the Fifth European Wireless Conference*, pages 563–569, Barcelona, Spain, February 2004. February.
- [15] Google. Google PowerMeter. <http://www.google.org/powermeter/>, 2007.
- [16] Bluetooth Special Interest Group. Bluetooth specification version 1.1 and 1.2, 2001.
- [17] D.J. Hammerstrom. Olympic peninsula gridwise study (2007), 2007.
- [18] Jyh how Huang. Cenwits: A sensor-based loosely coupled search and rescue system using witnesses. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 180–191. ACM Press, 2005.
- [19] Bret Hull, Vladimir Bychkovsky, Yang Zhang, Kevin Chen, Michel Goraczko, Allen Miu, Eugene Shih, and Samuel Madden Hari Balakrishnan. CarTel: A distributed mobile sensor computing system. In *Proceedings of the Fourth ACM SenSys Conference*, pages 125–138, Boulder, Colorado, USA, November 2006. ACM.
- [20] Sushant Jain, Rahul C. Shah, Waylon Brunette, Gaetano Borriello, and Sumit Roy. Exploiting mobility for energy efficient data collection in wireless sensor networks. *Mob. Netw. Appl.*, 11(3):327–339, 2006.
- [21] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li shiuan Peh, and Daniel Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet. In *Proceedings of the 10th international conference on Architectural support for programming languages and operating systems*, ASPLOS-X, pages 96–107, New York, NY, USA, 2002. ACM.
- [22] Randy Katz, David Culler, Seth Sanders, Sara Alspaugh, Yanpei Chen, Stephen Dawson-Haggerty, Prabal Dutta, Mike He, Xiaofan Jiang, Laura Keys, Andrew Krioukov, Ken Lutz, Jorge Ortiz, Prashanth Mohan, Evan Reutzler, Jay Taneja, Jeff Hsu, and Sushant Shankar. An information-centric energy infrastructure: The berkeley view. *Sustainable Computing: Informatics and Systems*, 2011.
- [23] J. Lebrun, Chen-Nee Chuah, D. Ghosal, and Michael Zhang. Knowledge-based opportunistic forwarding in vehicular wireless ad hoc networks. In *61st IEEE Vehicular Technology Conference*, volume 4, pages 2289–2293, 2005.
- [24] Microsoft. MS Hohm. <http://www.microsoft-hohm.com/>, 2009.
- [25] Alex (Sandy) Pentland, Richard Fletcher, and Amir Hasson. DakNet: Rethinking connectivity in developing nations. *IEEE Computer*, 37(1):78–83, January 2004.
- [26] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the Second ACM SenSys Conference*, pages 95–107, Baltimore, MD, USA, November 2004. ACM.
- [27] Curt Schurgers, Vlasios Tsiatsis, Saurabh Ganeriwal, and Mani Srivastava. Optimizing sensor networks in the energy-latency-density design space. *IEEE Transactions on Mobile Computing*, 1(1):70–80, January 2002.
- [28] Rahul C. Shah and Sumit Roy. Data mules: Modeling a three-tier architecture for sparse sensor networks. In *Proceedings of the IEEE SNPA Workshop*, pages 30–41, 2003.
- [29] Arun A. Somasundara, Aman Kansal, David D. Jea, Deborah Estrin, and Mani B. Srivastava. Controllably mobile infrastructure for low energy embedded networks. *IEEE Transactions on Mobile Computing*, 5(8):958–973, 2006.
- [30] Robert Szwedczyk, Alan Mainwaring, Joseph Polastre, John Anderson, and David Culler. An analysis of a large scale habitat monitoring application. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 214–226, New York, NY, USA, 2004. ACM.
- [31] Lide Zhang and Birjodh Tiwana, Zhiyun Qian, Zhaoguang Wang, Robert P. Dick, Zhuoqing Morley Mao, and Lei Yang. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *Proceedings of the eighth IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis*, CODES/ISSS '10, pages 105–114, New York, NY, USA, 2010. ACM.
- [32] UCLA CENS. Garbagewatch. <http://urban.cens.ucla.edu/projects/gabagewatch/>.
- [33] Tijs van Dam and Koen Langendoen. An adaptive energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the First ACM SenSys Conference*, pages 171–180, Los Angeles, California, USA, November 2003. ACM.
- [34] Hanbiao Wang, Deborah Estrin, and Lewis Girod. Pre-processing in a tiered sensor network for habitat monitoring. In *EURASIP JASP special issue of sensor networks*, pages 392–401, 2002.
- [35] Yu Wang and Chang hua Wu. Robot-assisted sen-

sensor network deployment and data collection. In *Proceedings of the Computational Intelligence in Robotics*, pages 467–472, 2007.

- [36] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 2001. ACM.
- [37] Wei Ye, John Heidemann, and Deborah Estrin. Medium access control with coordinated, adaptive sleeping for wireless sensor networks. *ACM/IEEE Transactions on Networking*, 12(3):493–506, June 2004. A preprint of this paper was available as ISI-TR-2003-567.
- [38] Wei Ye, Fabio Silva, and John Heidemann. Ultra-low duty cycle MAC with scheduled channel polling. In *Proceedings of the Fourth ACM SenSys Conference*, pages 321–333, Boulder, Colorado, USA, November 2006. ACM.