


Mobiscopes for Human Spaces

*Tarek Abdelzaher, Yaw Anokwa, Péter Boda, Jeff Burke, Deborah Estrin,
Leonidas Guibas, Aman Kansal, Samuel Madden, and Jim Reich*

Vol. 6, No. 2
April–June 2007

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.



Mobiscopes for Human Spaces

Mobiscopes extend the traditional sensor network model, introducing challenges in data management and integrity, privacy, and network system design. Researchers need an architecture and general methodology for designing future mobiscopes.

A *mobiscope* 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 affordably extend into regions that static sensors cannot, proving especially useful for applications that only occasionally require data from each

location. They represent a new type of infrastructure—virtual in that a given node can participate in forming more than one mobiscope, but physically coupled to the environment through carriers, including people and vehicles. Mobiscope applications include public-health epidemiological studies of human exposure using mobile phones and real-time, fine-grained automobile traffic characterization using sensors on fleet vehicles. Although mobility has proven critical in many scientific applications, such as Networked Infomechanical Systems for science observatories,¹ we focus on the challenges and opportunities mobiscopes pose in human spaces.

Mobiscopes complement static sensing systems by addressing the fundamental limitations created by fixed sensors. System

designers can't always place sensing devices with sufficiently high spatial density to accurately sample the field of spatially varying phenomena,

making it impossible to satisfy the spatial band-limiting guarantees that traditional sampling criteria require. Covering large areas can be challenging because of the need for long dwell times, the unavailability of wired power, the impracticality of battery replacement, the inability of any entity to install devices across the entire area, and the expense of purchasing and maintaining enough devices. Equally important, target sensor types might be unavailable or unaffordable in the form of autonomous instruments, further motivating mobile, human-in-the-loop instruments. For example, city-scale air quality measurements—which typically use costly mass spectrometers to measure pollutants—are expensive when using fixed-sensor infrastructures but could be substantially cheaper and could cover much larger areas if sensors were mounted on mobile nodes (for example, cars).

This combination of application demand and increasingly powerful wireless and sensing technology suggests that it's time to consider a general architecture for mobiscopes. To understand what's needed to build a unified system, we consider several broad classes of mobiscope. We discuss common architecture challenges, existing solutions, and major areas for future work.

Classes of mobiscopes

Early mobiscopes arose directly from widely available sensing modalities in networked devices. Examples include image sensors in mobile phones, GPS in phones and vehicles, and the increasingly diverse telemetry available in vehicles. We consider these as representatives of two broad categories of mobiscope.

Tarek Abdelzaher
*University of Illinois
at Urbana-Champaign*

Yaw Anokwa
University of Washington

Péter Boda
Nokia Research Center

Jeff Burke and Deborah Estrin
*University of California,
Los Angeles*

Leonidas Guibas
Stanford University

Aman Kansal
Microsoft Research

Samuel Madden
*Massachusetts Institute
of Technology*

Jim Reich
Palo Alto Research Center

Vehicular mobiscopes

One category is vehicular applications for traffic and automotive monitoring,² where a subset of equipped vehicles senses various surrounding conditions such as traffic, road conditions, or weather. These mobiscopes exploit the spatial oversampling often provided by dense vehicle traffic to produce useful information before all vehicles can send data. However, even when vehicular instrumentation achieves nearly 100 percent market penetration, subsampling the data intelligently will prevent network congestion and save storage and processing space.

Initial probe applications have already been deployed commercially. For example, Inrix (www.inrix.com) uses anonymous GPS data to provide real-time traffic measurements for freeways and local streets. Vehicular mobiscopes can query for certain traffic types. For example, the EZCab application uses vehicle-to-vehicle communication to find available taxi cabs.³ The probe-car concept can extend to other applications, such as augmenting the number of NavTeq or TeleAtlas vehicles used for street mapping or increasing the update frequency of Microsoft's street-level imagery capture for Virtual Earth. Probe cars can also acquire high-density maps of roadways and measure road conditions, weather, and pollution using sensors built into cars and phones.

Handheld mobiscopes

The second emerging category is mobiscopes that use handheld devices. Coarse-grained location information can inform studies ranging from the health impact of exposure to highway toxins to an individual's use of transportation systems. Researchers have proposed automated image and acoustic capture to provide user feedback on diet, exercise, and personal interaction as well as to identify and share real-time information about civic hazards and hotspots. An

interesting example is civic participation during a crisis, where individuals could exercise a loose form control over sensor placement.⁴ Users ranging from police officers to citizens could use their cell phone cameras to photograph trou-

ble spots in their neighborhood. Such a civic system could request that police officers document unexplored areas or intervene in trouble spots.

A similar concept of camera-based mapping can apply to tourism. For example, tourists at the Taj Mahal might share their photographs in virtual albums that potential visitors can then browse to see all perspectives of the mausoleum. Researchers have paid special attention to metadata management to facilitate such sharing.⁵

Common requirements

The applications we've mentioned share several important requirements that are also a priority for mobiscope operation and acceptance. For example, data persistence must be assured even when sensing nodes leave the data collection area or when no mobile nodes are present. At the same time, data access tends to be spatially correlated with the users' location and can change rapidly (somewhat predictably) as the user moves. The system can use data to make decisions in real time. It might use a human-in-the-loop as an actuator, sensor, interpreter, or responder. Because the system will exploit sensors and mobility sources already in the environment, social constraints on system behavior come into play. Many private and public entities will likely share ownership of sensors and the resulting data, so we can't assume trust, coordi-

nated deployment, and respect of users' privacy. The needed sensor data might be fragmented across multiple networks, and connectivity and user needs might change dynamically with motion. Furthermore, the metadata (such as sensor position, ori-

Heterogeneity is a fundamental and beneficial quality of mobiscopes, not just a problem to overcome.

entation, and calibration parameters) must also compatibly cross networks.

The commonality of problems across a wide range of envisioned mobiscope systems calls for a general architecture and design guidelines for future mobiscopes. This architecture will not only encourage component reuse and reduce development costs but also promote interoperability among future mobile sensing systems. The systems will need common interfaces to negotiate privacy settings, exchange data feeds, distill information, and perform coordinated actuation on the physical world. The literature discusses such architectures' components but often focuses on a specific application type—for example, an architecture for vehicle-to-vehicle live video streaming^{6,7} or architectural support for heterogeneity based on a MediaBroker.⁸

Mobility and sampling coordination

Fundamentally, mobiscopes' performance depends on mobility patterns of transporters—which are stochastic and whose movement patterns range from highly structured (such as road traffic) to less structured (such as foot traffic)—and on sensor densities that can vary widely over time and space. Several challenges arise from mobility. The network organization can be highly variable, both in terms of sensing coverage and radio connectivity. Nodes might not have con-

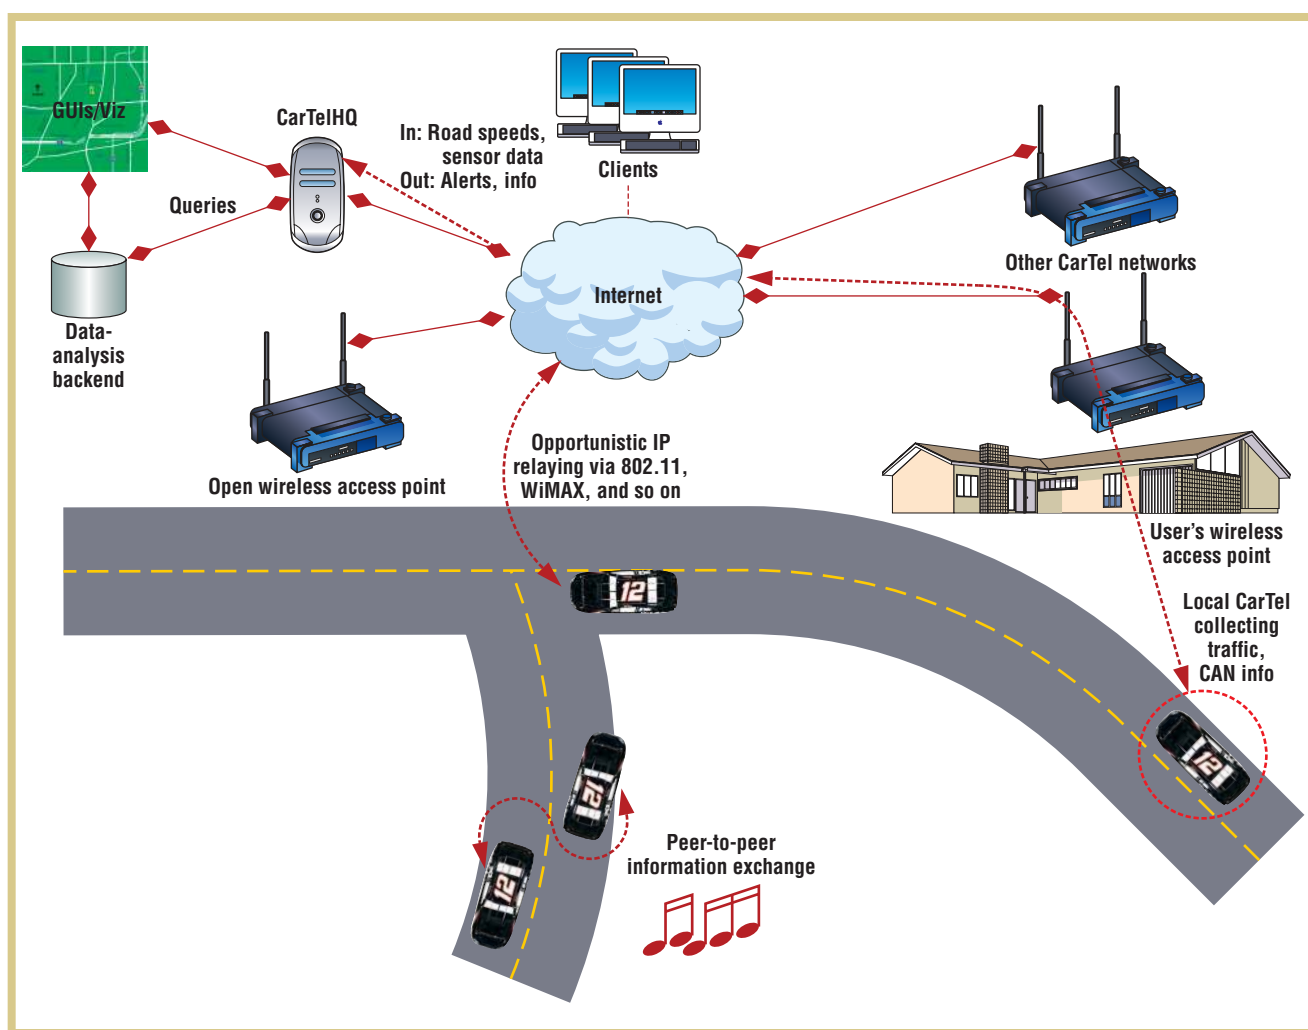


Figure 1. CarTel—a typical mobiscope.

nectivity at all times and locations when they collect data. Understanding data acquisition and distribution behavior in the network is therefore nontrivial.

Often, uncoordinated mobile nodes sense areas that other nodes have already visited, but they don't visit rarely traveled areas frequently enough. For example, in MIT's CarTel project (see figure 1), researchers have equipped cars with small embedded computers, wireless radios, and GPS sensors and cameras for measuring road speed and capturing road conditions.⁹ Members of the community drive the cars as they go about their daily business. Unfortunately, their daily business often takes them on the same routes, providing only sparse cov-

erage of back roads and outlying communities while the MIT campus is highly (and redundantly) covered.

Mobility also causes challenging dynamic behavior in sensor allocation and network topologies. In a typical mobiscope, particular regions in space will likely be of interest, but the motion of the sensors (especially when commuters or pedestrians are carrying them) might make coverage for those regions poor. For example, the midblock regions of roadways with traffic lights often contain no cars, so data might be unavailable from that region, creating a hole in the peer-to-peer network as well. The mobile sensing devices' availability can depend on user behavior and device characteristics. For

example, users might forget or turn off their phones, making them unavailable to the mobiscope, or a data service might not be available during a telephone conversation. We need solutions to these challenges. Sensors' availability can also change drastically with time, as with cars and pedestrians at rush hour compared to low availability at midnight.

Application adaptation

Applications must adapt to the network's available communication and sensing characteristics. For instance, they could buffer data when connectivity is unavailable or dynamically adapt their spatial scope. Mechanisms from disruption-tolerant networks and data-mulling

contexts will be increasingly important. Analytic foundations also help explain global network behavior and the extent to which adaptation can help maintain acceptable system performance. Research has addressed adaptation to some extent—for example, summarizing data in low-bandwidth environments¹⁰ or trading latency for bandwidth by using vehicles to carry large amounts of data.¹¹ Yet, substantial questions remain, particularly with respect to global resource optimization. It's unclear how to best relate diverse resources such as total data transfer bandwidth, data freshness, data propagation latency, and application-relevant metrics such as traffic conditions or time of day.

Actuated mobility

In some mobiscopes, it might be possible to task some or all of the nodes to visit a specific location to collect information on demand. We call such mobiscopes *actuated*, and the actuable nodes *actuators*. In an actuated mobiscope, server nodes can record areas that most need to be visited and can task actuators to visit those areas either one at a time or as part of a circuit. Actuated mobiscopes present numerous interesting research challenges related to determining the value of observing a particular location or collecting information versus the cost of sending an actuator to that location. One common way to address these challenges is to formulate them as an optimization problem that aims to maximize the utility of information collected in a particular time period or subject to an energy or cost constraint.¹² Distributed robotics has also addressed similar coverage problems.^{1,13}

Opportunistic connectivity

In networks without predictable mobility, opportunistic connectivity—where nodes happen to come into contact with each other or with network infrastructure (such as an open 802.11

network)—can substantially improve connectivity. This technique performs better than mechanisms that wait until nodes return to some “home” location where infrastructure connectivity exists and can be cheaper and potentially more efficient than solutions that rely on fixed cellular infrastructure (where upstream data rates are typically highly constrained). Important techniques worth exploiting include

- building low-level network protocols that can quickly identify and associate with nearby nodes (without, for example, long address acquisition or channel negotiation delays)⁹ and
- designing routing algorithms that can deliver data through such opportunistic connections (for example, by using limited-scope flooding or exploiting point-to-point long-haul connections in the network).

Prioritization

In many mobiscopes, it's likely that the system will capture more information than it can deliver in real time. One solution is aggregation,¹⁴ but another alternative is prioritization, which assigns dif-

ferent priorities to pieces of collected data and delivers data in priority order. In some cases, simple prioritization, where particular data types (such as emergency alerts¹⁵) are given greater importance, might be sufficient. However, a more coordinated form of prioritization is necessary when different nodes cover overlapping geographic areas to avoid wasting valuable bandwidth on redundant data reports. To avoid redundant reports, you need coordination—for example, before sending large collections of sensor readings to a server responsible for presenting data to the user, a mobile node might first send a compact summary of its data. The server can then examine the summary and determine what information is valuable given the data it has already received from other mobile nodes. CarTel employs a similar technique.¹⁶ In addition, many applications have data requirements that vary with time, location, and their need for frequent, low-latency data delivery. This might be summarized by a more complex utility function that can radically increase efficiency over fixed-priority approaches.¹⁷

Challenges and opportunities of heterogeneity

Mobiscopes will come into being through many different mechanisms. In some cases, hardware will be explicitly deployed as a single mobiscope to achieve a particular data-collection goal. In others, already deployed devices will be fed-

Heterogenous sensing systems are far more immune to weaknesses of any one sensing modality and more robust against defective, missing, or malicious data sources than even carefully designed homogeneous systems.

erated into an ad hoc collection through their owners' participation. In still different cases, virtual mobiscopes will be formed by correlating gathered data without edge devices.

Mobiscopes will thus take on various topologies and structures, federate devices

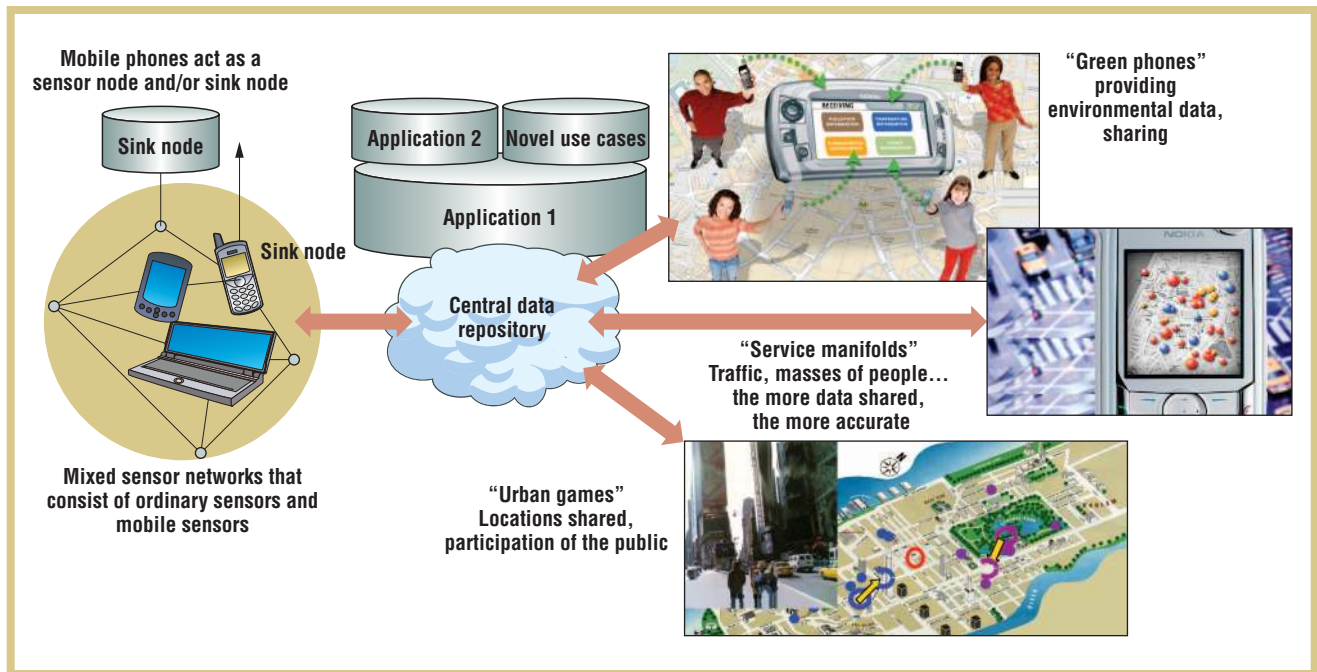


Figure 2. The Nokia SensorPlanet initiative. Data from various sources is collected in a central repository that users can query and process.

with different capabilities, and draw together components with varying levels of trust and credibility. We can classify heterogeneity in mobiscopes across several dimensions: structure and topology, transducer performance, data ownership, and data dissemination qualities (such as resolution). Irregular data streams are inevitable owing to sensors' mobile nature. We must prepare applications that can adapt task allocations and service levels to the available resources.

Heterogeneity is a fundamental, beneficial quality of mobiscopes, not just a problem to overcome. Heterogenous sensing systems are far more immune to the weaknesses of sensing modalities and far more robust against defective, missing, or malicious data sources than even carefully designed homogeneous systems. In some cases, the information the application needs is available only by fusing data from several different sensing modalities. Moreover, the same qualities necessary to support heterogeneous sensors also help the system adapt to new varieties of sensors that are developed over time or brought from other areas into the sensed space.

Heterogeneity of ownership

Unlike a sensor network that a single entity (such as a corporation or university laboratory) has assembled, mobiscopes must be federated from individually owned devices, held or worn by owners who might not be trustworthy and might not maintain their devices in good condition. This exposes the mobiscope to intentional or accidental injection of incorrect data or biased sampling. It also raises issues of data integrity, trust, security, and selective sharing that the architecture should address.

Heterogeneous data resolution and types

Applications can use the available data to derive and maintain metrics at multiple resolutions with varying coverage in space and time. For instance, imagine a system of wearable ski-slope monitoring devices that collect data from poll-mounted sensors.¹⁸ Coarser granularity measurements might be available for more area and time instances, while finer-granularity data is served only for specific regions (namely, the more populated slopes). Applications

can receive regularized grid interpolations derived from raw data. Simple interpolations might be sufficient for smoothly varying data such as temperature, while more complex known or learned dynamics models fill in the gaps in faster-varying or sparser data. In addition, buffering data and providing aggregates over multiple time-window sizes might allow the data from an irregularly sampled world to yield reasonably uniform coverage as data trickles in over time. Researchers have examined using model-based techniques to improve the quality of data despite missing values,^{19,20} but adapting this work to heterogeneous data of varying types remains a challenge.

Data heterogeneity also presents challenges when trying to integrate data from many different sensors. Nokia's SensorPlanet effort (www.sensorplanet.org; see figure 2) and the Microsoft SensorMap project invite users to submit data from their own sensing deployments. Although these efforts present a great potential resource, to effectively use the data, we need to a way to combine and query these different data sets.

Robustness through data heterogeneity

So, heterogeneity is both a difficulty and an asset. Heterogeneity of devices implies that different transducers (for example, image, acoustics, physical sensors, and user annotations) are available at different times at the nearby locations. However, mobiscopes' nature implies that any sensor fusion algorithm relies on the transporters and the network to determine when it will receive data and what type of data will arrive. This puts additional stress on the sensor fusion and estimation algorithms. This also suggests significant advantages for model-driven approaches such as Kalman filters²¹ or particle filters,²² which adapt well to irregular sampling and allow the use of heterogeneous sensor models and systems dynamics models. In some cases, system designers can even harness heterogeneity to increase the robustness to defective, malicious, or unavailable sensors.²³

Tackling data privacy

The topic of privacy hinges on complex policy issues that vary culturally and legally among societies, but fundamentally, it's about people's ability to control information flow about themselves. These issues are especially difficult in mobiscopes because the connection between the observed party and the sensor collecting information is implicit in the transporters' relationships and because entities collecting data often inadvertently reveal information about themselves. The nodes' distributed ownership and control makes policies even more difficult to define, let alone enforce.

Policy definition

Beyond data distribution and management, data privacy issues present important trade-offs between the need for selective sharing and the network information output's fidelity, configurability, usability,

and verifiability. The inability to publicly associate data with sources (for privacy reasons) could lead to loss of context, which reduces the network's ability to generate useful information. Conversely, revealing too much context can potentially thwart anonymity, violating privacy requirements.

For example, consider sampling ambient sound from a carried cell phone to map sound pollution profiles or as a proxy for exposure to highways and associated fumes. Similarly, consider the use of cell phone cameras as a form of civic engagement where citizens document concerns such as uneven sidewalks, lack of handicap access, and overflowing trash cans. Clearly, these data are more meaningful if the users send the associated GPS readings along with the primary data. However, transmission of GPS data could reveal the user's identity (for example, if the GPS trace ends up at a particular address most nights). It can also shed light on user activities.

This problem is often termed *location privacy*.²⁴ Ensuring location privacy is a cross-cutting challenge that has implications for routing²⁵ and energy management.²⁶ A mobile node isn't restricted to a known private space under a single user's jurisdiction but passes through multiple spaces. Furthermore, a mobile node isn't part of a dedicated application but might serve different applications at different times. So, the system might require effective user interfaces and, in some cases, automatic adaptation.

Local processing

Related to these concerns is an individual's need to keep something from

becoming data, not just reducing access to collected data. In many contexts, this means putting the selectivity and filtering capabilities on the end-user node itself rather than relying on post facto filtering. Another alternative is to give the user a private server space in which to review his or her data before release⁴—

**It's important to develop systems
where users can verify data's correctness
without violating the source's privacy.**

at the expense of latency and other usability aspects, of course.

Verification

Another related issue is trust and data integrity. It's important to develop systems where users can verify data's correctness without violating the source's privacy. Moreover, because distributed data management in mobiscopes relies on user cooperation, a challenge becomes introducing proper incentives that promote successful participation and prevent abusive access with the purpose of "gaming the system." In some cases (such as a peer-to-peer vehicle or handheld device network), privacy constraints, transience of communication between participants, and the sheer number of participants might actually make cryptographically authenticating the user's identity undesirable or impractical. In these cases, the correct solution might be to instead rely on redundancy in the sensor data to validate a data source anonymously.²⁷ The literature presents an early account of many other privacy (and security) issues in sensor networks.²⁸

Privacy preserving data mining

Although we can resolve several privacy challenges by employing a trusted authority that guarantees nondisclosure, the need to be trusted increases the bar-

rier to entry for anyone who might want to contribute aggregation and search functions to the sensor Web. This begs the question: Can a system achieve privacy-preserving aggregation and search without trusting the entity performing these operations? In this case, a user isn't willing to share his or her data (with the untrusted node) but might be interested

ple, a society of three), knowing an accurate average can reveal a lot about the individual values.

The choice of the best perturbation algorithm for a data set is a nontrivial problem that has received much attention.³² Researchers have also studied the extent to which additive noise improves privacy, revealing cases where private data

needs. A direct result of this information reduction requirement is the need for network storage as a key service because aggregation and filtering both imply a need to buffer information for a potentially long time. Users might need to use disruption-tolerant protocols and architectures³⁴ to diffuse data toward destinations and buffer data when no opportunities exist for making progress.

Information reduction in disruption-tolerant networks raises a host of research challenges. One such challenge is protocols that integrate opportunistic en route aggregation mechanisms with buffering. The interplay between storage and data reduction in mobiscopes also offers interesting directions for rethinking other basic network functions such as congestion control. For example, you can reduce information more aggressively as a congestion control mechanism to alleviate high storage use or to increase user privacy by keeping data local. A direct consequence of the need to process (for example, reduce) data inside the mobiscope is having to focus on network programming issues. So, a mobiscope architecture should present not only communication protocol and data management interfaces but also programming interfaces for in-network computation. In short, the advent of ubiquitous networks of mobile embedded devices shifts the fundamental networking paradigm, offering application programmers interesting challenges in network architecture, protocol design, and exported abstractions. Resolving these challenges to shape future network standards will be a most interesting task for the research community over the next decade.

Mobiscopes will be tightly coupled with their users, presenting significant human-factors design challenges and many sociocultural implications.

in the result of aggregation over the target community. An important field that solves this problem is privacy-preserving statistics and data mining.²⁹ Typically, the system uses additive random noise to perturb data without affecting the statistics to be collected. It can then use perturbed data, for example, to compute original data distributions³⁰ or construct decision trees for data classification³¹ without disclosing original values.

Consider a trivial example where a population wants to compute its average undetected number of speeding violations per month (as measured by speed sensors mounted in vehicles that correlate actual speed to the speed limit read from a map at the vehicle's current GPS location). A trivial solution is one where each sensor adds a zero-mean random number to the actual number of violations, sharing the resulting total. The disclosed total doesn't reveal the actual private data. Nevertheless, averaging the totals over a large population gives the true population average (because the added noise averages out). This has the additional benefit that if the population isn't large enough, the computed statistic isn't accurate. This is an advantage because in a small population (for exam-

can be reconstructed in violation of the intent from perturbation. For example, private data reconstruction is possible in the presence of high data correlation.³³

Another challenge is developing robust solutions (with respect to various attacks on privacy) that apply to statistical operations such as regression (used in the traffic example) and classification (used in learning). In general, because future mobiscopes' main goal will be information distillation from raw data, system designers will need theoretical foundations for obfuscating the raw data in a way that reconciles privacy requirements on individual measurements with the ability to compute certain aggregate properties of the collective.

Networking challenges

Integrating sensing-capable mobile devices into the networking infrastructure shifts the network's main utility from data communication to information filtering. For humans to make sense of the constantly increasing flow of data from sensors (and other sources), they must have tools to selectively filter, aggregate, and disseminate data on the basis of individual data consumers' expressed or statistically most likely

Human factors and social implications

Mobiscopes will be tightly coupled with their users. This presents significant human-factors design challenges and many sociocultural implications that extend beyond limited notions of privacy in data

transmission and storage. Both of these issues should become integral considerations for designing future mobiscopes.

Academics contemplating observing systems for various disciplines—including social sciences, public health, arts, and humanities—have considered social issues stemming from the autonomous observation of individuals by information technology.^{35–37} Mobiscopes are part of our maturing ability to silently watch ourselves and others, and their design and development must consider related social issues. Traditional responses to these concerns often postpone serious investigation because the technology is somehow considered immature. This is no longer true, so we propose four areas for the community to have a richer discussion of these observing technologies' social implications:

- explicitly considering broader policy precedents in information privacy as they apply to mobiscopes,
- extending popular education on information technology's new observation capabilities,
- facilitating individuals' participation in sensing their own lives, and
- helping users understand and audit their own data uploads.

For example, Rakesh Agrawal and his colleagues applied the Organisation for Economic Co-operation and Development international guidelines for data protection—collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability—to medical-record databases.³⁷ They demonstrated that implementing or supporting such ideals in technology presents rich research problems with few immediate solutions to practical concerns. Similarly, Harry Hochheiser compared the World Wide Web Consortium's Platform for Privacy Preference with the US data privacy policy.³⁸ Mobiscope system design should

begin with a similarly broad consideration of the existing policy base and the practical concerns of the people being sensed rather than focusing solely on narrower challenges, such as protecting sensitive data during transmission and storage.

Finally, mobiscope design can involve user interfaces of mobile devices, which haven't been present in traditional embedded systems. We also suggest that having a local UI provides a key opportunity for ambient and explicit feedback to the user on what data uploads are occurring from mobile devices at any given time. It can also help users configure their sensing participation and provide feedback on operational status. An effective interface might present real-time streams or historical examples of locally collected data to help users decide their desired privacy and sharing settings. Or it might present long-term features extracted from local or remote data over some geographic region.

For instance, an application providing both upload feedback and location-based information might display a stream of automobile traffic information from an aggregation service on a map by color or line thickness (instead of actual numbers of passing cars in a part of a city) or through simple classification such as *avoid*, *heavy traffic*, *easy*, and *practically no traffic* categories on a speech-enabled interface. At the same time, the display of a multimodal interface might remind users that the system was anonymously sharing their positions to contribute to the traffic map and give them the option to disable it. This alone presents many challenges in guaranteeing to the user that the system is actually doing what it says it's doing.

Existing research partially addresses some of the challenges mobiscopes pose—for example, several researchers have addressed privacy in sensor networks^{25,39} and examined network architectures for par-

ticular classes of mobiscopes.⁹ Unfortunately, this early work didn't address the broader architectural issues that must be considered before mobiscopes can be widely deployed. Much work must still be done on platforms and APIs that offer efficient, robust, private, and secure networking and sensory data collection in the face of heterogeneous connectivity and mobility. ■

REFERENCES

1. R. Pon et al., "Networked Infomechanical Systems (NIMS): Next Generation Sensor Networks for Environmental Monitoring," *IEEE MTT-S Int'l Microwave Symp. Digest*, June 2005, pp. 4.
2. B.S. Kerner et al., "Traffic State Detection with Floating Car Data in Road Networks," *Proc. 2005 IEEE Int'l Conf. Intelligent Transportation Systems (ITS 05)*, IEEE Press, 2005, pp. 44–49.
3. P. Zhou et al., "EZCab: A Cab Booking Application Using Short-Range Wireless Communications," *Proc. 3rd IEEE Int'l Conf. Pervasive Computing and Communications*, IEEE CS Press, 2005, pp. 27–38.
4. A. Parker et al., "Network System Challenges in Selective Sharing and Verification for Personal, Social, and Urban-Scale Sensing Applications," *Proc. 5th Workshop Hot Topics in Networks*, ACM Press, 2006, pp. 37–42.
5. M. Davis et al., "MMM2: Mobile Media Metadata for Media Sharing," *CHI 05 Extended Abstracts on Human Factors in Computing Systems*, ACM Press, 2005, pp. 1335–1338.
6. M. Guo, M.H. Ammar, and E.W. Zegura, "V3: A Vehicle-to-Vehicle Live Video Streaming Architecture," *Proc. 3rd IEEE Int'l Conf. Pervasive Computing and Communications*, IEEE CS Press, 2005, pp. 171–180.
7. M.D. Dikaiakos et al., "VITP: An Information Transfer Protocol for Vehicular Computing," *Proc. 2nd ACM Int'l Workshop Vehicular Ad Hoc Networks*, ACM Press, 2005, pp. 30–39.
8. M. Modahl et al., "MediaBroker: An Architecture for Pervasive Computing," *Proc.*



Tarek Abdelzaher is an associate professor in the Computer Science Department of the University of Illinois at Urbana-Champaign. His research interests include real-time computing, distributed systems, sensor networks, and control. He received his PhD in adaptive real-time systems from the University of Michigan. He is the editor in chief of the *Journal of Real-Time Systems* and an associate editor of *IEEE Transactions on Mobile Computing*, *IEEE Transactions on Parallel and Distributed Systems*, *ACM Transactions on Sensor Networks*, and *Ad Hoc Networks*. He's a member of the IEEE and ACM. Contact him at the Dept. of Computer Science, Univ. of Illinois at Urbana Champaign, 201 N. Goodwin Ave., Urbana, IL 61801; zaher@cs.uiuc.edu.



Yaw Anokwa is a PhD student in computer science at the University of Washington. His research interests are building, deploying, and evaluating low-cost technologies that positively impact the developing world and exploring the ubiquitous computing and human-computer interaction problems in that space. He has a BSc in computer science from Butler University and a BSc in electrical engineering from Indiana University-Purdue University Indianapolis. Contact him at Computer Science and Eng., Box 352350, Univ. of Washington, Seattle, WA 98195; yanokwa@cs.washington.edu.



Péter Boda is a principal scientist at the Nokia Research Center, Palo Alto. He heads SensorPlanet, a Nokia initiative on large-scale mobile-device-centric wireless sensor networks, which aims to set up a global open-research framework with the participation of leading universities. His research interests include data fusion of sensor information and the development of mobile applications where the physical and digital worlds merge to provide a better user experience, as well as advanced user interface solutions with fused contextual and sensor information accurately describing the user's intention. He received his LicTech from the Helsinki University of Technology. Contact him at the Nokia Research Center, 955 Page Mill Rd., Palo Alto, CA 94304; peter.boda@nokia.com; http://research.nokia.com/people/peter_pal_boda.



Jeff Burke is the executive director of REMAP, the Center for Research in Engineering, Media, and Performance at the University of California, Los Angeles. In collaboration with UCLA's Center for Embedded Networked Sensing, he's focusing on urban public space uses for mobile, embedded, and media technologies, including approaches that involve communities in system specification and design. Contact him at REMAP, Univ. of California, Los Angeles, 102 E. Melnitz Hall, Los Angeles, CA 90095-1622; jburke@ucla.edu.



Deborah Estrin is a professor of computer science at the University of California, Los Angeles, where she holds the Jon Postel Chair in Computer Networks. She's also the founding director of the university's NSF-funded Center for Embedded Networked Sensing. Her research interests include rapidly deployable and robustly operating distributed systems of physically embedded devices, with particular application to environmental monitoring. She received her PhD in computer science from the Massachusetts Institute of Technology. She is a Fellow of the IEEE, ACM, and AAAS (the American Association for the Advancement of Science). Contact her at the UCLA Computer Science Dept., 3531H Boelter Hall, Box 951596, Los Angeles, CA 90095-1596; destrin@cs.ucla.edu; <http://cens.ucla.edu/Estrin>.



Leonidas Guibas is a professor of computer science at Stanford University, where he heads the Graphics Laboratory's Geometric Computation group. He's also part of the university's AI Laboratory, Bio-X Program, and ICME Program. His research interests include computational geometry, geometric modeling, computer graphics, computer vision, robotics, ad hoc communication and sensor networks, and discrete algorithms. He received his PhD from Stanford University. He's an ACM fellow. Contact him at the Computer Science Dept., Gates Bldg., Stanford Univ., Stanford, CA 94305; guibas@cs.stanford.edu.



Aman Kansal is a researcher in Microsoft Research's Networked Embedded Computing group, where he's developed a shared and programmable mobile phone sensor network. His research interests include sensor-actuator networks, pervasive computing, wireless and mobile computing, and embedded systems. He received his PhD in electrical engineering from the University of California, Los Angeles. He's a member of the IEEE and ACM. Contact him at Microsoft Research, 1 Microsoft Way, Redmond, WA 98052; kansal@microsoft.com.



Samuel Madden is an assistant professor in the Electrical Engineering and Computer Science Department and a member of the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. His research interests include all areas of database systems, modeling and statistical techniques for value prediction and outlier detection in sensor networks, high-performance database systems, and networking and data processing in intermittently connected environments. Contact him at 32 Vassar St., Rm. 32-G938, Cambridge, MA 02139; madden@csail.mit.edu.



Jim Reich is a researcher at the Palo Alto Research Center, where he leads the Embedded Collaborative Computing Area. His research interests include collaborative signal processing for acoustic and video sensor networks, distributed tracking, resource-aware programming approaches, and control of object-mover airjet arrays and space vehicles. He received his MS in electrical and computer engineering from Carnegie Mellon University. Contact him at PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304; jreich@parc.com; www.parc.com/jreich.

2nd Ann. IEEE Int'l Conf. Pervasive Computing and Communications, IEEE CS Press, 2004, pp. 253–262.

9. B. Hull et al., "CarTel: A Distributed Mobile Sensor Computing System," *Proc. 4th Int'l Conf. Embedded Networked Sensor Systems*, ACM Press, 2006, pp. 125–138.

10. A. Talukdar, B.R. Badrinath, and A. Acharya, "Rate Adaptation

- Schemes in Networks with Mobile Hosts," *Proc. 4th Ann. ACM/IEEE Int'l Conf. Mobile Computing and Networking*, ACM Press, 1998, pp. 169–180.
11. E. Brewer et al., "The Case for Technology in Developing Regions," *Computer*, vol. 38, no. 6, 2005, pp. 25–38.
12. A. Meliou et al., "Data Gathering Tours in Sensor Networks," *Proc. 5th Int'l Conf. Information Processing in Sensor Networks*, ACM Press, 2006, pp. 43–50.
13. A. Howard, L. Parker, and G. Sukhatme, "Experiments with Large Heterogeneous Mobile Robot Team: Exploration, Mapping, Deployment, and Detection," *Int'l J. Robotics Research*, vol. 25, no. 5, 2006, pp. 431–447.
14. S. Madden et al., "TinyDB: An Acquisitional Query Processing System for Sensor Networks," *ACM Trans. Database Systems*, vol. 30, no. 1, 2005, pp. 122–173.
15. A. Festag, R. Schmitz, and R. Baldessari, "InterVehicle Communication—Improving Driver Control and Vehicle Safety," demo presented at *11th Ann. Int'l Conf. Mobile Computing and Networking (MOBI-COM 05)*, 2005.
16. Y. Zhang et al., "IceDB: Continuous Query Processing in an Intermittently Connected World," to be published in *Proc. 23rd IEEE Int'l Conf. Data Eng. (ICDE 07)*, IEEE Press, 2007.
17. A. Ghosh et al., "Variable-Resolution Information Dissemination," *Proc. 2nd Ann. IEEE Communications Society Conf. Sensor and Ad Hoc Communications and Networks*, IEEE Press, 2005, pp. 487–497.
18. S.B. Eisenman and A.T. Campbell, "Ski-Scape Sensing," *Proc. 4th Int'l Conf. Embedded Networked Sensor Systems*, ACM Press, 2006, pp. 401–402.
19. A. Krause et al., "Near-Optimal Sensor Placements: Maximizing Information while Minimizing Communication Cost," *Proc. 5th Int'l Conf. Information Processing in Sensor Networks*, IEEE Press, 2006, pp. 2–10.
20. A. Deshpande and S. Madden, "MauveDB: Supporting Model-Based User Views in Database Systems," *Proc. 2006 ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 2006, pp. 73–84.
21. H.W. Sorenson, *Kalman Filtering: Theory and Application*, IEEE Press, 1985.
22. S. Arulampalam et al., "A Tutorial on Particle Filters for On-Line Non-Linear/Non-Gaussian Bayesian Tracking," *IEEE Trans. Signal Processing*, vol. 50, no. 2, 2002, pp. 174–188.
23. W. Li et al., "Robust Statistical Methods for Securing Wireless Localization in Sensor Networks," *Proc. 4th Int'l Symp. Information Processing in Sensor Networks*, IEEE Press, 2005, pp. 91–98.
24. M. Li et al., "Swing & Swap: User-Centric Approaches towards Maximizing Location Privacy," *Proc. 5th ACM Workshop Privacy in Electronic Society*, ACM Press, 2006, pp. 19–28.
25. P. Kamat et al., "Enhancing Source-Location Privacy in Sensor Network Routing," *Proc. 25th IEEE Int'l Conf. Distributed Computing Systems (ICDCS 05)*, IEEE CS Press, 2005, pp. 599–608.
26. C. Ozturk, Y. Zhang, and W. Trappe, "Source-Location Privacy in Energy-Constrained Sensor Network Routing," *Proc. 2nd ACM Workshop Security of Ad Hoc and Sensor Networks*, ACM Press, 2004, pp. 88–93.
27. P. Golle, D.H. Greene, and J. Staddon, "Detecting and Correcting Malicious Data in VANETS," *Proc. 1st ACM Workshop Vehicular Ad Hoc Networks*, ACM Press, 2004, pp. 29–37.
28. H. Chan and A. Perrig, "Security and Privacy in Sensor Networks," *Computer*, vol. 36, no. 10, 2003, pp. 103–105.
29. V.S. Verykios et al., "State-of-the-Art in Privacy Preserving Data Mining," *ACM SIGMOD Record*, vol. 33, no. 1, 2004, pp. 50–57.
30. D. Agrawal and C.C. Aggarwal, "On the Design and Quantification of Privacy Preserving Data Mining Algorithms," *Proc. 20th ACM SIGMOD-SIGACT-SIGART Symp. Principles of Database Systems*, ACM Press, 2001, pp. 247–255.
31. R. Agrawal and R. Srikant, "Privacy-Preserving Data Mining," *Proc. 2000 ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 2000, pp. 439–450.
32. A. Evfimievski, "Randomization in Privacy Preserving Data Mining," *ACM SIGKDD Explorations Newsletter*, vol. 4, no. 2, 2002, pp. 43–48.
33. Z. Huang, W. Du, and B. Chen, "Deriving Private Information from Randomized Data," *Proc. 2005 ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 2005, pp. 37–48.
34. K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets," *Proc. 2003 Conf. Applications, Technologies, Architectures, and Protocols for Computer Communications*, ACM Press, 2003, pp. 27–34.
35. J. Rule, "Toward Strong Privacy: Values, Markets, Mechanisms, and Institutions," *Univ. of Toronto Law J.*, vol. 54, no. 2, 2004, pp. 183–225.
36. T.Y. Levin, U. Frohne, and P. Weibel, *CTRL [SPACE]: Rhetorics of Surveillance from Bentham to Big Brother*, MIT Press, 2002.
37. R. Agrawal et al., "Hippocratic Databases," *Proc. 28th Int'l Conf. Very Large Databases*, Morgan Kaufmann, 2002, pp. 143–154.
38. H. Hochheiser, "The Platform for Privacy Preference as a Social Protocol: An Examination within the US Policy Context," *ACM Trans. Internet Technology*, vol. 2, no. 4, 2002, pp. 276–306.
39. M. Gruteser and D.D. Grunwald, "Anonymous Usage of Location-Based Services through Spatial and Temporal Cloaking," *Proc. 1st Int'l Conf. Mobile Systems, Applications, and Services*, ACM Press, 2003, pp. 31–42.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.