

BikeNet: A Mobile Sensing System for Cyclist Experience Mapping

SHANE B. EISENMAN

Columbia University

EMILIANO MILUZZO, NICHOLAS D. LANE, and RONALD A. PETERSON

Dartmouth College

GAHNG-SEOP AHN

Columbia University

and

ANDREW T. CAMPBELL

Dartmouth College

We present BikeNet, a mobile sensing system for mapping the cyclist experience. Built leveraging the MetroSense architecture to provide insight into the real-world challenges of people-centric sensing, BikeNet uses a number of sensors embedded into a cyclist's bicycle to gather quantitative data about the cyclist's rides. BikeNet uses a dual-mode operation for data collection, using opportunistically encountered wireless access points in a delay-tolerant fashion by default, and leveraging the cellular data channel of the cyclist's mobile phone for real-time communication as required. BikeNet also provides a Web-based portal for each cyclist to access various representations of her data, and to allow for the sharing of cycling-related data (for example, favorite cycling routes) within cycling interest groups, and data of more general interest (for example, pollution data) with the broader community. We present: a description and prototype implementation of the

This article extends work published in the *Proceedings of the ACM Conference on Embedded Networked Sensor Systems (SenSys) 2007*.

This work is supported in part by Intel Corp., Nokia, NSF NCS-0631289, ARO W911NF-04-1-031, and the Institute for Security Technology Studies (ISTS) at Dartmouth College. ISTS support for this research is provided under award 60NANB6D6130 from the U.S. Department of Commerce and award 2005-DD-BX-1091 from the U.S. Bureau of Justice Assistance. The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of the National Institute of Standards and Technology (NIST), the U.S. Department of Commerce, or the U.S. Department of Justice.

Authors' addresses: S. B. Eisenman, Electrical Engineering Department, Columbia University, New York, NY; email: shane@ee.columbia.edu; E. Miluzzo, N. D. Lane, R. A. Peterson, Department of Computer Science, Dartmouth College, Hanover, NH 03755; G.-S. Ahn, Columbia University, New York, NY; A. T. Campbell, Department of Computer Science, Dartmouth College, Hanover, NH 03755.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2009 ACM 1550-4859/2009/12-ART6 \$10.00

DOI 10.1145/1653760.1653766 <http://doi.acm.org/10.1145/1653760.1653766>

ACM Transactions on Sensor Networks, Vol. 6, No. 1, Article 6, Publication date: December 2009.

system architecture based on customized Moteiv Tmote Invent motes and sensor-enabled Nokia N80 mobile phones; an evaluation of sensing and inference that quantifies cyclist performance and the cyclist environment; a report on networking performance in an environment characterized by bicycle mobility and human unpredictability; and a description of BikeNet system user interfaces.

Categories and Subject Descriptors: C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communications*; J.3 [**Computer Applications**]: Life and Medical Sciences—*Health*

General Terms: Design, Experimentation, Performance

Additional Key Words and Phrases: Applications, bicycling, recreation, systems

ACM Reference Format:

Eisenman, S. B., Miluzzo, E., Lane, N. D., Peterson, R. A., Ahn, G.-S., and Campbell, A. T. 2009. BikeNet: A mobile sensing system for cyclist experience mapping. *ACM Trans. Sensor Netw.* 6, 1, Article 6 (December 2009), 39 pages.

DOI = 10.1145/1653760.1653766 <http://doi.acm.org/10.1145/1653760.1653766>

1. INTRODUCTION

As evidenced by the wealth of bicycle-related sensors available at stores ranging from WalMart to specialty bike shops, there is substantial interest in the mainstream recreational cycling community in collecting data quantifying various aspects of the cycling experience. This mirrors a broader societal interest in diet [Reddy et al. 2007], health [Wood et al. 2008], and fitness metrics [Nike+iPod Sports Kit 2007]. Existing commercial bike-sensing systems targeting this demographic measure and display simple data such as wheel speed, and provide simple inferences such as distance traveled and calories burned. These systems have become increasingly more sophisticated and miniaturized. This trend is continuing and bicycles in the future will be sold with embedded fitness/performance-related sensing systems. However, a cyclist does not exercise in isolation and by ignoring important context items the data collected by current systems fails to capture a comprehensive picture of the cyclist experience. For example, limited by the imperfect human senses and often focused on the activity of cycling itself, without BikeNet the cyclist gains little empirical knowledge of important factors such as exposure to air and noise pollution, and danger due to car density.

Among recreational cyclists there is a spread in the level of interest about various characteristics of a ride. Some are competitive with their friends for the sake of bragging rights, and may want to initiate challenges to set up virtual competitions among geographically separated cyclists; some focus on health-related aspects such as personal fitness; many view bicycling as a time to relax while getting some moderate exercise and are most interested in finding routes that are safe and quiet; others want to simply archive statistics about their rides for later analysis [Beating Heart Blog 2007].

In this article, we design and implement the prototype of a system not only to give context to the cyclist performance as part of a user-targeted application (e.g., health and safety), but also to collect environmental data as part of communal projects (e.g., pollution monitoring/mapping). We quantify aspects of

cycling performance and environmental conditions that the mainstream recreational cyclist can appreciate and afford, akin to the Nike+iPod kit, a system [Nike+iPod Sports Kit 2007] for recreational runners that logs exercise history. Also, BikeNet facilitates the social sharing of cycling-related data (e.g., safer and quieter routes, or exciting and challenging routes) among cyclists seeking similar cycling experiences.

BikeNet contributes the first working mobile networked sensing system for bicyclists. Characteristics of our system include:

- Cyclist Performance/Fitness Measurement.* Like off-the-shelf solutions, BikeNet collects and stores data about the following baseline cycling performance metrics: current speed, average speed, distance traveled, calories burned. In addition, we collect and store the following advanced metrics: path incline, heart rate, galvanic skin response (a simple indicator of emotional excitement or stress level). All data sensed by the system is at least stamped with time and location metadata.
- Environment/Experience Mapping.* The system provides quantitative guidance to cyclists about the healthiness of a given route in terms of pollution levels, allergen levels, noise levels, and roughness of the terrain. These measurements, together with data from cyclist performance measurements, are correlated to create a holistic picture of the cycling experience.
- Long-Term Performance Trend Analysis.* Collected data persists beyond the ride on which it is collected. The system enables the upload of data traces into a personal repository allowing the cyclist to monitor her performance improvement or her exposure to health risks like automobile exhaust. This data that can be selectively shared with other individuals (e.g., cycling friends, physicians), or exported into a public database.
- Data Collection and Local Presentation.* BikeNet allows the cyclist to customize, via a profile of preferences, what data is collected by the system, when it is collected, where it is collected, and under what correlated conditions sensor data capture occurs (e.g., increase the sampling rate of the heart rate when the path incline is above a threshold). The profile also indicates how data is to be presented, both locally (e.g., on a handle bar-mounted cell phone LCD) on the bicycle when *en route* and through access and presentation methods once the data has been delivered to the backend repository.
- Data Query and Remote Presentation.* The system provides a Web-based portal [BikeView - the BikeNet Web Portal 2007] on the backend as a means to inject real-time queries into the system to request particular bicycling-related data of interest to the backend system user. Also, the portal can be used as a place to publish/share data with friends/competitors in the same interest group (similar to the Nike+iPod group challenge feature [Nike+iPod Sports Kit 2007]) about themselves and the paths they traverse for real-time or delayed display. In so doing, we provide a useful tool to network members of the cycling community through data of mutual interest.
- Disconnected Operation.* BikeNet utilizes an opportunistic networking paradigm, whereby mobile sensing platforms are tasked and data is muled

or uploaded according to the opportunities that arise as a result of the uncontrolled mobility of the cyclists. The BikeNet system operates in a *delay-tolerant sensing* mode by default, where cyclists go on trips, collect sensed data, and upload their data when they return to home, possibly using the assistance of data mules as discussed in Section 2.2.3. This default mode is akin to the Nike+iPod [Nike+iPod Sports Kit 2007] sensing system for runners. In this case, our BikeNet implementation uses a purely delay-tolerant mote-based solution. However, if the cyclist carries a cell phone, such as the Nokia N80 in our implementation, BikeNet automatically integrates the cell phone into the system as a mobile sensor gateway and offers real-time interaction between the backend and the cyclist in support of *real-time sensing*.

Our BikeNet implementation uses sensors embedded or interfaced with the Moteiv Tmote Invent [Moteiv Tmote Invent 2007] platform as well as camera and microphone sensors available on the Nokia N80 [Nokia N80 Mobile Device 2007] platform. All Tmote Invent and N80 platforms are networked over a common IEEE 802.15.4 short-range radio channel. The N80 additionally possesses IEEE 802.11g, Bluetooth and GSM cellular radio interfaces. With these additional radios the N80 can act as a mobile sensor gateway to support both real-time data uploading to the backend and real-time queries from the backend, over the wide area cellular network while the cyclist is en route.

In the following sections, we describe our experiences deploying a sensing system for cyclist experience mapping, leveraging opportunistic sensor networking principles and techniques [Campbell et al. 2006]. We discuss the system architecture, design, and implementation in Section 2. Section 3 describes our cyclist experience mapping application, including sensing accuracy and inference techniques, communication protocol performance, and feasibility results. Related work is discussed in Section 4 before concluding with a summary and a discussion of possible extensions in Section 5. The article presents an extended version of work [Eisenman et al. 2007] first presented at ACM SenSys 2007.

2. SYSTEM ARCHITECTURE AND DESIGN

BikeNet is a network characterized by mobile sensing and sparse radio network connectivity. Given these characteristics, and the application requirements for the system, we design the BikeNet system as an instantiation of the MetroSense architecture [Campbell et al. 2006]. The architecture offers a people-centric paradigm for large-scale sensing at the edge of the Internet using an opportunistic sensor networking approach. This approach leverages mobility-enabled interactions and provides coordination between people-centric mobile sensors and static sensors, as in the CitySense project [CitySense 2007], and includes edge wireless access nodes (i.e., SAPs) in support of sensing, tasking, and data collection. Figure 1 shows a pictorial overview of the BikeNet system. Details of a prototype implementation are included to make the architecture and design descriptions more concrete.

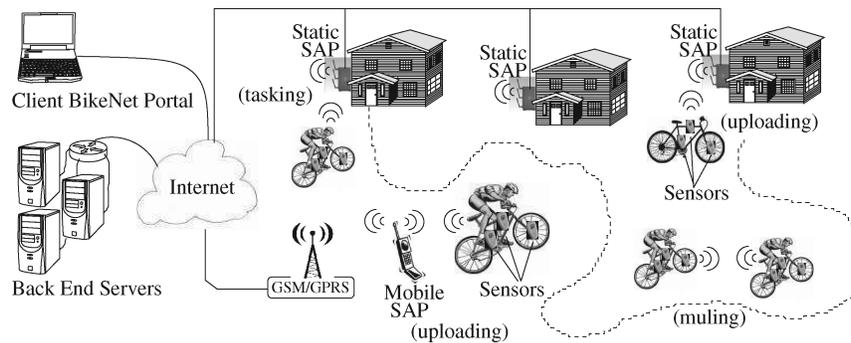


Fig. 1. BikeNet system overview. Sensors collect cyclist and environmental data along the route. Application tasking and sensed data uploading occurs when the sensors come within radio range of a static Sensor Access Point (SAP) or via a mobile SAP along the route. Sensed data muling can occur when cyclists come within mutual radio range. We collect data about the cyclist (heart rate, galvanic skin response), about the cyclist's performance (wheel speed, pedaling cadence, frame tilt, frame lateral tilt, magnetic heading), and about the cyclist's surroundings (sound level, carbon dioxide level, cars).

2.1 Hardware

The BikeNet system hardware is organized into three tiers, the backend server tier, the Sensor Access Point (SAP) tier, and the mobile sensor tier. In the following, we discuss the design and implementation of each tier, along with information on ruggedizing and calibration of the hardware.

2.1.1 Mobile Sensor Tier. The mobile sensor tier incorporates a number of bicycle-mounted and human-mounted Moteiv Tmote Invent [Moteiv Tmote Invent 2007] mobile sensing platforms. Together these sensors gather data concerning cycling performance, cyclist health and fitness, and the environment surrounding the cyclists' routes. The Tmote Invents mounted to a particular bicycle, along with those mounted to the human riding the particular bicycle, constitute a BAN. Intra-BAN communication occurs via short-range IEEE 802.15.4 radio. The BAN architecture is designed in a modular way such that sensing components can be added or subtracted simply according to user preferences (dynamically) set in software. Figure 2 shows a logical representation of the *Bicycle Area Network* (BAN), and Figure 3 shows a prototype sensor-enabled bicycle.

We use the native sensors provided by the Tmote Invent: a two axis accelerometer, a thermistor, a photodiode, and a microphone. The Nokia N80 offers a camera, microphone, and multiple radios that can be used as sensors. In follow-on work [Miluzzo et al. 2008a], we have focused on the mobile phone as the primary sensor and interface allowing for user annotation and interaction with her data. In our current BikeNet prototype, we also interface a number of additional sensors to the Invent. We process the accelerometer data to measure the angle of incline, and lateral tilt of the bicycle. To measure the angular velocity of the wheel and pedal, forward speed, and distance traveled, we attach a magnet-triggered reed relay mounted across the Invent's user button.

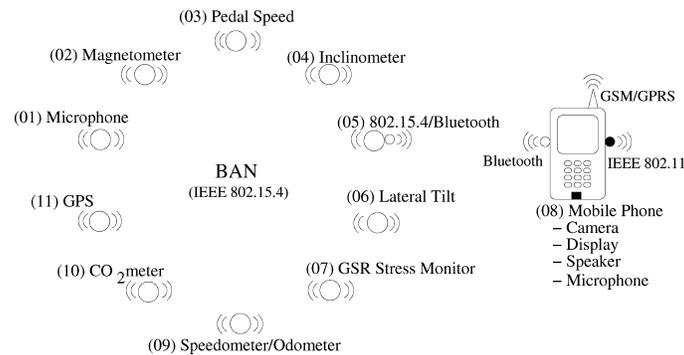


Fig. 2. Logical representation of bicycle area networking. Sensors share a common IEEE 802.15.4 channel. A mobile phone plays a dual architectural role depending on whether its cellular radio is active/connected. If connected to the cellular backend the mobile phone acts as a Mobile Sensor Access Point (SAP) facilitating real-time sensing; else it acts as a local member of a BAN engaged in delay-tolerant sensing.

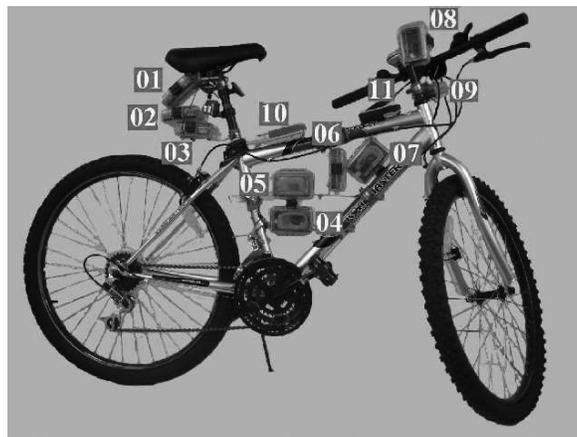
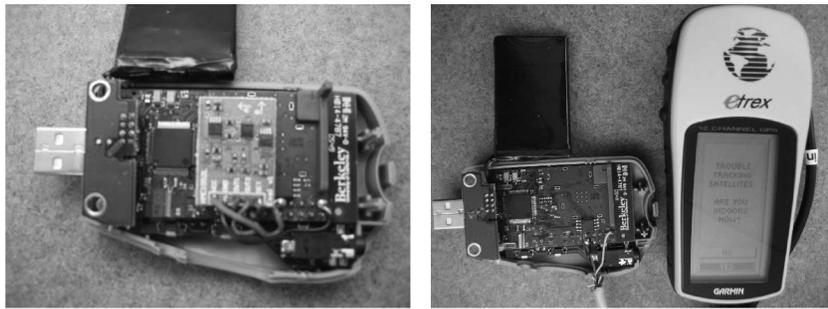
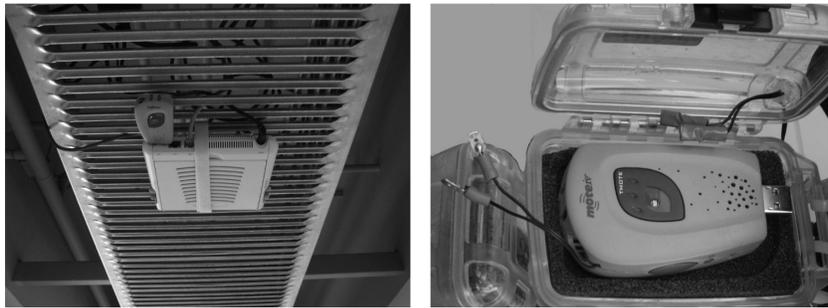


Fig. 3. Physical implementation of the BikeNet system. Numbered sensors installed on the bicycle map to the sensor types labeled in the logical BAN representation in Figure 2.

Every time the relay closes (every pedal/wheel rotation) a TinyOS [TinyOS 2007] interrupt event is generated. To measure direction and deviation with respect to the Earth's magnetic field, we add a dual axis magneto-inductive sensor (Honeywell HMC1052L) by connecting the sensor output to two ADC channels on the Invent and connecting a free I/O pin from the Invent's MSP430 microcontroller configured as output to act as a digital control line. We further process the magnetometer data for use as a metal detector, and in particular for automobile detection. To provide a common notion of absolute time and location within a BAN, we connect a Garmin Etrex 12 channel GPS unit (Figure4(b)) via the UART0 port of the Tmote Invent's MSP430. The Garmin Etrex provides time and location data at the fixed rate of once per two seconds via its RS232 interface. To measure the carbon dioxide levels in the environment



(a) A two-axis magnetometer is attached to a Tmote Invent via its ADC. (b) An external GPS unit is attached to a Tmote Invent via its UART0 port.



(c) A BikeNet static SAP is a WiFi AP with an Invent inserted in the USB port. (d) Waterproof OtterBox. Wires are fed through drilled holes that are then filled with silicone sealant. Wires have crimped connectors for easy disconnect.



(e) Ground-truth video/sound/photo helmet with (f) BikeNet mobile SAP implementation. The four N80s and GPS receiver, only for use in de-Nokia N80 Bluetooth radio associates with a bugging our system and validating our inference custom-built Bluetooth/802.15.4 gateway techniques.

Fig. 4.

surrounding the cyclist, we interface the standard Tmote Invent with the Telaire 7001 CO₂/Temperature Monitor, via an ADC port of the Tmote Invent's MSP430. To measure the galvanic skin response of the cyclist, we use an ArcherKit Biofeedback Monitor connected to a Tmote Invent. Wires connected to the fingers of the cyclist measure epidermal microcurrents.

The Tmote Invents, Nokia N80 mobile phones, and external sensors are powered using rechargeable batteries in our prototype. A commercial implementation could leverage ongoing work in energy harvesting (particularly from pedaling and frame vibration) to reduce the need for external recharging. This is outside the scope of the current work.

2.1.2 SAP Tier. The SAP tier offers high performance, high reliability, and secure gateway access from the sensor tier to the backend servers. This access allows sensed data to flow to the system repositories, and provides a point of command for the architecture to task available sensors with user application requests/queries. When possible, these gateways are symbiotically implemented on the back of existing network infrastructure by plugging a short-range radio module into the existing network element (e.g., IEEE 802.11 access point), allowing it to communicate with the sensor tier. SAPs can be static and wired directly to the Internet, or can be mobile and use a data service over a wide area radio access network to provide connectivity to the backend (e.g., mobile phone with GSM/GPRS). We study both tasking and uploading via both static and mobile SAPs in our implementation. SAPs are also equipped with sensors to provide context and validation for uploaded data.

The static SAP is implemented using an unmodified Tmote Invent plugged into the USB port of an Aruba AP-70 IEEE 802.11a/b/g access point. The Aruba is running a customized version of OpenWRT, an embedded Linux variant. The BikeNet SAP is implemented as an overlay of tools requiring only user privileges. Certain kernel module support is needed; modules are loaded at runtime if necessary. The tools distribution is cleanly encapsulated in a single tarball making symbiotic deployment of a BikeNet SAP on to a standard WiFi access point easy to manage. The mobile SAP is implemented using a Nokia N80 paired to a custombuilt Bluetooth/802.15.4 gateway via its Bluetooth radio. The N80 SymbianOS uses a serial device emulation of the Bluetooth SPP profile to read and write from the Bluetooth/802.15.4 gateway. The backend interface of the SAP uses GSM/GPRS to the BikeNet repository and backend services. This is done with a combination of SMS messages from the backend pushed to the phone, and TCP connections initiated by the N80 to transmit responses to a backend server that translates data uploads to SQL commands to insert data into the repository.

The use of a personal device like a cell phone as a mobile SAP gives rise to an interesting dual role for the N80 in our system. Architecturally, there is a clean separation between SAP and sensor tiers, but in the case of a mobile phone owned by the cyclist the BAN to which the cyclist belongs may have continuous access to the SAP services and resources whenever GPRS service is available. A mobile phone thus plays a dual architectural role depending on whether its cellular radio is active/connected. If connected to the cellular backend the mobile phone acts as a mobile SAP facilitating real-time sensing and access to backend services such as the *checkpoint* service discussed in Section 2.2.6; else, it acts as a local member of a BAN engaged in delay-tolerant sensing. GPRS pricing and performance also comes into play when using the cell phone as a mobile SAP.

2.1.3 *Server Tier.* Members of the backend are Ethernet-connected servers equipped with practically unbounded storage and computational power. These provide a number of services to the architecture, some of which are described in Section 2.2. In particular, it is to the backend servers that system users connect to submit application requests/queries for execution in the sensor tier, and to retrieve and visualize sensed data.

2.1.4 *Ruggedizing the Hardware.* Because of the outdoor nature of the BikeNet testbed we take steps to protect the Tmote Invents from the weather (e.g., rain, snow) by enclosing each in an OtterBox 1600 Case. The OtterBox comes with adhesive foam that is customizable to a degree that allows us to secure the Tmote Invents inside the cases without any slipping. A number of sensors require running wires from the Tmote Invent out of the OtterBox to other places on bicycle or cyclist (e.g., the WheelSensor’s reed relay is wired to the front fork of the bicycle). For these we drill holes through the OtterBox 1600 and fill the holes with silicone gel after passing through the wires to maintain waterproofing. We cut the wires inside the box and crimp/solder on connectors (see Figure 4(d)) to allow a quick disconnect of the Tmote Invents for recharging. Additionally, the Otterbox cases are securely fastened to the bicycle frame, using a system of steel mounting bars and steel hose clamps, since bicycling implies often severe vibration and jolting. The OtterBoxes are screwed to these mounting bars and the screw holes are sealed with silicone gel. In determining the geometry and placement of the mounting bars we have attempted to minimize vibration and unwanted degrees of freedom for the sensors (a picture of a sensor-enabled bicycle appears in Figure 3).

2.1.5 *Calibration/Validation.* Despite efforts to mount the accelerometers at perfect right angles (in two dimensions) with the ground, we find that calibration is required for each bicycle in order to correctly understand the measured values. Even if the error angle of the mounting bracket is small it can lead to a large skew in the calculated slope, because of the nonlinear nature of the inverse tangent function used to calculate the slope. Stationary calibration is done in the lab by matching the bicycle-mounted accelerometer outputs against a set of known inclines to derive a calibration curve for each device. To validate this static calibration in the field, we manual measure a 0.75km section of the road containing slopes from 0 to 7 degrees using a laser level (model TUV EPT-97A, 650nm) at 30m intervals. We receive excellent correlation between manual measurements and those made using the accelerometer (the TiltSensor role).

We find that calibration is also necessary for the magneto-inductive sensors due to the steel frame of the bicycle, and the steel mounting bars. This is done by executing a hard/soft iron calibration [Caruso 2000] for each bicycle, and adding the correction for the magnetic declination of Hanover, NH, USA.

We infer cyclist fitness level using a combination of the lateral tilt, slope, and pedal speed to wheel speed ratio. To check our inference technique against a more direct physiological measure of cyclist fitness, we use the Garmin Forerunner 301 Heart Rate Monitor/GPS. A positive correlation between our inferred

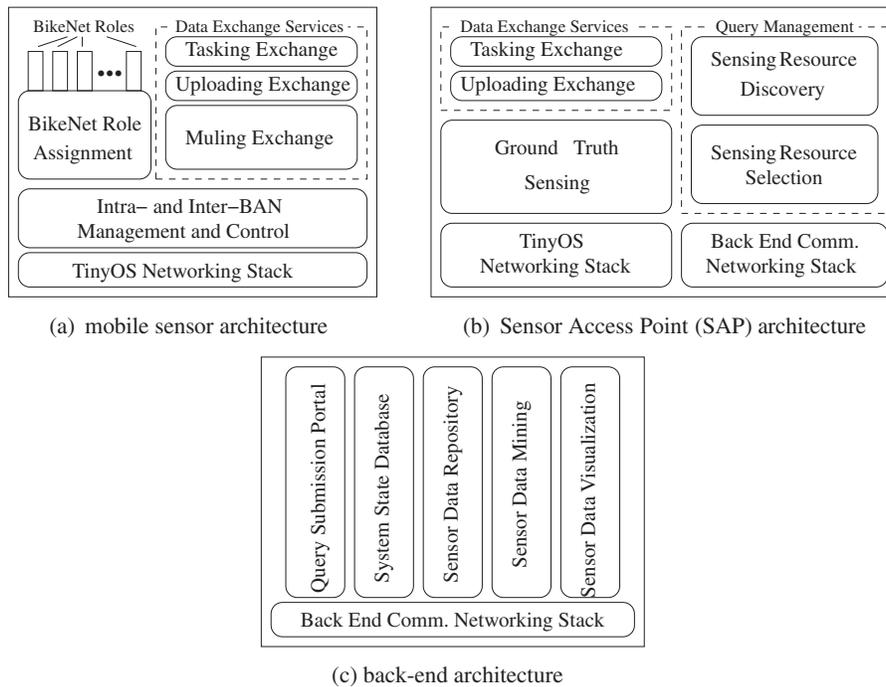


Fig. 5.

cyclist fitness level and that indicated by the actual cyclist heart rate validates our technique.

To provide richer context for the sensor measurements and inference we do in BikeNet, we attach four Nokia N80 phones on a bicycle helmet, that is, facing front, back, left, and right (see Figure 4(e)). Using continuous video capture (both visual and audio) throughout the ride we are able to validate that events sensed/inferred by BAN sensors are at least reasonable/probable and depending on the measurement type we can definitively validate the data (e.g., car passing the bike or not).

To validate detection-based inferences, we use a standard Tmote Invent programmed to write the (time, location) 2-tuple to the Flash every time the user button on the Tmote Invent is clicked. We term this the *ButtonMote* for ease of reference. For example, in testing the MetalDetector (Section 3.1.1) we manually click the ButtonMote user button every time we pass a parked or moving automobile or an automobile passes us, and compare the time/location-aligned MetalDetector trace with the ButtonMote trace to determine detection accuracy.

2.2 Software

Figures 5(a), 5(b), and 5(c) show how the BikeNet software system maps to the three tier hardware architecture, respectively defining the mobile sensor, SAP, and back-end software subarchitectures. In our implementation, communication between the SAP and back end subarchitectures is via either a TCP/IP

BAN Hardware	BAN Roles
Invent	PersonalNode
Invent using its accelerometer	TiltSensor
Invent using its accelerometer	LateralTiltSensor
Invent using its microphone	SoundSensor
Invent + magnetometer	MetalDetector, CompassSensor
Invent + reed relay mounted on pedal	PedalSensor
Invent + reed relay mounted on wheel	WheelSensor
Invent + GPS	SyncSprinkler
Invent + 802.15.4/Bluetooth Gateway + N80	CameraSensor, LocalDisplay
Invent + CO ₂ meter	CO ₂ Sensor

Fig. 6. Mapping between BAN hardware and logical roles.

stack (for static SAPs) or a GPRS/GSM stack (for mobile SAPs). Primary software elements are discussed in the following.

2.2.1 BikeNet Role Assignment. For purposes of modularity the functional requirements within a BAN are divided into logical *roles*. The *PedalSensor* and *WheelSensor* roles measure the angular velocity of the pedal and front wheel, respectively. From these the current and average speed, distance traveled, pedaling cadence and gear ratio are measured or inferred. The *TiltSensor* role measures the angle of incline of the bicycle frame with respect to the gravitational force vector, allowing for real-time slope calculation and a mapping of the terrain along a cyclist's route. The *LateralTiltSensor* role measures the lateral angle of incline of the bicycle frame. The *CompassSensor* role measures the instantaneous angle of the bike frame with respect to the Earth's magnetic field, allowing for a form of dead reckoning when used in combination with speed and distance information obtained from the *WheelSensor* role. The *MetalDetector* role measures distortions in the the Earth's magnetic field caused by nearby ferromagnetic metals, allowing inference of the amount of passing automobile traffic. The *SyncSprinkler* role provides a common absolute notion of time and location to all members of the bike area network via periodic short-range broadcasts. The *LocalDisplay* role provides a means to locally display sensed data. The *CO₂Sensor* role measures the carbon dioxide content in the atmosphere surrounding the bicycle, allowing the system to infer whether the cyclist is passing through an urban area (more CO₂ from auto exhaust) or a rural area (less CO₂ due to plant respiration). The *SoundSensor* role measures the volume of noise in the environment surrounding the cyclist, and is used for voice-triggered sensing and audio annotation of a cyclist's ride. The *CameraSensor* role provides triggered capture of an image, or a video clip of specified duration. The *PersonalNode* role provides control via short-range radio over the other sensing roles, including executing user preferences within the BAN (e.g., required sensors, sampling parametrization), and signaling the start and stop of a cycling trip. Each cyclist necessarily possesses a *PersonalNode*, but all other roles are optional, depending on the sensing preferences of the cyclist. Figure 6 shows the mapping between BikeNet roles and the sensing hardware, where each row represents a different (set of) devices on a fully equipped prototype bicycle.

We assume that each cyclist possesses a mobile personal computing device (e.g., Tmote Invent, Nokia N80, radio-equipped Apple iPod) at all times that can be tasked by the SAP to take on the PersonalNode role. In our prototype system, each cyclist carries a Tmote Invent preconfigured with the PersonalNode role. The PersonalNode role includes a list of user preferences that dictate what additional sensing roles are desired to quantify the cyclist fitness/performance/environment. These sensing roles are split into two lists, required and preferred, that are included into a *hello* beacon periodically broadcast by the PersonalNode. The *hello* beacon also includes the required sensing parameterization (e.g., sample rate). Each available mobile sensing platform (i.e., Tmote Invent) that receives the beacon replies with a *hello reply* if its sensing capabilities match either a required or preferred role requested in the *hello* beacon. The *hello reply* indicates which role(s) the respondent is offering to fill. However, a recipient of the *hello* that is already associated with another PersonalNode will not reply. Upon receiving a *hello reply*, the PersonalNode first registers the respondent and the role(s) it is offering to fill, and then sends a *hello reply ack* to complete the association. The *hello reply ack* contains a list of identifiers reflecting the current BAN membership. Subsequent *hello* beacons sent by the PersonalNode do not request sensing roles that are already being filled by associated mobile sensing platforms. If a *hello reply ack* is not received in response to a *hello reply*, the *reply* is retransmitted up to N times. If after N times the *ack* is not received, then it is assumed that mobility has carried the PersonalNode and potential sensor out of range and the partial association state is purged.

Though generally a mobile sensing platform may, depending on sensing capabilities, be able to take on more than one BikeNet role, for our current prototype implementation we allow only one role per sensing platform to work within hardware limitations of the Tmote Invent. For example, there are a limited number of free configurable I/O pins and ADC channels available for external sensors on the Tmote Invent, and a shortage of Flash/RAM. In the future with more capable hardware, we will be able to condense our current implementation, assigning multiple roles to a single sensing platform and greatly reducing BAN complexity and costs (i.e., monetary, radio congestion, energy consumption). In the rest of the article, unless stated we treat the logical sensing role and the mobile sensing platform to which the role is assigned as synonymous.

2.2.2 Intra-BAN and Inter-BAN Management.

Localization and synchronization. The SyncSprinkler role provides, via a periodic broadcast within the BAN, periodic samples of the instantaneous absolute time and location. In our implementation these are obtained from a GPS unit. The SyncSprinkler controls its transmission power to limit the scope of its beacons, thereby maintaining a higher location accuracy for all broadcast recipients. All BAN members' time estimate is updated externally with the values contained in the SyncSprinkler broadcasts, and internally via a local clock set to provide higher time resolution between received SyncSprinkler broadcasts.

Sensing control. When the PersonalNode has established associations with sensing platforms (i.e., Tmote Invents) sufficient to meet all the roles specified by the user preferences, an LED on the PersonalNode indicates a “Ready” state. A button click on the PersonalNode when in this “Ready” state sends a *start* message broadcast from the PersonalNode indicating that the ride is beginning and sensors should start collecting data with their prescribed parametrization. This message is acted on by mobile sensing platforms that are associated with that PersonalNode, moving both the PersonalNode and the associated mobile sensors into the “Started” state. If associated sensors do not receive a *start* message within a timeout period, the association times out and the mobile sensors are free to associate with another PersonalNode at that time. A subsequent PersonalNode button click while the PersonalNode is in the “Started” state sends a *stop* message broadcast, signaling the end of the ride. The *stop* message causes mobile sensors associated with that PersonalNode and in the “Started” state to cease sensing.

Event-triggered sensing vs. continuous sensing. Sensing is set up to occur either continuously or only when triggered by other events. In the continuous case, the user preferences executed by the PersonalNode parametrize the sensing capture (e.g., sampling rate, duration, local processing functions) that takes effect immediately upon receiving the *start* message of the sensing control protocol. Continuous sensing in BikeNet is appropriate for roles such as the Tilt-Sensor where terrain mapping should be continuous. On the other hand, some sensing operations may be too energy expensive for a mobile sensing platform to do continuously, or may not have meaning except under certain contexts (e.g., certain locations of interest, or under certain sensed circumstances). Triggers are defined by dynamically updatable user profiles executed by the PersonalNode that specify the conditions under which sensing should occur.

The BikeNet implementation support of triggered sensing includes methods to define and submit sensing triggers and actions to the PersonalNode for execution within the BAN. Upon receiving the triggered sensing definition, the PersonalNode breaks apart the conditions that must be met for the action to take place, and reliably transmits each condition (e.g., “slope > 5 degrees”) to the BAN member suited to evaluate the condition (e.g., the TiltSensor). When a condition evaluates to true, the BAN member signals the PersonalNode. The PersonalNode initiates the action when all conditions for a given triggered action are met. We are currently focusing on triggered photography, video, and audio capture using the camera and microphone on the N80, when certain conditions in the BAN are met, but we also implement support for a number of other actions such as sending data to be displayed on the LocalDisplay, sensing something at a different parametrization than the current one, playing a sound on the N80 or Tmote Invent speaker, transferring sensed data from one Tmote Invent to another, and blinking LEDs.

Real-time feedback/display. The local display protocol is used by the LocalDisplay to query other BAN members for values to display. The LocalDisplay is provided by a handlebar-mounted N80 mobile phone, via the Bluetooth/802.15.4 gateway shown in Figure 4(f). The Tmote Invent’s hardware design shares the

SPI bus between radio and flash, and the same physical microcontroller pins are used for UART0. Since the Bluetooth-to-Serial converter is connected to the Tmote Invent UART0 port, this precludes simultaneous radio communication and display updating. Hence, BikeNet uses a simple TDMA-like time slot assignment on top of the TinyOS CSMA MAC to improve communication between the roles generating sensor data and the LocalDisplay. The LocalDisplay periodically broadcasts a query for data and the sensor roles register the first LocalDisplay they hear a query from as the only LocalDisplay they will reply to thereafter. This association times out after a period if no queries are heard from the LocalDisplay. The data that a given sensor role returns to the LocalDisplay is a matter of user policy, but a typical display includes speed, distance traveled, bike frame tilt angle, pedal RPMs, and time of day. To support flexibility in the user configuration of the display, data is represented in the packet with (type,length,value) format.

2.2.3 Data Exchange Services. Three types of data exchange occur in the BikeNet system: tasking exchange, uploading exchange, and muling exchange. The tasking and uploading data exchanges take place between mobile sensing platforms and SAPs. The muling data exchange takes place only between members of the mobile sensing tier (e.g., Tmote Invents). As default, BikeNet uses a delay-tolerant mode where a BAN's PersonalNode mules data for sensing roles in its BAN (up to the limits of its available storage) and uploads the data via wire or a wireless upload protocol. Inter-BAN muling and in situ uploading via either mobile SAPs or opportunistically encountered static SAPs support queries from backend user applications that may want data in a more timely manner.

In the BikeNet tasking exchange, a SAP interacts with available mobile sensing platforms (e.g., Tmote Invents) to first instantiate a PersonalNode programmed with a cyclist's BAN preference profile. Based on this profile, the PersonalNode assembles a BAN by tasking other available mobile sensing platforms with the required sensing roles as discussed in Section 2.2.1. Aside from this BAN bootstrapping, the tasking exchange also includes the handling of user queries/requests for data by backend system users, received via the SAP. The PersonalNode responds to these queries by invoking the necessary continuous or triggered sensing (Section 2.2.2) within its BAN.

In the muling exchange, sensed data is transferred between mobile sensors outside of the radio range of either a mobile or static SAP. A simple muling protocol is implemented on every Tmote Invent, but the option to activate muling (i.e., spend Flash space to carry others' data) is set by cyclist preference. The protocol uses an *advertisement-accept-data* exchange, where the *advertisement* specifies the amount of data the provider wants to have muled, the *accept* message indicates the amount of data the consumer is willing to mule (based on Flash constraints), and the *data* message represents a burst of data packets from the producer to the consumer. In addition, Stop-and-Wait ARQ with a maximum of three retransmissions provides for reliable transfer of the data packet burst. If a producer still receives no acknowledgment after three retransmissions of the same packet it will assume the session is over and begins advertising

anew. Our implementation includes support for replication of sensed data (i.e., via the muling exchange) but the replication of muled data is not allowed. Restricting the right to replicate to the data origin allows it to maintain control over the number of copies of its data that are circulating and also to vet (in terms of trustworthiness) all candidate mules.

In the uploading exchange, when a BAN comes within the radio range of a mobile or static SAP, the Tmote Invents composing the BAN attempt to upload sensed data to the backend data repository. The upload protocol message exchange is identical to that of the muling protocol. When a SAP receives data packets, they are forwarded (in both the mobile SAP and static SAP cases) to the backend repository. The decision to accept new upload sessions is made based on channel congestion around the SAP.

2.2.4 Ground-Truth Sensing. In the BikeNet sensing system, SAPs are equipped with certain sensors and can provide *ground-truth* measurements. Ground-truth¹ sensing refers to a trusted, high fidelity, always accessible stream of data. One use of ground-truth data is as a filter applied to data uploaded from a sensor before the data is passed by the SAP to the backend repository. The ground-truth filter can be applied to validate or invalidate uploaded data when the uploaded data samples and ground-truth data samples have a high expected correlation (e.g., temperature sampled at the same location and at the same time, samples triggered by the same set of circumstances). Further, ground-truth sensing can be used to satisfy queries coming from a back end system user that have only coarse location context requirements. Ground-truth data is also used to satisfy queries coming from a BAN in the radio range of a SAP. In this case the BAN can ask for readings from the SAP's ground-truth sensors, for example, as part of a self-calibration routine.

2.2.5 Query Management. The query management component on the SAP handles queries both from the backend system user, and from the PersonalNode of a BAN. It invokes a sensing resource discovery routine to determine what sensing resources are available to meet the sensing request. The routine checks both any ground-truth sensors on the SAP itself (Section 2.2.4) and available sensing resources on any BANs that may be within radio range of the SAP. Once a list of available sensing resources is compiled, the SAP invokes a sensing resource selection routine to decide which resources will be tasked in order to satisfy the request, and invokes a tasking routine to execute the necessary request (i.e., a simple function call if the resource is on-SAP, or via the tasking exchange (Section 2.2.3) if the resource is in a BAN in radio range. In the BikeNet implementation, we have experimented with handling queries to the SAP, originating both from the backend and from a BAN in radio range, for ground-truth data. In particular, using a cellular phone as a mobile SAP, we have experimented both with event-triggered capture of images, sound and videos requested by the BAN (e.g., an audio annotated ride); and with direct requests from the backend BikeNet Web portal for image, sound, and video samples.

¹While related in principle, this notion of ground-truth sensing should not be confused with sensing for experimental validation and debug (e.g., using the quad capture video helmet in Figure 4(e)).

2.2.6 Route Checkpoints. Route checkpoints are virtual location-specific storage entities designed to allow cyclists to store and retrieve location-specific data. A checkpoint is meant to maintain performance-related characteristics about the rides of individual cyclists in a particular location (e.g., average speed up a particular hill). Conceptually, each time a cyclist approaches a checkpoint (as determined by the BAN's GPS receiver) her previous statistics for that checkpoint are automatically downloaded via the cellular data connection from the checkpoint server on the backend, notifying the cyclist of past performance via the LocalDisplay, perhaps motivating the cyclist to improve. Additionally, average performance characteristics across multiple cyclists may be available for display (according to policy), enabling a form of competitive challenge between cyclists that ride along the same route even if they are not cycling together. The current characteristics are automatically uploaded to update the checkpoint record (see the following) when she completes the checkpoint. Checkpoint updates may be done in real time via the cellular data channel of the PersonalNode, or opportunistically during a SAP rendezvous (e.g., when the cyclist returns from her route).

Checkpoints may be created by the cyclist through a user interface on the LocalDisplay by indicating the start and end of the checkpoint via button presses. Alternatively, checkpoints may be created via a Web interface before the ride by defining the GPS coordinates defining the checkpoints start and stop locations. Either way, the checkpoint is registered in the checkpoint server as a record containing the creator's user ID, the start location, stop location, and a privacy policy indicator of whether the checkpoint is private (only those approved by the creator can interact with the checkpoint) or public (all cyclists may interact with the checkpoint). Performance characteristics, including average slope angle, average speed, average gear ratio, and distance, for the checkpoint are appended to this record on a per-user basis (pursuant to the privacy policy) and are recalculated with every upload.

In practice, a list of checkpoints may be preloaded onto a cyclist's PersonalNode (e.g., cell phone) prior to the ride. Alternatively, partial lists of checkpoints can be swapped in and out along the route, that is, load a list of upcoming checkpoints and unload checkpoints the cyclist has bypassed. Whether preloaded or dynamically fetched, interaction with a particular checkpoint is governed by the privacy policy set when the checkpoint is created.

2.2.7 Query Submission Portal. The BikeNet backend includes a Web portal (*BikeView* [BikeView - the BikeNet Web Portal 2007]) containing a graphical presentation of a cyclist's data, but also allowing for the real-time querying of BANs using a GPRS connection via the N80, if the cyclist is using such a device for her PersonalNode. The user can select the BAN of interest and assemble a query to submit to the query manager component of that SAP using a collection of pull down menus. A final mouse click transmits the query over the cellular network to the selected mobile SAP. We implement the ability to query a BAN's location, capture a camera image, and sample the microphone via this portal interface by sending SMS over GPRS to the N80.



Fig. 7. BikeView portal for data display and query submission. A CO₂ map of Hanover, NH, USA streets on a summer weekday afternoon is shown.

2.2.8 Sensor Data Storage, Processing and Visualization. The sensor data repository provides a location for the long-term storage of cyclist experience data on a per-cyclist basis and also provides a convenient location for the aggregation of all long-term trace data for all participating cyclists. Access to particular data is a matter of the policy that each cyclist registers with the repository (or a separate access control entity). The sensor data mining component provides a set of standard statistical functions and reusable calculations/data transformations that a user (e.g., cyclist) can invoke to control the retrieval and presentation of data. For BikeNet we use a number of data interpretation and inference tools and techniques, including scatter plots to look for data correlation, fast Fourier transforms (FFTs) to look for periodicity, running averages to smooth data to look for trends, and interpolation to align samples according to distance. For example, based on noise observed in the raw data from the Tilt-Sensor (see FFT of tilt data converted to WAV format and analyzed in Audacity, Figure 9), we apply a smoothing method to filter spurious vibration from the data before calculating the slope of the cyclist's path. Currently, the data handling is done in a nonautomated way by storing raw data streams in flat text log files, processing these files using Awk scripts to extract data of different types and apply methods for smoothing, averaging, and scaling. Further, we develop scripts to transform data values into *BikeView* [BikeView - the BikeNet Web Portal 2007] visualizations. With BikeView (see Figure 7), we present summarized collected data sorted by user, and sorted by ride within each user account (akin to the presentation of "My Runs" on the Nike+iPod Web page [Nike+iPod Sports Kit 2007]). Detailed sensed data can be obtained by simple mouse hovers and clicks over the graphic representations of different rides. The vision is to provide backend sharing between users facilitated by dynamic creation of

group pages that are visible to all users in the group and to which all group members can publish data.

2.2.9 System State Database. The system state database contains both static state information (e.g., Tmote Invent physical address and sensing capabilities, PersonalNode human custodian information) and dynamic state information (e.g., last known position of a mobile SAPs, and BANs, SAP load average) about elements in the network. In particular, the system state database tracks which BANs are currently in radio range of mobile and static SAPs. This facilitates proper query routing from the back end query submission portal to particular BANs, for BAN-specific user queries. The information is also valuable more generally for debugging and management of the network.

2.3 Incentivizing BikeNet

People-centric sensing systems [Eisenman 2008] like BikeNet have the distinguishing property that at least some of the sensors are owned and maintained by the plebes, rather than by a scientific, industrial, or government establishment. To facilitate the mass adoption of people-centric systems, users must be incentivized to participate. BikeNet operates simultaneously as a personal sensing system, a group sensing system, and a municipal sensing system, depending on how queries to sample data enter the system and how the resulting sensed data is shared. In the following, we discuss how to incentivize BikeNet at these three levels, focusing on the issues of resource consumption and user privacy.

2.3.1 Resource Consumption. In the sampling phase, resources such as battery energy and local storage are certainly consumed from the cyclist's cell phone when it assumes the PersonalNode role, and from other sensors embedded in the bicycle itself. However, the sampling itself is undertaken for the direct benefit of the cyclist so there is no special need for incentivizing this phase. The MetroSense architecture [Campbell et al. 2006] on which BikeNet is based uses mobile sensors (e.g., BANs) in part as a publicly accessible sensing platform where externally submitted queries may run. While this is not the focus operating regime for BikeNet, in this case a quid pro quo scheme (e.g., serve a query, get a token to make a query) with the necessary supporting AAA services may suffice. Note that for BikeNet in particular, the net energy cost of sensing may be low: the batteries of embedded BAN sensors may be recharged by harvesting energy from the pedaling of the cyclist, and the user's cell phone can be recharged daily by the user.

For the upload phase, muling may be used to help upload the data to the server tier. Since muling consumes energy, storage, and radio bandwidth in the handling of others' data, system users may be hesitant to participate. In our implementation of muling for BikeNet, the user profile loaded into the PersonNode role indicates the amount of storage the user is willing to use for muling others' data. Beyond this amount, the BAN ignores muling solicitations from others. This limit on storage also acts as a coarse control on radio bandwidth and energy spent for others. Note that BikeNet operates fine in the case that all nodes

refuse to allocate any storage for muling. The penalty is that upload delay is increased, but this may not be an issue. We anticipate that most users will be most interested in viewing their data locally on the bicycle, and after they have completed the route and returned home (e.g., a la Nike+iPod). Beyond this, we conjecture existing proposals for incentivizing peer-to-peer networks [Gupta and Somani 2004] and mobile ad hoc networks [Zhang et al. 2007] can be applied to data muling in BikeNet.

2.3.2 Privacy and Trust. Given that the sensed data in BikeNet originates from the cyclist and her surroundings, a careful consideration of privacy risks to the system users is required. Further, since an important aspect of BikeNet is the sharing of collected data with others, the degree to which data originating from others can be trusted should be considered.

As a personal sensing system, these issues are addressed using existing technologies. In this mode the end consumer of the data is the cyclist herself so data trust is not an issue, and the main privacy threat comes from the muling process. This threat can be addressed by using standard encryption techniques proposed for the sensor networking domain [Karlof et al. 2004] [Luk et al. 2007], or more advanced techniques depending on the computational capabilities of the PersonalNode.

Once the data resides in the cyclist's secure BikeNet data repository, she may choose to make selected data available to, for example, members of a cycling interest group to which she belongs, and she may further share data with the public as part of a community action initiative, for example, on pollution. When users export data in this way across the Internet, we make the reasonable assumption that existing technologies such as SSL or IKE are used as appropriate to provide for data security. Contributions to such public forums may be anonymous to protect user privacy (e.g., routes and times), or may be identifiable to allow for ratings-based data trust systems, a la Amazon.com. Addressing this tension between user privacy and trust on the Internet is outside the scope of this article, but has received much attention in the academic literature (e.g., Pai et al. [2008], Johnson et al. [2007]) and the public press (think uproar over Facebook privacy policy). A recent report by the OCLC [De Rosa et al. 2007] gives a largely nontechnical overview of many central issues in this area.

3. SYSTEM EVALUATION

We build five fully equipped BikeNet bicycles, implement all of the aforementioned sensing roles using Tmote Invent motes and Nokia N80 mobile phones, build a number of static and mobile SAPs, and implement a functional back-end Web portal offering query submission and data retrieval services. In this section, we present selected results from several groups of experiments respectively targeted at: quantifying the cyclist experience from sensed data collected about a single cyclist and his environment; looking at performance aspects of key BikeNet subsystems; and measuring the real-time performance of a deployed system across the Dartmouth campus and in adjacent areas of the town of Hanover, NH, USA. We use a common path that we call the ground-truth

route. This route includes a variety of urban cycling terrain, including built-up busy roads in the town center with lots of cars and pedestrian traffic and quiet back roads with little or no traffic. The route exposes cyclists to a variety of flat terrain, gradual down hill, and steep uphill sections. Typically the ground-truth route takes 25–30 minutes to ride and is nearly 5km long. The experiments are conducted at rush hour and in the middle of the day when there is less traffic and activity. We conducted many experiments over the period August 2006–August 2007 collecting a typical dataset of 0.8MB per ride per bike. We record the runs using video from the video helmet (Figure 4(c)) that collects quad-directional video of a ride for ground-truth validation of our correlation/inference methods (not part of standard BikeNet equipment). BikeView [BikeView - the BikeNet Web Portal 2007] contains an example of one such video recording.

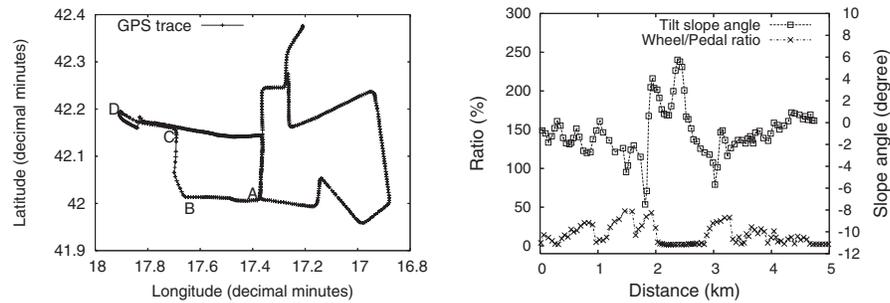
3.1 Cyclist Experience Mapping

3.1.1 Inference and Cyclist Fitness Sensing. In this section we present a series of plots characterizing cyclist behavior and the environmental conditions encountered during a ride. We collect data from each of the sensing roles mentioned in Section 2.2.1, and apply fusion techniques and trend analysis to extract additional information from the raw data.

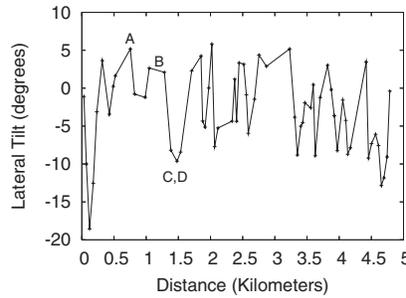
Figure 8(b) shows the measured slope profile of the route calculated from TiltSensor readings versus distance (the Wheel/Pedal ratio curve is explained later). The slope s is calculated according to $s = \arctan(x/y)$, where x and y are the TiltSensor’s measured x- and y-channel accelerometer readings, respectively. We register accurate measurements when the bike is stationary; error increases with speed and terrain roughness due to unfiltered vibrations and cyclist behavior. The slope profile which matches the manually measured ground-truth road segment (refer to Section 2.1.5) well (less than 10% deviation from the ground-truth slope measurements).

Figure 8(c) shows the lateral tilt plotted versus distance. The lateral tilt is calculated in the same manner just described for the TiltSensor. A cyclist’s aggressiveness in turning is inferred. From the plot we correlate the increases in lateral tilt magnitude shown on the y-axis, with corner turns expected from the mapped GPS trace shown in Figure 8(a). Positive lateral angle indicates a right-side lean whereas negative angles indicate left-side leans. In Figure 8(c), we label (viz. A, B, C, D) a sample of the lateral tilts that can be correlated with corner turns in the cyclist GPS trace (see Figure 8(a)), where at A, B, C, and D the biker makes, respectively, a right, a right, a left, and a left turn. The sharp left tilt (almost -20 degrees) is due to mounting the bicycle at the start of the ride.

The quantitative aspects of the cyclist fitness include the slope of the road/trail that the cyclist covers on his ride, the speed profile of the cyclist, the gear used when traveling up a given slope, and the location of the route. Figure 8(b) shows the slope profile of the road traversed on the cyclist’s trip, and the ratio of the tire/wheel speed to the pedal speed. This ratio infers the approximate gear the bicycle is in at a given point in the route, and provides a notion of the fitness of the cyclist. This indicator is most accurate when the cyclist is



(a) The mapped GPS trace of the cyclist route, comprising roads in the vicinity of Dartmouth College in Hanover, NH, USA. (b) Cyclist fitness is inferred by correlating the gear ratio inferred from the wheel/pedal ratio and the measured road slope.



(c) Lateral tilt plotted versus distance. Changes in lateral tilt are correlated with turns along the route shown in Figure 8(a).

Fig. 8.

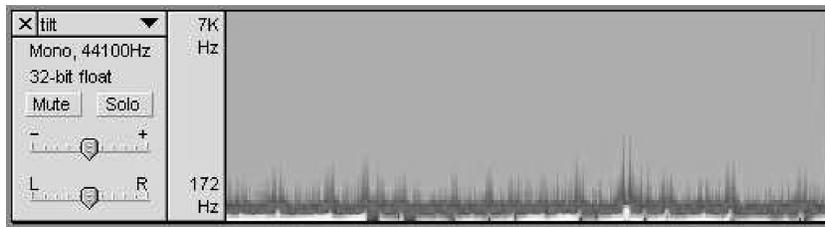
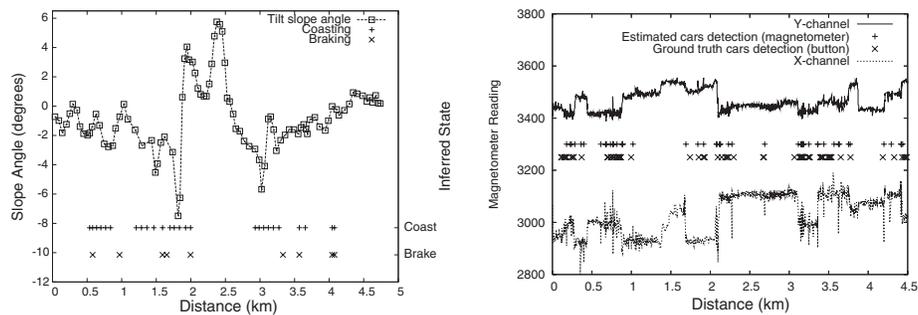


Fig. 9. We used the free audio application Audacity to compute the FFT's by turning our sensor data into a WAV file. Based on the vibration noise observed in the tilt data, we applied filtering before calculating the road slope.

going uphill, since when coasting downhill the pedals may not be moved much. A strong cyclist can use a higher Wheel/Pedal ratio when climbing hills. In Figure 8(b), intervals where the cyclist changes gears to climb hills are evident (from roughly 1 to 1.25km and from roughly 2 to 2.7km), where the wheel/pedal ratio is nearly 1.

Knowledge of the sensed path slope combined with the measured pedal speed and wheel speed allows us to infer when a cyclist is coasting or braking. On a



(a) Periods of coasting and strong braking can be inferred from relationships between pedal RPM, wheel RPM, and road slope. (b) From measured magnetic field distortions, we use a thresholding method to infer locations with dense patches of cars.

Fig. 10.

given bicycle there is a finite discrete set of pedal speed to wheel speed ratios possible when the bicycle chain is engaged with a gear and providing thrust to the bicycle. The cardinality of this set is equal to the number of gears the bicycle has. If the measured ratio of pedal speed to wheel speed does not match one of the allowable values we can infer that the cyclist is coasting. Braking can be inferred in a similar fashion to coasting. It is likely a cyclist is braking if the measured wheel speed slows while the slope is negative (downhill). Further, braking is likely when going uphill if the measured wheel speed slows faster than dictated by the slope of the hill. However, this is more challenging to detect since inference of uphill braking is also dependent on unknown quantities such as the combined mass of the bicycle and the cyclist, and the route surface composition/coefficient of rolling friction.

Figure 10(a) shows a plot of the road slope versus distance along the ground-truth route. Applying the simple inference technique of observing decreasing speed when the slope is negative (downhill), we infer sharp braking intervals at 1.6km, 3.3km, and 4.1km, which are verified by our known cyclist behavior. In these cases, we see a sharp decrease in wheel speed concurrent with a sustained downhill slope. Similarly, if the pedal speed is near zero and the wheel speed is high, we can easily infer the cyclist is coasting. In the figure, we infer periods of coasting from roughly 1.25 to 2km and 2.7 to 3.3km.

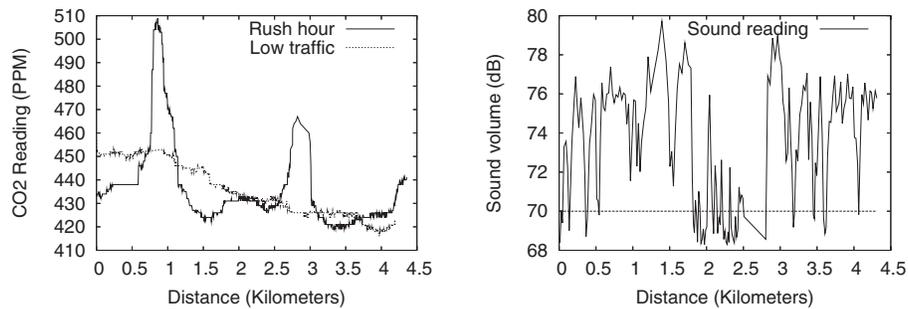
Aside from route topography and personal performance metrics, cyclists are interested in the ambience and safety of a route as a determinant in the overall enjoyment of the cycling experience. We take steps towards quantifying the ambience in terms of automobile traffic, air quality, and sound level. The presence of vehicles is often undesirable for cyclists who have concerns about safety, noise, or pollution. To infer automobile traffic along the cyclist route (Figure 8(a)), each BAN is equipped with a MetalDetector. When the MetalDetector passes close to any large metal body the Earth's magnetic field is deformed and the presence of a car is inferred. To collect ground-truth data for the experiment to compare against the inference from the MetalDetector, the car rendezvous event is manually logged by the cyclist with a ButtonMote click. We include the cases when the bike passes a car (parked or moving), and when

a car passes the bike. We find that detection of cars more than 2m away is unreliable with our hardware, so we do not try to log when a car passes in the oncoming lane (about 3m away). Each click generates a record of GPS time and location information. While, due to skew in button press times among various cyclists in the experiment, this method is somewhat error prone with respect to the exact location of observed cars, the logged button clicks do give an excellent idea of car density along the route. Figure 10(b) shows the raw x-channel and y-channel readings of a MetalDetector’s magnetometer plotted versus the distance covered along the ride. ButtonMote events and positions of inferred cars are overlaid on the same plot. The inference algorithm is run against both x and y channel data and works as follows. First, the exponentially weighted moving average of the magnetometer reading is calculated. If the difference between the current value and the moving average is greater than a threshold, and the current value is a peak (greater than both the preceding and succeeding values), a car is inferred. The threshold values for x and y channels and the moving average weight are learned by training with the readings from 0 to 1 kilometers. These trained values are then used along the rest of the route from 1–4.5 kilometers. While discrepancies between the ground-truth clicks and the output of the detection algorithm exist, we note that our aim is not counting the exact number of cars but answering the question, “Where may it be dangerous to bicycle due to lots of cars?”. Figure 10(b) shows our simple detection technique yields excellent overlap with the ground-truth in this regard.

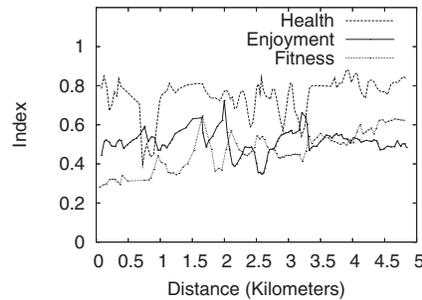
To provide a measure of air quality along the cyclist’s route we conduct experiments using a sensor measuring the level of carbon dioxide in the air surrounding the cyclist. Figure 11(a) shows a trace of the carbon dioxide sensor readings along the route shown in Figure 8(a) for two different cases, namely, rush hour and low traffic. The peaks in the rush hour case occur when the biker was cycling on downtown roads with a considerable presence of cars. In fact, variation in carbon dioxide levels measured on roadways is likely to be the result of automobile exhaust. While carbon dioxide has low toxicity at all levels we recorded during our experiments, it can act as a predictor of other noxious automobile exhaust constituents such as hydrocarbons, nitrous oxides, and particulate matter. Thus, from readings of the carbon dioxide sensor we can infer how enjoyable the traveled route is for a cyclist from the standpoint of pollution. The portal snapshot in Figure 7 shows a CO₂ map of the Hanover streets on a summer weekday afternoon.

Another way to detect the presence of high vehicle density, and to characterize the ambience of a route, is by measuring the sound volume. Sound in decibels is plotted versus distance in Figure 11(b) for a ride along the ground-truth route. The sound volume peaks near 80dB when the route passes through the main intersections of town where the automobile traffic is more prevalent.

3.1.2 Interpreting Cyclist Experience. In Figures 8(b)–11(b), we present a large amount of raw data and first-level inferences. In this section, we introduce two example metrics to help cyclists and other system users understand and make use of the types of data that a BikeNet system provides. The metrics



(a) CO₂ level along the ground-truth route. Large spikes as the cyclist passes through the center of town at rush hour. (b) Sound level along the ground-truth route. Even small town traffic exceeds the long-term health threshold (70dB).



(c) Digesting the data: health and performance metrics to visualize a cyclist's experience.

Fig. 11.

are weighted combinations of various sensor data types. In the metrics introduced shortly (health and performance), we constrain ourselves to incorporating sensors for which we collect data in our prototype BikeNet implementation, though there are other appropriate sensors that might reasonably be added. The weights (e.g., \mathbf{a} in the expression for Health that follows) comprise two subweights: the user-defined preference/importance and the normalizing factor for each element. The user-defined preferences reflect relative personal sensitivities (e.g., a cyclist with asthma might weight the CO₂ higher) to the elements composing a given index score. With the second subweight, we normalize each element (e.g., CO₂) according to its maximum dynamic range measured along routes about which sensed data has been collected so far. At the backend, or on a user's local display, index scores can be plotted versus distance to see the variations across the route to identify critical/interesting sections. For example, the plot of health in Figure 11(c) shows a large dip in the health index just before 1km where the CO₂ spikes (refer to Figure 11(a)). Secondly, users can compare the average index value among different routes at different times to identify the most favorable routes for a given aim (e.g., joy-riding, exercise). These index values can be mapped to colors and routes can be visualized as a color-coded playlist. As the number and coverage of route segments are built

up, a lookup service that returns the most healthy route at the desired time between two endpoints becomes possible. By sharing index values for routes of interest, and the user-defined preference weights, cyclists are able to learn from each other about where the good cycling is.

Health. Air pollution and its effect on public health is of great interest in many urban communities. In Austria, France, and Switzerland, by measuring particulates specifically from motorized traffic the effect of air pollution on public health is estimated to account for >20,000 adult deaths, more than 290,000 episodes of bronchitis in children, and more than 500,000 asthma attacks each year [Kunzli et al. 2000]. Noise pollution is also a factor in urban areas. According to the Environmental Protection Agency’s Office of Noise and Abatement Control in order to protect from hearing loss, one should not be exposed to more than 70dB for an extended period of time. Meanwhile the average city traffic is 85dB and in larger cities like New York, the noise level often exceeds 90db. 87% of America’s city dwellers are exposed to noise so loud it has the potential to degrade hearing capacity over time [Orlando 2007]. Even in the small town of Hanover, NH, USA (see Figure 11(b)) the noise level is often above 70dB on the main streets at certain times. BikeNet sensing supports not only communal pollution mapping, but on a more personal level it supports the categorization of cycling routes according to their potential impact on a cyclist’s health. We define a *health index* that combines data that indicate safety, noise, and air pollution (either directly or through inference) as follows.

$$Health = 1.0 - a_1 * CarDensity - a_2 * CO_2Level - a_3 * SoundLevel$$

From the raw values obtained from the MetalDetector we infer the density of cars along the route, along with raw values of CO₂Sensor and sound levels from the SoundSensor, we derive values for the health index of routes that a cyclist travels. A higher CO₂ level and derived car density imply there are more cars near the cyclist, creating an unpleasant experience due to exhaust, noise, and increased danger, driving the health index down. Similarly an increase in noise level indicates more traffic, people, wind, shouting, etc., reducing the health index. Figure 11(c) shows the details versus distance of how the characteristics of the route affect the cyclist, highlighting areas that should be avoided on future rides. We use an equal user-defined preference weight of $\frac{1}{3}$ for each of the three elements that are included in the score. The dynamic ranges (i.e., measured difference between max and min) for car density, CO₂ level, and sound level are 12, 100, and 70, respectively. Therefore, we use $a_1 = \frac{1}{3} * \frac{1}{12}$, $a_2 = \frac{1}{3} * \frac{1}{100}$, and $a_3 = \frac{1}{3} * \frac{1}{70}$. The average value of the health index over the entire route is 0.746 and the standard deviation is 0.096. As previously mentioned, this average value can be used to help a user rank his routes according to his own preferences and also to share with his peers.

Performance/Fitness. Some cyclists’ primary purpose in riding is for exercise or for competition. For these riders we calculate a *performance index*, using the values obtained from the WheelSensor, PedalSensor, and TiltSensor. We compute a unitless measure of performance using the following equation.

$$Perf. = b_1 * HillAngle + b_2 * WheelSpeed / PedalSpeed + b_3 * Distance.$$

When HillAngle is positive the performance index goes up; when it is negative the index goes down. When the wheel/pedal ratio is high this indicates the bike is in a higher gear (the wheel goes further with fewer pedal turns) and the index increases. The further a rider travels (larger *Distance*) the higher the performance index. The plot in Figure 11(c) shows the performance of a cyclist traveling the ground-truth route. The weights are set according to the same rationale as for the health expression, with values: $b_1 = \frac{1}{3} * \frac{1}{13}$, $b_2 = \frac{1}{3} * \frac{1}{80}$, and $b_3 = \frac{1}{3} * \frac{1}{5}$. With these user weights, the average value of the performance index over the entire route is 0.253 with a standard deviation of 0.094.

User-defined Metrics. Part of the appeal of the BikeNet system to the casual cyclist, is the ability to have fun defining new ways to visualize the cycling experience. Therefore, beyond health and performance, a cyclist is free to define her own metrics or index scores depending on how she likes to categorize her rides. One such frivolous index might only have one weighted component, such as the following definition of the *enjoyment* metric, adopted by many of the experiment participants.

$$Enjoyment = 1.0 - c_1 * HillAngle$$

When the HillAngle is positive cycling is more difficult, reducing the enjoyment metric. When the HillAngle is negative the cyclist can coast, resulting in an increased enjoyment index. Figure 11(c) shows a plot of the enjoyment index for a given rider of the course of the ground-truth route. Generally for these subjective types of index scores, the selection of weighting factors is cyclist dependent since each person has different levels of like/dislike/tolerance for annoyances, pollutants, etc. The weighting factor chosen here is $c_1 = 1 * 0.03$, indicating that the cyclist dislikes steep rides (a negative weight would indicate the opposite). Accordingly, the enjoyment index goes down during the steep uphill parts of the ride. The average joy of the route is 0.510 with a standard deviation of 0.063.

3.2 Social Cycling

There is a strong social element to bicycling, with many cyclists often riding in groups. If all of the cyclists have BikeNet sensors then we would expect certain types of sensors to return very similar readings for all the cyclists in the group. This correlation amongst readings facilitates a number of backend data sharing possibilities, including: automatic calibration checks of sensors, anomaly detection and noise reduction in sample data using shared information, and sharing of sensor data with those who are missing sensors (e.g., the more expensive ones.)

We examine the correlation between the sensor data in one of the multicycle experiments to explore this facet of BikeNet. In this experiment, the cyclists travel in tandem and usually remain within about 10 meters of each other though there are occasional separations of as much as 50 meters. The route taken by the cyclists is the same as shown in Figure 8(a). We compare results

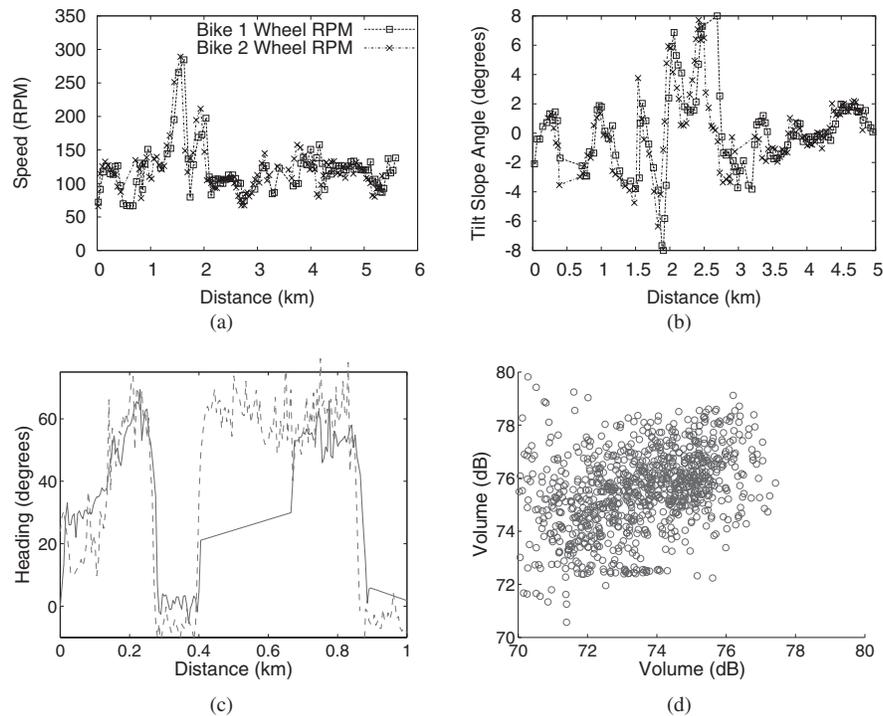


Fig. 12. Comparison of data sampled from two bicycles traveling in tandem with a separation of about 10m. The speed profiles (a), tilt slope profiles (b), and compass heading profiles (c) are highly correlated as the cyclists traverse the same terrain with a small time lag. The distance-aligned sound data shown in (d) is less correlated between the two bikes due to the highly temporal nature of sound events. However, gross statistics such as average volume are still fairly well correlated.

from a representative two bicycles in the following plots. Figure 12(a) shows the speed profile of two bikes, demonstrating a strong correlation in speed, albeit with a small lag, throughout the experiment as expected given the bikes traverse the same terrain. Figure 12(b) shows the tilt profile of the same bikes, again demonstrating a strong correlation in tilt. Figure 12(c) shows that compass readings from the bikes are also strongly correlated. Sharing data among members of a cycling group allows the system to “patch” the trace the compass that stops reporting data between 0.4 and 0.65km with data from another in its group.

Figure 12(d) on the other hand shows that the sound sensors on the two bikes are only somewhat correlated (perfect correlation would be reflected in all points along a 45 degree line, i.e., $y = x$). This graph is computed from the sound data from both bikes, interpolated so that the samples are aligned at the exact same distance along the path. Thus we are comparing the sounds at the same location (but not necessarily at the same time.) The weak correlation is due to several factors including multipath reflections of sound, wind noise variations over time, one or the other riders yelling comments, inverse square fading of volume with distance, cars passing each cyclist at different times,

and phase differences from small differences in each bike's path. While specific event signatures may be distorted or lost through intergroup sharing, certainly even without further processing the correlation between bicycles is good enough that the average volume is sharable, for example, for use in route noise mapping.

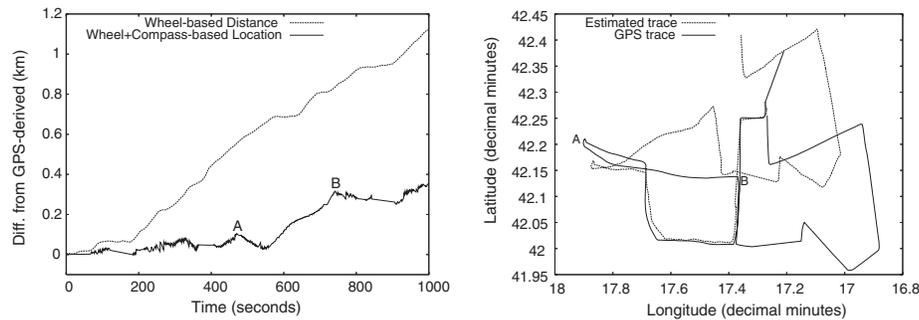
Some sensor data exhibits little or no correlation at all. The lateral tilt of each bike is dependent on the way each cyclist rides and varies a lot from bike to bike. GSR is tied to the stress response and is mostly unique for each cyclist. Gear ratio is loosely correlated depending on several factors including the strength and condition of the cyclists, the gear ratios available on each bike, the wheel size, and the choice of each cyclist as to when he pedals and when he coasts. Pedal speed has a similar loose correlation for the similar reasons. CO₂ data is highly correlated between riders in a flock since the exhaust of cars becomes well mixed by the turbulence created by the cars motion. GPS data is strongly correlated between cyclists depending on how close together they ride.

Thus, the sharing of some types of sensor data is feasible and likely to be quite useful. In the course of this study we learn that the sample rates for several of our sensors reduce the potential usefulness of shared data. For example, we want to try correlating GSR to pedal RPMs to see if the periodicity of the pedaling shows up due to the shifting pressure the cyclist put on the GSR electrodes. However the 1Hz sample rate of the GSR output is too slow to capture this phenomenon if it exists because people often pedal at somewhere near 1Hz. Coming up with methodologies for examining the correlation of data that may be skewed in time or location, and which may have only a fuzzy sort of correlation, is an open challenge. For example, correlating CO₂ concentration with car detections from the magnetometer is complicated by the fact that the CO₂ sensor reacts slowly to changes in CO₂ concentration, taking about a minute to respond to changes.

3.3 Services Performance

3.3.1 Cyclist Localization. In the BikeNet system, each data record is tagged with location (and time) to provide context. Section 2.2.2 describes the BikeNet localization design. While GPS is a natural solution for the implementation of the SyncSprinkler role, other lower cost and lower energy-consuming implementations warrant an investigation given that every BAN may not be equipped with a GPS receiver. Further, GPS is limited in its application since it must have vision of at least three satellites. In particular, it is not always functional in some common bicycling environments (e.g., among tall buildings in a city, under dense overgrowth (tree canopy)). In the following, we present the accuracy of a WheelSensor in measuring total distance traveled by the cyclist. We also present the accuracy of using the combination of WheelSensor and CompassSensor for measuring localization (i.e., dead reckoning). In both cases, a known starting location is required. Both results are shown with respect to the GPS solution.

We present the absolute error of these localization alternatives in Figure 13(a). As seen from the curve labeled "Wheel", the deviation from the



(a) Deviation from GPS of measured distance by (b) Map of estimated absolute location using WheelSensor; deviation from GPS of absolute location WheelSensor+CompassSensor using dead reckoning from a known starting location.

Fig. 13.

GPS-derived distance grows continuously over the course of the ride. A closer look at the data shows that the WheelSensor is consistently underestimating the distance traveled. The WheelSensor technique requires that the circumference of the tire be known; we initialized the WheelSensor with the value we measured with a tape measure. We conjecture that over the course of a ride small errors in this measurement and changes in circumference from compression of the tire due to the weight of the cyclist accumulate to cause the continuously increasing error. As seen from the curve labeled “Wheel+Compass”, the deviation in absolute localization from GPS is rather small until about the half-way point of the route (at label A in Figures 13(a) and 13(b)). However, the performance between label A and label B shows one limitation of this approach, without further corrective measures. In this interval, the cyclist is riding next to a large truck, causing the inferred track to deflect to the left. When the cyclist reaches the intersection (label B), it turns right and leaves the truck. At this point, the track again closely follows the GPS track, as seen from the minimal error growth from label B to the end of the trace. However, due to the deflection, the “Wheel+Compass” track is translated north and west with respect to the GPS track (see Figure 13(b)).

While the WheelMote may be made accurate enough if the weight of the rider is taken into consideration to correct for error introduced by tire compression, from Figure 13 it is evident that exclusive use of the WheelSensor+CompassSensor for localization is likely not accurate enough to replace GPS. However, it might still be used as a complementary solution to interpolate between occasionally received GPS updates. This is especially useful in scenarios where the BAN’s GPS receiver is integrated with the cyclist’s cell phone (e.g., Nokia N95), and energy consumption is a concern. For example, our initial experience [Miluzzo et al. 2008a] with the Nokia N95 is in line with that of other researchers [Gaonkar et al. 2008] and those in the blogosphere where the lifetime of the standard battery shrinks from approximately 3 days to less than 7 hours when the GPS receiver is left on continuously. We therefore duty cycle the GPS, acquiring a signal for 150 seconds and sleeping for 850 seconds.

Applying this strategy to the location error curve shown in Figure 13(a) reduces the average error from 125.8m (no GPS correction) to 64.1m (10min GPS duty cycle). Clearly, the GPS duty cycle could be shortened to improve accuracy, but at the cost of increased energy consumption by the GPS hardware.

3.3.2 Data Uploading and Muling. Throughout the experiments, sensed data is always stored in the local Flash memory of the Tmote Invent. Ultimately, this data must be transferred to the backend data repository. Depending on bicycle and cyclist mobility this transfer may happen directly from BAN to SAP using the upload protocol (refer to Section 2.2.3), or may happen indirectly via the muling protocol (refer to Section 2.2.3). Each BAN member knows the identifiers of the other mobile sensing platforms in its BAN as a result of the role assignment protocol (see Section 2.2.1), and will only mule data for mobile sensors not in its BAN. The exception to this rule is that the PersonalNode does mule data on behalf of its associated BAN members. There is no intra-BAN coordinator for either upload or muling exchanges; BAN members transfer their data independently and contend using CSMA for the wireless channel. There are two main scenarios to consider: (i) the bicycle enters in range of a SAP in which case the mobile sensors in the BAN upload their data to the SAP directly or (ii) the bicycle is traveling out of range of the SAP, and must rely on probabilistic mobility of other people or BANs to mule the data to a SAP.

Since the Stop-and-Wait ARQ reliable transfer mechanism used in the muling and upload protocols is well known we omit any evaluation of this mechanism per se. Rather, we aim to characterize the opportunistic sensor networking environment provided by our initial prototype implementation of BikeNet. In Figures 15(a), 15(b), and 15(c), we show results from a multibicycle experiment where each cyclist follows a prescribed path; the paths intersect giving rise to interbicycle muling opportunities. The paths followed by each of the four bicycles used in the experiment follow the perimeter of a central grassy area called “the Green” commonly used for recreation at Dartmouth College. The area of the Green is approximately 150m by 100m. Two cyclists ride clockwise around the Green and the other two cyclists ride counter clockwise. The transmission range of the Tmote Invent reaches fewer than 50m on average so the connections among the cyclists are intermittent. However, there are data exchange opportunities when the cyclists who are moving in opposite directions pass by each other. After ten minutes of circling around the Green, each cyclist leaves the square in turn at fifteen minute intervals and parks the bicycle within the radio range of the SAP installed at the Sensor Systems Lab in the Computer Science building which is 250m away from the northeast corner of the Green (i.e., out of range).

Characterizing interbicycle communication. To understand the expected data transfer opportunity during the bicycle rendezvous periods, we run a controlled experiment using two bicycles where Bike A transfers CompassSensor data as fast as possible to Bike B. We conduct ten independent trials where Bike A and Bike B travel towards each other starting at opposite ends of the path. We record the average speed of each bike and then average each bike’s speed

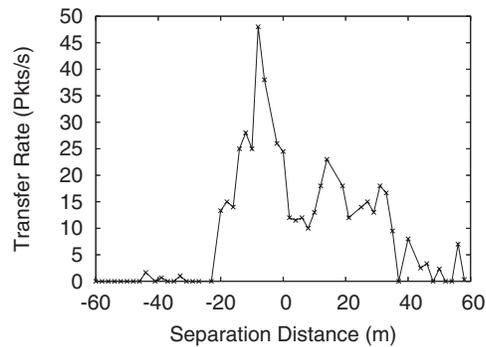


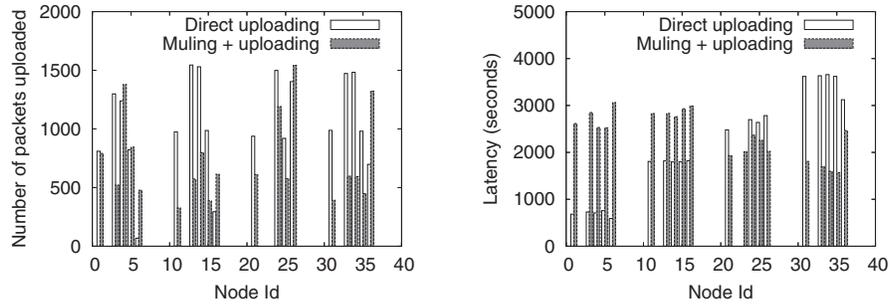
Fig. 14. Characterizing interbicycle communication opportunity during a mobile rendezvous. Transfer throughput exhibits an asymmetric pattern likely due to radio absorption by the cyclists' bodies.

over ten runs. The distance separating the bikes is reported as (position of Bike A - position of Bike B), where the position is measured using markings along the path. A profile of the packet transfer rate averaged over the ten bike-to-bike rendezvous is shown in Figure 14. The average speed of the reference bike and other bicycle are recorded as 3.45 meters per second and 2.55 meters per second, respectively. The average packet transfer rate recorded for an average contact time of 10.75 seconds is 19.8 packets per second. This results in an average reliable transfer of 556 bytes between the bikes.

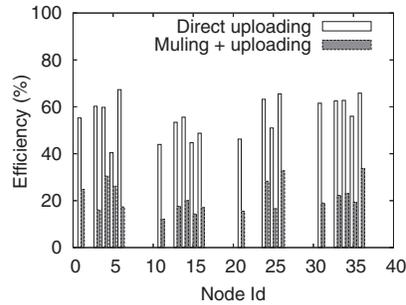
The throughput profile in Figure 14 exhibits an asymmetry whereby when the bicycles are approaching each other (positive separation distance) the throughput grows gradually to a peak of 50 pkts/sec, while shortly after the bicycles pass each other (zero distance) the throughput falls precipitously to zero. While the exact signature shown in the figure depends on the body size of the cyclist and the exact trajectories of the cyclists, the transfer rate trace likely describes qualitatively all people-centric rendezvous. In our particular experiment the sensor is attached to the handlebar on the front of the bike. In follow-on work [Miluzzo et al. 2008b], we show that the drop is due to the additional attenuation presented by the body of the cyclist rather than other issues (e.g., we discounted Doppler shift as a factor).

Muling performance. In Figure 15(a), we show the number of data packets directly uploaded and the number of data packets muled to the SAP. Data records have unique identifiers to allow filtering of duplicates between muled and directly uploaded packets. The x-axis of the plot shows the identifiers of the Tmote Invents in each BAN. The Tmote Invents with identifiers 1 through 6 belong to BAN-1, 11 through 16 belong to BAN-2, 21 through 26 belong to BAN-3, and 31 through 36 belong to BAN-4. This x-axis is the same for Figures 15(b) and 15(c). Figure 15(a) shows that the rendezvous intervals in the experiment support a substantial amount of muling exchange. Overall, more data packets are directly uploaded than muled, but a considerable amount of data are muled before uploading (e.g., almost half of BAN-1's data).

Figure 15(b) presents the latency of direct uploading and muling. We measure the latency as the time difference from the time the sensor data packet is



(a) Data uploaded directly and after muling for the twenty nodes in the four BANs. (b) Delivery delay comparison between direct uploading and muling + uploading.



(c) Delivery efficiency comparison between direct uploading and muling + uploading.

Fig. 15.

generated to the time the packet is uploaded to a SAP, either by the originator of the packet or by a mule. The result clearly shows the benefit of muling. The data from BAN-3 and BAN-4 are muled by the Tmote Invents on BAN-1 and BAN-2 which enter within radio range of the SAP earlier than BAN-3 and BAN-4. As a result, the data from BAN-3 and BAN-4 are delivered earlier with muling than with direct uploading. In particular, data from BAN-4 experience an average of 1500 seconds less delay with muling than with direct uploading.

Muling implies a performance penalty due to the additional radio transmissions that are required on the path from origin to mule(s) and mule(s) to SAP. Figure 15(c) illustrates the transfer efficiency for both muling and direct uploading as a function of the packets transferred. Here we define efficiency as the ratio of data bytes transferred to the total bytes sent (including data packet replicas and retransmissions). A higher efficiency for a given transferred packet reflects a number of possible factors, including a higher-quality radio link and less congestion between sender and receiver, both of which lead to fewer retransmissions. The data origin always stores a copy in its local Flash hoping to upload the data directly to a SAP. In our experiments, the muling replication degree is one, meaning we allow the originating Tmote Invent to transfer only a single copy of the data to another mobile sensor via the muling protocol. Further, we do not allow multihop muling; only the originating node may replicate data.

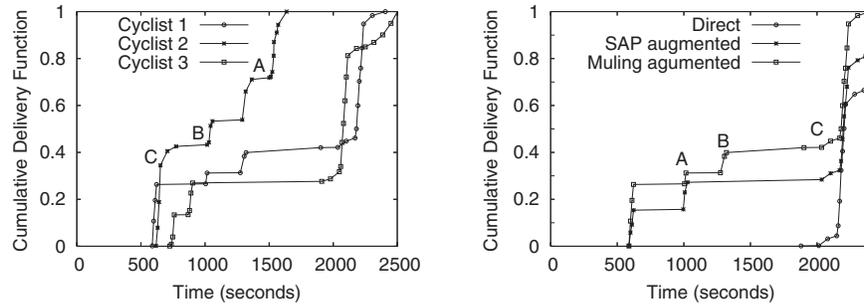
Figure 15(c) shows that the muling efficiency is less than uploading efficiency, as expected. While the uploading efficiency ranges from 40%–68%, the muling efficiency ranges from 12%–33%. The difference between the muling efficiency and the uploading efficiency represents the cost for improving the data delivery delay. The effect of replication degree on performance is multifaceted. On the one hand, replication improves the probability for delivery and tends to give improved delivery delay performance, but it also lowers efficiency and limits storage space for unique data. When determining the replication degree, one must carefully consider the system requirements in terms of average delivery delay, and the costs involved in providing such delay performance. Data replication is a problem in many delay-tolerant networks. We defer to the wealth of research studies on the topic (e.g., Ip et al. [2007], Wang et al. [2005], Chen et al. [2006]). In the BikeNet system, the bulk of collected data is highly delay tolerant since the cyclist herself is the primary consumer of the data (i.e., it is fine if the data arrives at the same time as the rider). Similarly, data via direct upload is likely sufficient for secondary consumers of the cyclist’s data, such as cycling interest groups and community action groups. However, BikeNet provides for the real-time upload of high-priority data (e.g., emergency signals, or periodic location) via the Nokia N80 data connection.

3.4 Town-Scale System Scenario

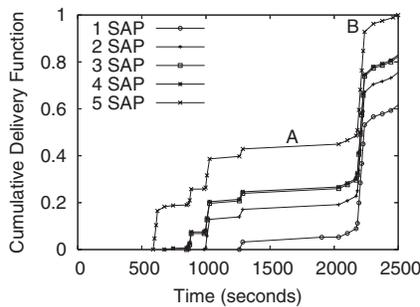
We build five sensor bikes and implement a small-scale BikeNet testbed with seven static SAPs at a number of points across the Dartmouth College campus and in the town of Hanover to validate and evaluate an operational BikeNet system. In what follows, we present results from data collected by a group of three cyclists on the morning of November 20, 2006. We have collected a significant amount of data from over 50 different BikeNet experiments starting in summer 2006 but here only present data from a single-shot experiment with the three cyclists. The three cyclists’ routes and the times they started their rides are preplanned. Cyclists 1 and 2 live near each other and ride much of the way toward campus from the town together.

Before getting to the campus they rendezvous with cyclist 3 before cyclists 1 and 3 depart toward the library while cyclist 2 heads to the Computer Science building. The longest journey time for a cyclist is 40 minutes. The results presented in this section provide insights into how live sensor data collected by each of the bikes is either muled or directly uploaded to a passing SAP, in an opportunistic manner.

We first consider the time taken for each cyclist to upload its data into the backend data repository via direct upload to a SAP and present the results in Figure 16(a). Note that in this case the data uploaded to a SAP by a cyclist’s BAN includes data originating in the local BAN and any data muled on behalf of other BANs. Figure 16(a) shows the cumulative delivery function versus the trip time. We define *cumulative delivery function* as the cumulative fraction as the trip proceeds of the total data packets delivered to the backend. For example, at 520 seconds into the experiment cyclist 2 has delivered 40% of its data including any data it may mule on behalf of cyclists 1 and 3. The initial lack of any data



(a) Data delivery traces for three cyclists using direct upload from BAN to SAP. (b) Comparison of different delivery methods for single cyclist's data.



(c) Evaluation of delivery delay versus number of SAP encountered en route.

Fig. 16.

delivery evident in the plots is due to the initial absence of SAPs along the route from the homes of cyclists 1, 2, and 3. Cyclists 1 and 2 encounter a SAP along Main Street in Hanover at approximately 520 seconds into their ride. Even though only a modest number of SAPs are deployed we can see from the plot that all bikes are capable of delivering sizable fractions of their data before their trips end. For example, cyclists 1 and 2 deliver approximately 50% and 40%, respectively, of their data to the backend repository before the half-way point of the trip time. From Figure 16(a), we observe that data is transferred between bikes and SAPs in quick bursts. The amount of data transferred is strongly influenced by the short contact times which are a product of the short-range radios used in the BikeNet experiments.

Figure 16(b) shows the delivery of sensor data that is generated by only a single cyclist—in this case cyclist 1. We consider three scenarios of possible transfer between cyclist 1 and the data repository: *Direct*, which is where cyclist 1 keeps all its data, does not replicate, and only uploads to the SAP at the Computer Science building (the cyclist's destination); *SAP augmented*, which is where cyclist 1 opportunistically transfers data to SAPs it encounters along the way with no help from mules; and finally, *Muling/SAP augmented*, which exploits muling and opportunistic use of SAPs to transfer cyclist 1's data to the repository. Figure 16(b) shows the performance of these three types of

communication across the trip time. From the plot we can observe the direct benefit of muling: cyclist 1 is disconnected for approximately 15 minutes between points A and C in the plot of the cumulative delivery function, but at the intermediate point B the delivery of cyclist 1's packets continues.

Next, we evaluate the impact of the incremental addition of SAPs to the system on the average delay of data delivered from the three cyclists to the backend repository. As the number of SAPs increases from one to five, we plot the cumulative delivery function versus time. We run five trials at each SAP level, each with the three cyclists riding prescribed routes from their respective homes to the Computer Science department. Figure 16(c) shows the cumulative delivery functions for each of the five cases, truncated at the time when 100% of data is delivered by all cyclists for the case of five SAPs. The plot shows that at the time all data is uploaded in the five SAP case, only 78% and 70% of data is uploaded in the three SAPs and one SAP cases, respectively. When the cyclists return to the Computer Science building they become stationary and upload their remaining data. This is reflected in the steep step in all the curves at point B in Figure 16(c), representing a large delivery of the remaining data from the bicycles to the SAP. In contrast, the flat portions of the curves, for example, in the area of point A, represent periods when cyclists 2 and 3 are disconnected from the network with no other mobile sensors acting to mule data to the SAPs. From Figure 16(c), we can also observe that the addition of a new SAP yields a nonuniform improvement in the data delivery performance. The impact of adding SAPs to the system on the data delivery delay is highly dependent on many factors, including the SAP deployment density and the location of the SAPs in relation to routes frequented by the cyclists.

4. RELATED WORK

A number of companies (e.g., Garmin Edge 305 and Garmin Emap, Garmin International [2007] and CicloSport HAC5 [2007]) have begun to offer products that integrate data from multiple sensors on a single user display, including biometric, advanced cyclo-performance and GPS location data, showing the interest in quantifying performance. These range from rather limited \$40 devices to very capable \$500 devices [CicloSport HAC5 2007]. Products with a slightly different focus offer integrated hardware and software solutions (e.g., Garmin Edge 305 and Garmin Emap, Garmin International [2007]) to help cyclists with preride route planning, and in-ride navigation cues via predownloaded maps combined with real-time GPS data. Others (e.g., tri [2007]) offer offline planning software packaged with an online GPS tracking service available via a select set of cellular providers. BikeNet goes beyond any of these commercial offerings by adding environmental sensors to give context to performance. Further, the BikeNet system incorporates a dual-mode wireless networking approach for data delivery to a backend analysis and visualization tier, providing usable and understandable information to users. By providing this dual-mode architecture, we aim to support both delay-tolerant query/delivery as well as expedited query/delivery models. BikeNet supports delay-tolerant interaction for the sake of the biker doing postride analysis of his entire ride, and real time

interaction for more time-sensitive operations by community consumers of the data.

The wearable computing and personal area networking fields have produced numerous examples [Feldmeier and Paradiso 2004; DeVaul et al. 2003; Zimmerman 1996; Gyselinckx et al. 2005] of wireless networks that operate on and near the human body and interact with the wearer's surroundings or other people's devices. There has also been work in delay-tolerant networking [Scott et al. 2006; Kansal et al. 2004; Wang and Wu 2006] to improve data transfer in networks that are often disconnected as BikeNet is. BikeNet synthesizes these ideas, using opportunistic rendezvous among personal computing devices (e.g., cell phones), embedded bicycle sensor networks, and sensor access points. BikeNet adds an implementation of backend data storage, analysis, and visualization services, to translate raw sensor data streams into meaningful information about personal health and performance and community health. BikeNet provides a complete opportunistic sensing system targeted at a cyclist experience application, and adds a new dimension to delay-tolerant networking by investigating both human-to-bike and bike-to-bike data transfer.

Mobile sensing systems have been proposed in other application contexts. Zebranet [Liu et al. 2004] monitors zebras wearing sensor collars using a mobile jeep-mounted radio gateway. The Cartel project [Hull et al. 2006] provides a mobile communications infrastructure based on car-mounted communication platforms exploiting open WiFi access points in a city, but unlike BikeNet it does not integrate either a sensing or sensed data analysis component. SATIRE [Ganti et al. 2006] presents a software architecture for smart clothing, similar to the BikeNet three tier architecture. The MetroSense Project [MetroSense Project 2007] proposed a mobile sensing system for skier-based sensing [Eisenman and Campbell 2006]. More recently, an airborne wireless sensor network composed of Micro-Aerial Vehicles (MAVs) has been implemented [Allred et al. 2007] to enable low-cost-high granularity atmospheric sensing of toxic plume behavior and storm dynamics, and provide a three-dimensional view for monitoring ecological systems. BikeNet differs in its application scope and its inclusion of bicycles as mobile sensing platforms for personal and environmental sensing. Further, we have developed BikeView, our community oriented Web-based portal for visualization and sharing of cycling health and performance information.

Tightly bound to the domain of people-centric sensing is the issue of privacy [Johnson et al. 2007]. Though the bulk of sensed data collected by the BikeNet system may seem innocuous (especially since it may be later shared with the community anyway), concerns about personal performance, and especially location tracking, must be addressed. While backend data sharing is regulated by the user, in-network techniques such as muling along with the possibility of wireless snooping [Saponas et al. 2007] mandate some privacy solution. While we do not integrate privacy protection into our prototype, we note the work of others can likely be leveraged. We conjecture that concepts such as virtual walls [Kapadia et al. 2007] can be used when deciding what information to reveal in real time via the backend or to peer BANs. Lightweight encryption (e.g.,

TinySec [Karlof et al. 2004], MiniSec [Luk et al. 2007]) may be used for intra-BAN communications and to protect against data prying on the part of mules.

5. CONCLUSION

In this article, we have presented the detailed design, implementation, and evaluation of the BikeNet mobile sensing system, adding to the growing body of work exploring opportunistic sensor networking techniques. BikeNet represents the first comprehensive mobile sensing system quantifying the cyclist experience. BikeNet provides for the collection and analysis of personal performance and communal environmental sampling. BikeNet supports two modes of operation in support of delay-tolerant and real-time sensing, and collected data can be presented both locally to the cyclist or to others via backend services. The BikeView portal concept promotes social networking among cyclists, and the broader community. Initial results are encouraging and demonstrate some of the value that mobile wireless sensor networks can bring to our lives, including how we are impacted by our environment and how we can regulate our activity patterns to improve our quality of life. While our current experiments have concentrated on sensing for the cyclist and bicycle-mounted sensors we conjecture that the BikeNet system could be implemented on other vehicles such as cars with little modification to the software system.

ACKNOWLEDGMENTS

The authors thank X. Zheng and H. Lu for their help with data collection, P. Boda and C.-Y. Wan for their support of this project, and P. Corke for his helpful suggestions.

REFERENCES

- ALLRED, J., HASAN, A. B., PANICHSAKUL, S., PISANO, W., GRAY, P., HUANG, J., HAN, R., LAWRENCE, D., AND MOHSENI, K. 2007. Sensorflock: An airborne wireless sensor network of micro-air vehicles. In *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems*. ACM.
- BEATING HEART BLOG. 2007. <http://turbulence.org/Works/beatheheart/blog/>.
- BIKEVIEW - THE BIKE NET WEB PORTAL. 2007. <http://bikenet.cs.dartmouth.edu>.
- CAMPBELL, A. T., EISENMAN, S. B., LANE, N. D., MILUZZO, E., AND PETERSON, R. A. 2006. People-Centric urban sensing. In *Proceedings of the 2nd Annual International Conference on Wireless Internet*. IEEE.
- CARUSO, M. J. 2000. Applications of magnetic sensors for low cost compass systems. In *Proceedings of the IEEE Position Location and Navigation Symposium*. IEEE. 177–184.
- CHEN, L., YU, C., SUN, T., CHEN, Y., AND CHU, H. 2006. A hybrid routing approach for opportunistic networks. In *Proceedings of the SIGCOMM Workshop on Challenged Networks*. ACM, 213–220.
- CICLOSPORT HAC5. 2007. <http://www.ciclosportusa.com/hac.html>.
- CITYSENSE. 2007. <http://www.citysense.net/>.
- DE ROSA, C., GAUDER, B., LIMES, R., AND CELLENTANI, D. 2007. Sharing, privacy and trust in our networked world. OCLC rep. 170923242.
- DEVAILLANT, R., SUNG, M., GIPS, J., AND PENTLAND, A. 2003. Mithril 2003: Applications and architecture. In *Proceedings of the 7th International Symposium on Wearable Computers*. 4.
- EISENMAN, S. B. 2008. People-Centric mobile sensing networks. Ph.D. dissertation, Columbia University.
- EISENMAN, S. B. AND CAMPBELL, A. T. 2006. Skiscape sensing. In *Proceedings of the 4th ACM Conference on Embedded Networked Sensor Systems*. 401–402.

- EISENMAN, S. B., MILUZZO, E., LANE, N. D., PETERSON, R. A., AHN, G.-S., AND CAMPBELL, A. T. 2007. The bikenet mobile sensing system for cyclist experience mapping. In *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems*. 87–101.
- FELDMEIER, M. AND PARADISO, J. A. 2004. Giveaway wireless sensors for large-group interaction. In *Extended Abstracts on Human Factors in Computing Systems*. 1291–1292.
- GANTI, R., JAYACHANDRAN, P., ABDELZAHER, T., AND STANKOVIC, J. 2006. Satire: A software architecture for smart attire. In *Proceedings of the 4th International Conference on Mobile Systems, Applications, and Services*. 110–123.
- GAONKAR, S., LI, J., CHOUDHURY, R. R., COX, L., AND SCHMIDT, A. 2008. Micro-Blog: Sharing and querying content through mobile phones and social participation. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services*.
- GARMIN EDGE 305 AND GARMIN EMAP, GARMIN INTERNATIONAL. 2007. <http://www.garmin.com>.
- GUPTA, R. AND SOMANI, A. K. 2004. Pricing strategy for incentivizing selfish nodes to share resources in peer-to-peer (p2p) networks. In *Proceedings of the 12th International Conference on Networks (ICON'04)*. IEEE, 624–629.
- GYSELINCKX, B., HOOFF, C. V., RYCKAERT, J., YAZICIOGLU, R., FIORINI, P., AND LEONOV, V. 2005. Human++: Autonomous wireless sensors for body area networks. In *Proceedings of the 27th Conference on Custom Integrated Circuits*. 13–19.
- HULL, B., BYCHKOVSKY, V., ZHANG, Y., CHEN, K., GORACZKO, M., MIU, A., SHIH, E., BALAKRISHNAN, H., AND MADDEN, S. 2006. Cartel: A distributed mobile sensor computing system. In *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*. 125–138.
- IP, Y.-K., LAU, W.-C., AND YUE, O.-C. 2007. Forwarding and replication strategies for dtn with resource constraints. In *Proceedings of the IEEE 65th Vehicular Technology Conference*. 1260–1264.
- JOHNSON, P., KAPADIA, A., KOTZ, D., AND TRIANOPOULOS, N. 2007. People-centric urban sensing: Security challenges for the new paradigm. Dartmouth Tech. rep. TR2007-586.
- KANSAL, A., SOMASUNDARA, A. A., JEA, D. D., SRIVASTAVA, M. B., AND ESTRIN, D. 2004. Intelligent fluid infrastructure for embedded networks. In *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services*. 111–124.
- KAPADIA, A., HENDERSON, T., FIELDING, J. J., AND KOTZ, D. 2007. Virtual walls: Protecting digital privacy in pervasive environments. In *Proceedings of the the 5th International Conference on Pervasive Computing*. 169–179.
- KARLOF, C., SASTRY, N., AND WAGNER, D. 2004. Tinysec: A link layer security architecture for wireless sensor networks. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*. 162–175.
- KUNZLI, N., KAISER, R., MEDINA, S., STUDNICKA, M., CHANEL, O., ET AL. 2000. Public-health impact of outdoor and traffic-related air pollution: A European assessment. *The Lancet* 356, 9232, 795–801.
- LIU, T., SADLER, C., ZHANG, P., AND MARTONOSI, M. 2004. Implementing software on resource-constrained mobile sensors: Experiences with impala and zebranet. In *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services*. 259–269.
- LUK, M., MEZZOUR, G., PERRIG, A., AND GLIGOR, V. 2007. Minisec: A secure sensor network communication architecture. In *Proceedings of the 6th International Conference on Information Processing in Sensor Networks*. 479–488.
- METROSENSE PROJECT. 2007. <http://metrosense.cs.dartmouth.edu>.
- MILUZZO, E., LANE, N. D., FODOR, K., PETERSON, R., LU, H., MUSOLESI, M., EISENMAN, S. B., ZHENG, X., AND CAMPBELL, A. T. 2008a. Sensing meets mobile social networks: The design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys'08)*. ACM, New York, 337–350.
- MILUZZO, E., ZHENG, X., FODOR, K., AND CAMPBELL, A. T. 2008b. Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks. In *Proceedings of the 5th European Conference on Wireless Sensor Networks*. 479–488.
- MOTEIV TMOTE INVENT. 2007. <http://www.moteiv.com/>.
- NIKE+iPOD SPORTS KIT. 2007. <http://www.apple.com/ipod/nike/>.
- NOKIA N80 MOBILE DEVICE. 2007. <http://www.nokia.com/nseries/index.html>.
- ORLANDO, J. 2007. http://www.culturechange.org/issue19/vehicle_noise.htm.

- PAI, S., MEINGAST, M., ROOSTA, T., BERMUDEZ, S., WICKER, S., MULLIGAN, D. K., AND SASTRY, S. S. 2008. Confidentiality in sensor networks: Transactional information. *IEEE Secur. Privacy Mag.*
- REDDY, S., PARKER, A., HYMAN, J., BURKE, J., ESTRIN, D., AND HANSEN, M. 2007. Image browsing, processing, and clustering for participatory sensing: Lessons from a dietsense prototype. In *Proceedings of the 4th Workshop on Embedded Networked Sensors (EMNETS'07)*. 13–17.
- SAPONAS, T. S., LESTER, J., HARTUNG, C., AND KOHNO, T. 2007. <http://www.cs.washington.edu/research/systems/privacy.htm>.
- SCOTT, J., HUI, P., CROWCROFT, J., AND DIOT, C. 2006. Huggle: A networking architecture designed around mobile users. In *Proceedings of the 3rd Annual Conference on Wireless On-Demand Network Systems and Services*. 78–86.
- TINYOS. 2007. <http://www.tinyos.net/>.
- TRIMBLE. 2007. Adventure planner, trimble outdoors. <http://www.trimbleoutdoors.com/biking.aspx>.
- WANG, Y., JAIN, S., MARTONOSI, M., AND FALL, K. 2005. Erasure-Coding based routing for opportunistic networks. In *Proceedings of the ACM SIGCOMM Workshop on Delay-Tolerant Networking*. 229–236.
- WANG, Y. AND WU, H. 2006. Dft-msn: The delay fault tolerant mobile sensor network for pervasive information gathering. In *Proceedings of the 25th International Conference on Computer Communications*. 1–12.
- WOOD, A., STANKOVIC, J., VIRONE, G., SELAVO, L., HE, Z., CAO, Q., DOAN, T., WU, Y., FANG, L., AND STOLERU, R. 2008. Context-Aware wireless sensor networks for assisted-living and residential monitoring. In *IEEE Netw* (Special Issue).
- ZHANG, Y., LOU, W., LIU, W., AND FANG, Y. 2007. A secure incentive protocol for mobile ad hoc networks. *Wirel. Netw.* 13, 5, 569–582.
- ZIMMERMAN, T. 1996. Personal area networks: Near-Field intrabody communication. *IBM Syst. J.* 35, 3, 4.

Received May 2008; revised December 2008; accepted August 2009