Mobility Changes Everything in Low-Power Wireless Sensornets

Prabal Dutta and David Culler
Computer Science Division
University of California, Berkeley
Berkeley, California 94720
{prabal,culler}@cs.berkeley.edu

Abstract

The system and network architecture for static sensornets is largely solved today with many stable commercial solutions now available and standardization efforts underway at the IEEE, IETF, ISA, and within many industry groups. As a result, many researchers have begun to explore new domains like mobile sensor networks, or mobiscopes, since they enable new applications in the home and office, for health and safety, and in transportation and asset management. This paper argues that mobility invalidates many assumptions implicit in low-power static designs so the architecture for micropower mobiscopes is still very much an open research question. In this paper, we explore several mobile sensing applications, identify research challenges to their realization, and explore how emerging technologies and real-time motion data could help ease these challenges.

1 Introduction

Researchers recently outlined a vision for monitoring human spaces using *mobiscopes* [1]. A mobiscope, according to this work, "is a federation of distributed mobile sensors into a taskable sensing system that achieves high-density sampling coverage over a wide area through mobility." Mobiscopes extend the traditional sensor network model and introduce new challenges in data management, data integrity, privacy, and network system design. Because of these new challenges, the authors contend that researchers need new architectures and methodologies for designing future mobiscopes. The paper identifies two distinct classes of mobiscopes – *vehicular* and *handheld* – and presents their requirements, sampling problems, heterogeneity challenges, data privacy issues, networking challenges, and social implications.

This paper considers a third class of mobile sensor networks – *micropower mobiscopes* – or simply *μ-mobiscopes*. While vehicular mobiscopes can draw power from a vehicle's electrical system and handheld mobiscopes are recharged daily, μ-mobiscopes are more like low-power sensor networks in that they are embedded in the environment and cannot be recharged easily. μ-mobiscopes may be integrated into clothing and accessories, or attached to everyday objects in the home, office, factory, farmhouse, or vehicle. To be truly unobtrusive, however, μ-mobiscope batteries must be smaller than the AA batteries used in many prototypical sensor nodes today, and these smaller batteries must either last for a long time or be automatically recharged *in situ*.

Unlike many traditional sensor networks, the things that μ-mobiscopes monitor – people, animals, packages, and vehicles – exhibit unpredictable motion. Mobility invalidates many assumptions implicit in low-power static designs, so the techniques and protocols now in standardization for static networks are not well-suited to the needs of low-power, mobile systems. Mobility, for example, breaks link estimators and synchronous MAC protocols, and makes asynchronous neighbor discovery and routing fundamentally more challenging.

Realizing the vision of μ-mobiscopes requires advances in platform, link, network, transport, and storage layers of the system stack, and new thinking on programming models and security mechanisms. Although these problems are well-studied for static sensor networks, and they are clearly articulated for vehicular and handheld mobiscopes, the issues remain relatively unexplored for emerging μ-mobiscopes.

In this paper, we explore enabling technologies, example applications, and research challenges for future μ-mobiscopes and suggest that real-time knowledge of motion could be exploited to address some of the challenges that mobility itself raises. We propose a new metric called *Motion ObserVed* (MOV) that indicates node-level movement, discuss how it might be gathered on a nanopower energy budget, and explore how it could be used to address a range of sensor tasking and networking issues for μ-mobiscopes.

2 Enabling Technology

New motion sensing and energy harvesting technologies that are just emerging from research labs, or on the verge of commercial release, may help enable the transition from static to mobile sensor networks. This section explores promising technologies for μ-mobiscopes, including how the MOV metric could be implemented.

2.1 Ultra Low-Power Motion Detection

In many mobile sensing applications, awareness of a node's own movement is beneficial. In some cases, motion detection and capture is the whole *raison d'être*. In other cases, motion may be correlated with important external events, so detecting motion can trigger additional sensing, processing, or communications. For example, a sensor monitoring periodic limb movement disorder in an adult might capture high-speed motion data only after detecting a sudden jerk but ignore more benign motions. A gesture detection algorithm might process motion data in search of a "double tap" only after experiencing the shock of a single tap. Or, a loss deterrent system might start to beacon only after the sensor detects motion but otherwise keep its radio turned off to conserve energy.

Although motion detection (and capture) can improve the performance of these applications, the challenge comes in acquiring this motion data in an energy-efficient manner. Many mobile sensing applications use MEMS accelerometers to capture motion data. Although these sensors are improving quickly, and integrating increasingly sophisticated motion processing algorithms in hardware, they still draw 100 μA or more for modest sampling rates [2]. Piezo-based accelerometers are available as well, and although they can detect only the dynamic component of acceleration, they do offer an order of magnitude reduction in power draw [19]. Vibration switches built using small, conductive balls that repeatedly open and close a circuit in response to ambient vibration can be augmented with nanopower (sub-μA) circuitry to detect movement or time-integrate motion data [24]. Finally, vibration tabs are cantilevers constructed from piezo material, so they generate a voltage in response to motion, and this effect is magnified with a small proof mass on the end of the cantilever [20].

Going forward, we expect to see ultra low-power accelerometers with integrated electronics, programmable thresholds, onboard gesture recognition, and buffered digital outputs. Analog Devices, for example, has announced a 3-axis digital accelerometer that draws as little as 25 μA while sampling at 25 Hz, offers tap/double-tap, activity/inactivity and free-fall detection, and has a 32 sample output FIFO to ease the microcontroller processing burden [3].

2.2 Practical Wireless Power Transfer

A problem with many mobile systems today is that their batteries must be recharged or replaced periodically. Battery maintenance is a pure hassle and it is one reason why ubiquitous computers are not so ubiquitous today. Having to recharge a laptop, cell phone, wireless headset, digital camera, and music player is annoying enough. Adding even a few mobile sensors, let alone dozens or hundreds, would quickly make life unbearable.

Even many wireless power transfer systems are not much of a help. Near-field induction is a broadly used technique for recharging everyday devices like a toothbrush but it requires the source and load to be physically close. Near-field resonant induction works at greater distances and higher coupling efficiencies, and is used to power many medical device implants, but it requires the source and load to be tuned to each other.

Fortunately, new developments in RF-based wireless power transfer are poised to make maintenance-free operation for ultra low-power devices practical. For example, researchers recently demonstrated harvesting 60 μW from a TV station operating 4.1 km away broadcasting at an effective radiated power of 960 kW on channel 48 at 674-680 MHz, received using a 5 dBi (3.1$\times$ gain) antenna [23]. The harvested power is approximately 25% of that predicted by the Friis equation, which relates received power (P_r) to transmitted power ($P_t = 960$ kW), antenna gains at the transmitter ($G_t = 1$) and receiver ($G_r = 3.1$), wavelength ($\lambda = 0.44$ m), and transmission distance ($R = 4.1$ km),

$$P_r = P_t G_t G_r \left(\frac{\lambda}{4\pi R} \right)^2$$

Scaling the above terms suggests that a 1 W source could provide 20 μW of power to a device with a 0 dBi omni-directional antenna from 4 m away. There is some evidence that such systems can be built and may be commercially available in the near future [22]. Today, the FCC limits unlicensed transmitters to specific bands and maximum power levels, and restricts continuous broadcast at even modest power levels. This situation raises the question of whether the FCC might sanction a small slice of spectrum for wireless power transfer with restrictions better suited to μ-mobiscope applications.

In the past, researchers have explored harvesting mechanical energy but these systems currently have several drawbacks: they are relatively large; they must operate near their mechanical resonant frequency; and they must be mechanically well-coupled to a vibrating source. Recent research, however, shows ample mechanical energy is available in many locations on the body and that with micropower budgets and adaptively-tuned generators, mechanical harvesting might be feasible [27].

3 Example Applications

μ-mobiscopes are useful for monitoring the actions and interactions of people and things, and the movement of people and things through space, as the examples show.

3.1 Activity Recognition

Human activity recognition is an important application of μ-mobiscopes well-studied in the ubiquitous computing literature with applications in elder care, medical monitoring, sports, and smart buildings. Techniques as diverse as monitoring air flow in homes [21], carrying pager-sized mobile sensing platforms [7], and outfitting attire with smart sensors have been explored [14].

Views on the best approach are divided: some argue that lightweight infrastructure-based techniques are best because they are minimally invasive while others believe the more-the-merrier: that multiple, body-worn sensors provide the fidelity needed for detailed activity recognition and tracking. However, to be viable, as the number of body-worn sensors increase, their size must decrease, their maintenance must converge with daily patterns or disappear altogether, and their communications must be unobtrusive and configuration-free.

Today, it is possible to integrate a microcontroller, radio, antenna, flash memory, accelerometer, and rechargeable 10 mAhr lithium-polymer battery in an area no larger than a clothing tag and just a few millimeters thick. Unfortunately, with current asynchronous MAC protocols that achieve 0.3 to 1% duty cycles, a radio that runs at 25 mA in idle, and accelerometers that draw 25 to 100 μA to operate, a node will exhaust its battery in a few days – far short of the desired lifetime for many applications. Some have proposed deep duty cycling – turning off the radio and accelerometer after detecting stillness and turning the accelerometer on for a few hundred milliseconds every few seconds to check for stillness [14].

Duty cycling sensors this way is good for power management but bad for some applications: a sudden leg spasm, a quick finger tap, or the shock of two railcars coupling may not be detected. The nanopower vibration wakeup circuits we advocate, however, can detect trigger sampling within milliseconds of detecting motion, vibration, or shock.

3.2 Persistent Spatial Queries

Persistent spatial queries are a modern interpretation of the old saying, "When in Rome, do as the Romans do." In this approach, physical spaces have their own long-standing queries – local customs, so to speak – and nodes moving through these spaces execute the queries – that is, adhere to the local customs.

Consider the problem of tracking building occupancy. If occupants carried tags, then a building-wide persistent spatial query could determine occupancy in real-time. Each tag would receive, execute, and disseminate a SQL `COUNT`-like query so that all nodes within multihop radio range could participate. Variations on this theme could compute the preferred lighting level on `AVERAGE` among the population in a conference room, take inventory of the `DISTINCT` objects in a warehouse, determine the `MAX` temperature in the datacenter, obtain a `COUNT` of the raised hands in a classroom, or other similar aggregates over a time-varying population of nodes. These applications breathe new life into the aggregate queries first proposed for sensing applications with well-known node populations [17].

The prior examples involve aggregates because the identities of nodes may not be known *a priori*, but it is easy to imagine queries over space for individual objects as well. Imagine, for example, persistent queries for commonly misplaced objects like one's keys or wallet.

Some of the challenges in building such systems have been explored in prior work. Mobile nodes will have to quickly, efficiently, and asynchronously discover the presence and absence of network neighbors [10] and determine whether any neighbors have pending traffic [11]. Nodes will need to allocate some energy to queries, track the usage of a particular query [13], and limit the amount or rate of energy usage allocated to hosting queries [15] so as not to exhaust the limited supply.

3.3 Mobile Interactions

Prior work identified three common patterns in mobile-to-mobile or mobile-to-static interactions called *talking*, *docking*, and *flocking* [10]. In the *talking* pattern, two mobile nodes encounter each other, communicate, and part ways [16]. In the *docking* pattern, a mobile node discovers a static node at a rendezvous point [26]. In the *flocking* pattern, two or more nodes move as a group [4].

Consider the flocking example more closely: a wristwatch warns its owner that she is about to leave the house without her keys, wallet, or cell phone. The application uses RFID readers situated around the house to read passive tags attached to objects, transmit the readings to a body-worn personal server, and forward alerts to a wristwatch. This application might be implemented very differently, however, if tags were active, could detect motion, and harvest power wirelessly. Tags would stay in a deep sleep – drawing perhaps a microamp – when not in motion, but still able to detect movement. When moved, tags would wake up, enable their radios, attempt to contact members of their flock, and raise an alert if any were missing, and in the process, obviate the need for RFID readers and personal servers.

4 Research Challenges

μ-Mobiscopes couple low-power wireless sensornet technology with motion sensing electronics and ubiquitous computing usage models. At the hardware level, new motion sensors, energy harvesting techniques, and flash storage make low-power, mobile embedded devices feasible and imminent. At the systems level, however, mobility raises new challenges, as prior work has argued [1] and explored [6, 16, 18]. Prior work has also explored the challenges with intermittent power in the context of computational RFIDs [5]. In this paper, we focus on the challenges particular to μ-mobiscopes, a relatively unexplored computing class.

4.1 Hardware Platforms and Abstractions

New hardware platforms that integrate low-power motion detection sensors, wireless power transfer systems, and tiny batteries must be designed and built to evaluate the ideas in this paper. These platforms must be able to almost completely turn off their peripherals as continuous currents of even tens of microamps may be too high in some applications. One questions is how these new hardware capabilities should be abstracted and multiplexed. One application or service may request a callback upon any motion while another may want to wait for a few seconds of activity while a third awaits a double tap gesture. In some case, hardware may be able to support the request while in other cases, signal processing software may be required. How should the OS expose these capabilities? In addition to the nodes themselves, wireless power transfer infrastructure will be required. An open question is whether special-purpose hardware will be needed [22] or if RFID readers will suffice [23].

4.2 Fleeting Connectivity

Low-power link layers are greatly affected by mobility so it is an area that has seen some recent exploration. One critical question is how two or more low-power nodes that are rarely awake or co-located discover each other quickly, reliably, and securely when they are nearby? Prior work has explored this question [10] but not with the knowledge that a node is moving or stationary. What might change? Link estimation raises another challenge in mobile networks because history may be a poor guide in predicting future link performance. Recent research suggests a node wait briefly before retransmitting an unacknowledged packet, which may be appropriate in low-power static networks, but problematic in mobile ones [25]. More generally, issues of link estimation become ones of link detection, prediction, and termination where agility may be preferred to stability.

Wireless nodes choose routing peers at the network layer and these choices are already influenced by many cost metrics today including hopcount, expected number of transmissions (ETX) [8], and expected transmission time (ETT) [9]. We suggest adding a Motion ObserVed (MOV) metric to this list. Nodes could include the MOV metric in routing beacons and keep track of their neighbors' MOV values. Peers with small MOV values might offer stable paths while peers with large MOV values may not. A node might prioritize route-through traffic from a peer with a large MOV value over one with a small value because the first node might soon disappear.

The MOV metric could come from several sources, including the low-power sensors previously discussed, but it could be generated by the link layer by tracking the changes in the received signal strength as well. One drawback with a signal strength-based approach is that it requires the radio to be turned on and actively transmitting or receiving, which may not be desirable. The particulars of the MOV metric are open questions: What units? Computed over what timeframe? Absolute or relative to radio range or transmission power?

Intermittent connectivity raises many well-known transport layer challenges [12]. The same issues arise in μ-mobiscope applications, but the challenges are exacerbated in this regime for several reasons. First, connectivity may be fleeting and not simply intermittent. Such limited connectivity requires that the network stack offer *prioritized* data transfer. These ideas have been explored in the context of application-specific software but not within a larger networking architecture. Second, fleeting connectivity suggests that nodes should upload data as fast as possible, and probably in bursts, suggesting that TCP's end-to-end flow control may be too slow and that high-speed, bundle-oriented, hop-by-hop message transfers are more appropriate. Third, even with hop-by-hop transfers, nodes may not be able to afford to repeatedly resend failed transfers, as might happen when a node loses connectivity with a neighbor. Data should still move toward the destination, perhaps not as bundles, but as partial ADUs, fragments, or segments.

4.3 Security and Privacy

μ-mobiscopes raise a number of security questions that are similar in spirit to ones raised by active networks, especially for pervasive spatial queries: How should these queries be expressed? How should DoS or energy deprivation attacks be prevented? How should the queries be disseminated? In addition to these security questions, mobile sensors have the potential to leak large amounts of personally identifiable information, just like Nike+iPod devices and RFIDs. Balancing security, privacy, and usability continues to present a challenge.

5 Conclusion

This paper argues that low-power, mobile sensing applications are imminent but that mobility invalidates many assumptions implicit in today's low-power, static sensor network designs. We suggest that real-time knowledge of motion can help mitigate many challenges that mobility itself raises and we propose a new metric, *Motion ObserVed* (MOV), describe how it might be gathered on a nanopower energy budget, and explore how it could address some of the problems that mobility creates.

References

[1] ABDELZAHER, T., ANOKWA, Y., BODA, P., BURKE, J., ESTRIN, D., GUIBAS, L., KANSAL, A., MADDEN, S., AND REICH, J. Mobiscopes for human spaces. *IEEE Pervasive Computing 6*, 2 (2007), 20–29.

[2] ANALOG DEVICES. ADXL330: Small, Low Power, 3-Axis, ± 3 g iMEMS Accelerometer, Sept. 2006.

[3] ANALOG DEVICES. ADXL345: Three-Axis, $\pm 2/4/8/16$ g Digital Accelerometer, Nov. 2008.

[4] BORRIELLO, G., BRUNETTE, W., HALL, M., HARTUNG, C., AND TANGNEY, C. Reminding about tagged objects using passive RFIDs. In *Ubicomp* (2004), pp. 36–53.

[5] BUETTNER, M., GREENSTEIN, B., SAMPLE, A., SMITH, J. R., AND WETHERALL, D. Revisiting smart dust with RFID sensor networks. In *Hotnets-VII: Proceedings of the 7th ACM Workshop on Hot Topics in Networks* (Oct. 2008).

[6] BYCHKOVSKY, V., CHEN, K., GORACZKO, M., HU, H., HULL, B., MIU, A., SHIH, E., ZHANG, Y., BALAKRISHNAN, H., AND MADDEN, S. The cartel mobile sensor computing system. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems* (New York, NY, USA, 2006), ACM, pp. 383–384.

[7] CHOUDHURY, T., BORRIELLO, G., CONSOLVO, S., HAEHNEL, D., HARRISON, B., HEMINGWAY, B., HIGHTOWER, J., KLASNJA, P., KOSCHER, K., LAMARCA, A., LANDAY, J. A., LEGRAND, L., LESTER, J., RAHIMI, A., REA, A., AND WYATT, D. The mobile sensing platform: An embedded system for activity recognition. *IEEE Pervasive Magazine-Special Issue on Activity-Based Computing 7*, 2 (April 2008), 32–41.

[8] DE COUTO, D., AGUAYO, D., BICKET, J., AND MORRIS, R. A high-throughput path metric for multi-hop wireless routing. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking* (2003), pp. 134–146.

[9] DRAVES, R., PADHYE, J., AND ZILL, B. Routing in multi-radio, multi-hop wireless mesh networks. In *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking* (2004), pp. 114–128.

[10] DUTTA, P., AND CULLER, D. Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications. In *SenSys '08: Proceedings of the 6th international conference on Embedded networked sensor systems* (Nov. 2008), pp. 71–83.

[11] DUTTA, P., MUSALOIU-E., R., STOICA, I., AND TERZIS, A. Wireless ACK collisions not considered harmful. In *HotNets-VII: The Seventh Workshop on Hot Topics in Networks* (Oct. 2008).

[12] FALL, K. A delay-tolerant network architecture for challenged internets. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications* (2003), ACM Press, pp. 27–34.

[13] FONSECA, R., DUTTA, P., LEVIS, P., AND STOICA, I. Quanto: Tracking energy in networked embedded systems. In *OSDI'08: Proceedings of the Eighth USENIX Symposium on Operating System Design and Implementation* (Dec. 2008).

[14] GANTI, R. K., JAYACHANDRAN, P., ABDELZAHER, T. F., AND STANKOVIC, J. A. Satire: a software architecture for smart attire. In *MobiSys '06: Proceedings of the 4th international conference on Mobile systems, applications and services* (New York, NY, USA, 2006), ACM, pp. 110–123.

[15] JIANG, X., TANEJA, J., ORTIZ, J., TAVAKOLI, A., DUTTA, P., JEONG, J., CULLER, D., LEVIS, P., AND SHENKER, S. An architecture for energy management in wireless sensor networks. In *WSNA'07: International Workshop on Wireless Sensor Network Architecture* (Apr. 2007).

[16] LIU, T., SADLER, C. M., ZHANG, P., AND MARTONOSI, M. Implementing software on resource-constrained mobile sensors: experiences with Impala and ZebraNet. In *MobiSys '04: Proceedings of the 2nd International conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2004), ACM, pp. 256–269.

[17] MADDEN, S. R., FRANKLIN, M. J., HELLERSTEIN, J. M., AND HONG, W. TAG: a Tiny AGgregation Service for Ad-Hoc Sensor Networks. In *OSDI'02: Proceedings of the ACM Symposium on Operating System Design and Implementation* (Dec. 2002).

[18] MALINOWSKI, M., MOSKWA, M., FELDMEIER, M., LAIBOWITZ, M., AND PARADISO, J. A. Cargonet: a low-cost micropower sensor node exploiting quasi-passive wakeup for adaptive asychronous monitoring of exceptional events. In *SenSys '07: Proceedings of the 5th International Conference on Embedded Networked Sensor Systems* (New York, NY, USA, 2007), ACM, pp. 145–159.

[19] MEASUREMENT SPECIALTIES. ACH-04-08-05: Accelerometer Application Note 01800024-000 Rev A, Mar. 2005.

[20] MEASUREMENT SPECIALTIES. MiniSense 100 Vibration Sensor, Mar. 2005.

[21] PATEL, S. N., REYNOLDS, M. S., AND ABOWD, G. D. Detecting human movement by differential air pressure sensing in hvac system ductwork: An exploration in infrastructure mediated sensing. In *Pervasive* (2008), J. Indulska, D. J. Patterson, T. Rodden, and M. Ott, Eds., vol. 5013 of *Lecture Notes in Computer Science*, Springer, pp. 1–18.

[22] POWERCAST: WIRELESS POWER FOR A WIRELESS WORLD. Wireless Sensor Power Distribution Systems, 2008. `http://powercastco.com/wireless_power/wireless_sensors.html`.

[23] SAMPLE, A., AND SMITH, J. Experimental results with two wireless power transfer systems. In *RWS'09: Proceedings of the IEEE Radio and Wireless Symposium* (Jan. 2009).

[24] SIGNALQUEST PRECISION MICROSENSORS. SQ-SEN-200: Omnidirectional Tilt and Vibration Sensor, Jan. 2009.

[25] SRINIVASAN, K., KAZANDJIEVA, M. A., AGARWAL, S., AND LEVIS, P. The β-factor: measuring wireless link burstiness. In *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems* (2008), pp. 29–42.

[26] WARK, T., HU, W., SIKKA, P., KLINGBEIL, L., CORKE, P., CROSSMAN, C., AND BISHOP-HURLEY, G. A model-based routing protocol for a mobile, delay tolerant network. In *SenSys '07: Proceedings of the 5th International Conference on Embedded Networked Sensor Systems* (New York, NY, USA, 2007), ACM, pp. 421–422.

[27] YUN, J., PATEL, S., REYNOLDS, M., AND ABOWD, G. A quantitative investigation of inertial power harvesting for human-powered devices. In *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing* (New York, NY, USA, 2008), ACM, pp. 74–83.