

# Using Tomography for Ubiquitous Sensing<sup>\*</sup>

Stacy Patterson<sup>1</sup>  
sep@cs.ucsb.edu

Bassam Bamieh<sup>2</sup>  
bamieh@engineering.ucsb.edu

Amr El Abbadi<sup>1</sup>  
amr@cs.ucsb.edu

<sup>1</sup>Department of Computer Science, <sup>2</sup>Department of Mechanical Engineering  
University of California, Santa Barbara  
Santa Barbara, CA 93106, USA

## ABSTRACT

By embedding sensors in mobile devices, it is possible to exploit the ubiquitous presence of these devices to construct applications for large-scale sensing and monitoring of environmental phenomena. To this end, we present Environmental Tomography, a novel approach in which mobile devices participate in the collection of aggregate sensor readings along roads or sidewalks, and these aggregates are used to reconstruct an estimate of the contaminant distribution throughout a region. We demonstrate how our data collection process preserves user location privacy and is robust to sensor and location reading errors. We also show how the estimation process can be formulated as a convex optimization problem that incorporates the physical dynamics of the phenomenon of interest. We study the performance of Environmental Tomography using various road network layouts and realistic models of pollution. Results indicate that estimates generated from path aggregates are of comparable accuracy to estimates generated from significantly greater numbers of individual sensor readings.

## Categories and Subject Descriptors

J.9 [Computer Applications]: Mobile Applications—*Pervasive computing, Wireless sensor networks*

## General Terms

Design, Measurement

## Keywords

sensor networks, pollution, privacy, convex optimization

## 1. INTRODUCTION

The ubiquitous presence of mobile phones presents an opportunity for large-scale sensing of environmental phenomena. This can easily be achieved with a participatory

<sup>\*</sup>This work was funded in part by NSF grant IIS 02-23022.

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM GIS '08, November 5-7, 2008, Irvine, CA, USA

Copyright 2008 ACM ISBN 978-1-60558-323-5/08/11 ...\$5.00.

sensing approach, where pollution sensors are embedded in the mobile phones, and phone users contribute by collecting readings of pollution levels. Since modern mobile phones are GPS-enabled, it is possible to record not only the concentration of a sensed phenomenon, but also the exact location of the sensor reading. This data and location information can be used to create detailed models of the pollution distribution.

The success of such an application depends on the willingness of users to participate in it, and therefore, the application should be implemented in a way that addresses user concerns. Mobile devices have limited power and storage capacities, and users may only be willing to contribute a small fraction of these resources to a sensing application. So, a sensing application should have a small per-device resource footprint. Additionally, users move in independent, unpredictable patterns, and a sensing application cannot expect that a device take sensor readings in predefined locations, nor can it expect that sensor readings can be taken at every point in a region. Therefore, the application should be able to use data collected from typical user routes to generate the pollution estimates. Finally, by directly reporting location information along with sensor readings, users are forced to reveal their locations. Users may not be willing to participate in a program that requires them to divulge this private location information. Any successful ubiquitous sensing application implementation must preserve user location privacy.

In this work, we present *Environmental Tomography*, a novel approach to environmental sensing and spatial data modeling that meets the challenges presented by the mobile users and the mobile network. Tomography and tomographic reconstruction have long been used in medical imaging techniques such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) [10]. For example, in a CT scan, two dimensional X-rays are taken in multiple directions, and a three-dimensional image is reconstructed from these two dimensional projections. Similarly, Environmental Tomography consists of two phases, a data collection phase in which the network of mobile devices computes aggregate values, or projections, of sensor readings along roads and sidewalks, and a reconstruction phase in which the aggregates are used to generate an estimate of the distribution of the sensed phenomenon.

We utilize a data collection approach that is designed to exploit the global characteristics of the mobile network. We assert that, while individual user mobility patterns are not necessarily predictable, it is possible to predict the pat-

terns of the network as a whole. Specifically, during certain times of the day, namely rush hour, there is a high density of users, and therefore of mobile devices, along roads and sidewalks. Aggregate sensor readings can be collected by routing a query message from device to device along these roads. Messages can be passed by the devices themselves using short-range communication protocols such as 802.11 or bluetooth, or they can be transmitted using more powerful communication radios in the vehicles in which the users are traveling. Each participating device only takes a few sensor readings, and the data collection process does not require any of the device’s permanent storage. Therefore, the individual device resource requirements are low. Additionally, since sensor readings are aggregated, there is no need to report the location of any individual participating device. Thus, the privacy of user locations is preserved.

To generate an estimate of the underlying phenomenon from the projections, we develop a tomographic reconstruction technique that can be posed as a convex optimization problem with an objective function that takes the physics of the underlying phenomenon into account. This approach enables us to efficiently generate accurate estimates of the phenomenon using the limited number of aggregates that are available. We verify the validity of our approach through extensive simulations using physically accurate models of environmental phenomena.

The remainder of this paper is organized as follows. In Section 2, we present our system model and architecture. Section 3 describes the data collection process. In Section 4, we present the tomographic reconstruction technique, formalize the tomographic reconstruction problem, and describe the solution method. In Section 5, we describe the pollution models used to validate our solution, and in Section 6, we give experimental results. We conclude in Section 7 with a discussion of related work and future research directions.

## 2. SYSTEM ARCHITECTURE

The goal of Environmental Tomography is to generate an estimate of the distribution of a physical phenomenon over a finite two-dimensional region, the *sensing region*. For example, if an accident in a factory results in a hydrogen chloride leak, Environmental Tomography can be used to determine the amount of the gas that is present throughout a city and identify dangerous areas. In this initial work, we restrict our interest to the sensing and estimation of ground level pollution. We also restrict our study to phenomena that diffuse slowly with respect to the data collection process, i.e. we assume that the distribution is relatively static for the duration of the execution of the queries used in an instance of tomographic reconstruction. These assumptions accurately describe many substances including carbon monoxide, sulfur dioxide, hydrogen chloride, and ammonia.

A high level view of the system architecture is shown in Fig. 1. The system is comprised of the mobile devices in the region to be modeled. We assume that these devices are equipped with environmental sensors as well as GPS capabilities, and that the devices have bluetooth or 802.11 radios that allow them to communicate with other nearby devices. We place no requirements on the movement pattern or availability of any individual device, and the set of devices that participates in the system can change over time. While we make no assumptions about individual behavior, we do make use of an observation about the global behavior

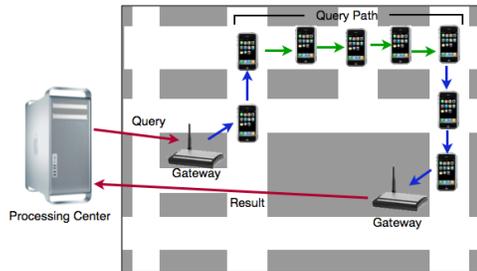


Figure 1: System Architecture

of the mobile network. Specifically, we assume that the network of mobile devices is dense along roads and sidewalks, as is observed in urban areas during rush hour traffic, so that the local communication channels in the devices form a connected network along these paths.

In addition to the mobile devices, our system contains several infrastructure devices. Gateway machines are installed throughout the region. These gateways have bluetooth or 802.11 radios so that they can communicate with mobile devices that are within range. They also have reliable communication channels to the processing center. The processing center is the control center of the system. It initiates communication with the gateways to request sensor information from the mobile network, it collects the query results, and it performs tomographic reconstruction to generate the estimate of the sensed phenomenon. The data collection process is described in the next section. Tomographic reconstruction is explained in Section 4.

## 3. DATA COLLECTION

As in medical imaging, our data collection process consists of computing projections along paths through the sensing region. In a CT scan, each projection is computed by taking an X-ray of the body. The result of each X-ray is a 2-dimensional image, where each point in the image corresponds to the density of the 3-dimensional region through which the X-ray passed. X-rays are taken along a 180 degree axis of rotation, and the collection of 2-dimensional projects are used to reconstruct a 3-dimensional image of the body.

In Environmental Tomography, data collection involves computing one dimensional projections of the two dimension region we wish to estimate. Each one dimensional projection is the sum of sensor readings taken at fixed points along a specified path, and each of the devices along the path contribute to the data collection process by taking a sensor reading, adding the reading to a partial sum stored in the message, and passing the message to the next device along the path. We explain this process in more detail below.

The path along which the query message will travel and the locations at which sensor readings should be taken are specified by a *query*. The simplest example of a query specification is the specification for a straight line query path,  $(start, end, \delta)$ . The path is defined by *start* and *end* GPS coordinates. The  $\delta$  value defines the distance between locations at which reading should be taken. We call these locations *sampling points*. It is not necessary that the query path be a straight line, and in fact, more complex query path trajectories can be defined using multiple line segments or

### Data Collection Protocol

1. The processing center creates a query and sends the query message to a gateway that is closest to the starting coordinates of the query path.
2. The gateway introduces the message into the wireless network by sending it to a nearby mobile device, and the query is routed to the starting point of the query path using a greedy routing protocol.
3. When the message reaches the starting point, a sensor reading is taken.
4. The query is routed along the query path using trajectory-based forwarding. While a device has the query message, if it is on (or near enough to) a specified sampling point, it takes a sensor reading and adds this value to the partial sum in the message.
5. When the message reaches the query path end coordinates, it is routed to the nearest gateway using greedy routing.
6. The query result is sent back to the processing center.

Figure 2: The data collection protocol

a parametric representation. Given the sophisticated mapping and directions systems available for mobile devices, we anticipate that it will also be possible to specify queries on a higher semantic level, for example,

*Take a sensor reading every 25 meters along Interstate-95 between mile marker 100 and mile marker 110.*

When a device receives a query message, it can determine the exact location that readings should be taken and the path along which the message is intended to travel.

The data collection process is initiated by the processing center which creates query specifications for each of the paths on which data is to be collected. In general, it is advantageous to collect data from as many paths as possible, so the processing center should issue queries for any roads along which the mobile network is expected to be dense. Results for each of these queries are computed by devices in the mobile network using the data collection protocol given in Fig. 2. The flow of the query message through the network is illustrated by the arrows in Fig 1.

To route messages within the mobile network, we rely upon two well-established geographic routing techniques for ad-hoc networks, greedy routing and trajectory-based forwarding. In geographic or position-based routing [25, 9], rather than establishing routes in the network, each device keeps track of its location and the locations of its neighbors, and it uses this location information to make all routing decisions. A greedy routing protocol, such as Greedy Perimeter Stateless Routing (GPSR) [12] can be used to send a message to a destination that is specified by a set of coordinates. Each device forwards the message to the neighbor

that is closest to the destination. Since these are greedy algorithms, the protocols also provide fallback mechanisms to route around holes and avoid local minima. In trajectory-based forwarding protocols [17, 3, 28], the message does not have a destination, but rather has a specified trajectory or path. Each device selects the next-hop so as to keep the message as close to the trajectory as possible. As both of these protocols are stateless, mobile devices need only to know the location of the gateways in order to participate in the data collection process. A participating device can store a list of gateway locations that is periodically updated by the application software, or the device can use a secure location based service, for example [15], to retrieve the location of the nearest gateway without revealing its location to the service.

The data collection process has a number of benefits:

- *Scalability:* The majority of communication takes place between neighboring mobile devices, and therefore, communication load on the processing center and gateways is minimized.
- *Low Per-Device Resource Usage:* The data collection process requires minimal computational power on each device and does not require any permanent storage on the devices. Each device only takes a few sensor readings and sends and receives a small number of query messages.
- *Exploits User Mobility Patterns:* Sensor readings are only taken where the mobile network is naturally dense, along crowded streets and sidewalks.
- *Location Privacy:* Since sensor readings are aggregated, locations of individual readings are not recorded. The processing center does not require any information about individual device locations.

While location information can be kept secret from the processing center, to route messages within the mobile network, devices must reveal their locations to neighboring devices as well as to the gateways. One could argue that this poses no additional risk to privacy because neighboring devices in the wireless network are physically near each other, and therefore locations are not secret. If stronger privacy guarantees are desired, the data collection process can also employ a location-based routing scheme that preserves the anonymity of participants [21].

Once the processing center has received the query results from the gateways, it uses tomographic reconstruction to generate an estimate of the distribution of the sensed phenomenon from the aggregate measurements. We explain this process in the next section.

## 4. TOMOGRAPHIC RECONSTRUCTION

In this section, we formally define the tomographic reconstruction problem as a convex optimization problem. We then show how the problem can be converted to a form that can be easily and efficiently solved with readily available convex optimization software.

### 4.1 Problem Formulation

We begin by formalizing the data collection process. Let  $S \subset \mathbb{R}^2$  be the two-dimensional region over which the estimate of the sensed phenomenon is to be generated. Let

$C : S \rightarrow \{\mathbb{R}^+ \cup 0\}$  be the underlying physical distribution, i.e. for every point  $(x, y) \in S$ ,  $C(x, y)$  is the concentration of the phenomenon at that point. Initially, we assume that all sensor readings are accurate, so if  $(x, y)$  is a sampling point,  $C(x, y)$  is also the value recorded by the sensor at that sampling point. Later, we will show how this assumption can be relaxed.

Aggregate sensor readings are collected along  $P$  query paths. Let  $(x_i^j, y_i^j)$  denote the  $j^{\text{th}}$  sampling point along the  $i^{\text{th}}$  query path. The data collection process can be expressed as a linear operator  $\mathcal{A}$  on the linear space of functions  $C : S \rightarrow \{\mathbb{R}^+ \cup 0\}$ . For any density distribution  $C$ , the vector of aggregate measurements  $m$  is given by

$$\mathcal{A}(C) := \begin{bmatrix} \sum_j C(x_1^j, y_1^j) \\ \vdots \\ \sum_j C(x_P^j, y_P^j) \end{bmatrix} = \begin{bmatrix} m_1 \\ \vdots \\ m_P \end{bmatrix}. \quad (1)$$

Tomographic reconstruction involves solving the inverse problem; given the vector of aggregate measurements  $m$ , find an estimate of the underlying distribution  $\hat{C}$ , that is consistent with the query results

$$\mathcal{A}(\hat{C}) = m. \quad (2)$$

Since  $\mathcal{A}$  is a linear operator, solving for the estimate  $\hat{C}$  in Equation 2 amounts to solving a system of linear equations.

In Environmental Tomography, there are severe restrictions in both the choice and number of paths due to the location of roads and sidewalks. Therefore the linear system is *underdetermined*; there is not enough information to yield a unique solution to Equation 2, and, in fact, there are infinitely many distributions that satisfy the equation.

In the case of an underdetermined system, one must define some criterion that identifies the optimal solution from the set of feasible solutions. A standard least squares solution finds the solution with minimum norm [14], where the  $L^2$  norm of the estimate  $\|\hat{C}\|_2$  is minimized. However, in the case of a physical phenomenon such as a plume of sulfur dioxide or a cloud of gaseous ammonia, there is no compelling argument for minimizing the  $L^2$  norm. We instead propose a minimization criterion that takes into consideration the physical dynamics of the phenomenon.

The motivation for our optimization criterion is based on the observation that distributions of gases and pollutants typically follow the dynamics of the advection-diffusion equation [19],

$$\begin{aligned} \frac{\partial}{\partial t} C(x, y, t) &= \left( u_x \frac{\partial}{\partial x} + u_y \frac{\partial}{\partial y} \right) C(x, y, t) \\ &+ \left( D_x \frac{\partial^2}{\partial x^2} + D_y \frac{\partial^2}{\partial y^2} \right) C(x, y, t). \end{aligned} \quad (3)$$

Advection, or dispersion due to wind, is determined by the wind velocities in the direction of each axis,  $u_x$  and  $u_y$ .  $D_x$  and  $D_y$  are the diffusion constants for each direction. We assume that diffusion is uniform in both directions, and therefore  $D_x = D_y =: D$ . If the correct diffusion constants, wind velocities, and initial conditions (source location and quantity) were known, it would be possible to design optimization criteria that give accurate consideration to both the advection and diffusion processes. However, it is not practical to assume that this information will be available, and our experimental results indicate that even small errors in

these parameters will yield large errors in the estimate generated by the optimization process. So, instead we choose an optimization criterion that depends only on the diffusion dynamics. In other words, we assume a simplified model of the distributed based only on the diffusion equation,

$$\frac{\partial}{\partial t} C(x, y, t) = D \left( \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) C(x, y, t).$$

Let  $\Delta := \left( \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)$  be the two-dimensional Laplace operator. If we assume that a diffusive distribution  $C$  is quasi-static, then the time-variation term  $\partial C / \partial t$  is expected to be small. This assumption consequently implies that  $\Delta C$  should be small. We measure the “size” of  $\Delta C$  using its  $L^2$  norm

$$\|\Delta C\|_2 := \langle \Delta C, \Delta C \rangle = \langle C, \Delta^2 C \rangle. \quad (4)$$

This is a non-negative quadratic form on the space of all functions  $C : S \rightarrow \mathbb{R}$ , and we set up the tomographic reconstruction problem so as to minimize this form.

Note that the optimization of Equation 4, subject to the constraints in Equation 2, is a weighted least squares problem with a well-understood solution. However, this solution may not necessarily satisfy the physical constraint that the estimate  $\hat{C}$  is non-negative at every point in space. To remedy this problem, we explicitly incorporate non-negativity constraints into the optimization problem. Experimentally, the addition of the non-negativity constraint results in large improvements in the accuracy of the estimates.

We now summarize the above discussion by formally stating the tomographic reconstruction problem as a convex optimization problem over the space of functions  $C : S \rightarrow \mathbb{R}$ ,

$$\begin{aligned} &\text{minimize} && \|\Delta C\|_2 \\ &\text{subject to} && \\ &&& \mathcal{A}(C) = m \\ &&& \forall (x, y) \in S, \quad C(x, y) \geq 0, \end{aligned}$$

where  $\mathcal{A}$  is the linear measurement operator given by Equation 1.

Since  $\Delta$  is a linear operator and  $\|\cdot\|_2$  is a norm,  $C \mapsto \|\Delta C\|_2$  is a convex function [2]. The linear constraints  $\mathcal{A}(C) = m$ , and the positive cone constraint  $\forall (x, y) \in S, \quad C(x, y) \geq 0$  both yield convex constraint sets. The net constraint set, the intersection of the two, is thus a convex set. The overall problem is therefore a convex optimization problem which implies the existence of a global minimum solution [2].

## 4.2 Computational Procedure

To convert the convex optimization problem to a form that can be solved using a conventional tool such as Matlab, we first must translate the problem to a discrete form. This conversion requires a discrete representation of the distribution  $C$  and discrete approximations of the objective function and constraints.

Consider a  $M \times N$  grid overlaying the sensing region  $S$ . Let  $\mathbf{C} = [c_{(i,j)}]$  be the  $M \times N$  matrix where each entry  $c_{(i,j)}$  is the value of the distribution  $C$  at the point corresponding to the  $i^{\text{th}}$  row and  $j^{\text{th}}$  column in the grid over  $S$ .  $\mathbf{C}$  is a discrete approximation of  $C$ , and the goal of our reconstruction problem is now to find a matrix  $\hat{\mathbf{C}}$  that estimates  $\mathbf{C}$ .

In order to discretize the system of constraints given by Equation 2, we require a discrete linear operator that approximates  $\mathcal{A}$ . For simplicity, we assume that the sampling

points are a subset of the grid points in the discrete representation of the region. In the next section, we show how to accommodate the relaxation of this assumption. Let  $\mathbf{c}$  be the vector that is formed by concatenating the rows of  $\mathbf{C}$ . We define the  $MN \times P$  matrix  $A$ , where  $P$  is the number of query paths as follows. Each row in  $A$  corresponds to a query path. The value in the  $k^{\text{th}}$  column of the row is 1 if the  $k^{\text{th}}$  component of  $\mathbf{c}$  is included in the query result. Otherwise, the value is 0. Using this  $A$  matrix, the system of linear constraints is approximated by the equation

$$A\mathbf{c} = \mathbf{m}.$$

For the objective function, we use a standard second order finite difference approximation of  $\Delta$ , which is defined as follows

$$\Delta C(i, j) \approx (\mathbf{c}_{i,j-1} - 2\mathbf{c}_{i,j} + \mathbf{c}_{i,j+1}) + (\mathbf{c}_{i-1,j} - 2\mathbf{c}_{i,j} + \mathbf{c}_{i+1,j}).$$

Letting  $\mathbf{L}$  be the matrix representation of this approximation, we have

$$\Delta C \approx \mathbf{L}\mathbf{c}.$$

The  $L^2$  norm on  $\Delta C$  is approximated by the standard Euclidean norm on the vector  $\mathbf{L}\mathbf{c}$ .

The discrete version of the convex optimization problem is then given by the following,

$$\begin{aligned} & \text{minimize} && \|\mathbf{L}\mathbf{c}\|_2 \\ & \text{subject to} && \\ & && A\mathbf{c} = \mathbf{m} \\ & && \mathbf{c}_k \geq 0 \quad \text{for } k = 1 \text{ to } (MN). \end{aligned}$$

This is a convex optimization problem with  $MN$  variables, one for each grid point,  $P$  linear equality constraints, one for each query path, and  $MN$  inequality constraints. Using this form, we can efficiently perform our tomographic reconstruction with any convex optimization solver.

In the limit of an infinitely fine grid, the solution to the discrete problem is equivalent to the solution to the continuous problem. Therefore, one should use as fine a grid as the available computing system can support.

### 4.3 Incorporating Errors

The problem formulation and solution outlined in the previous sections rely on the assumptions that all sensor readings are accurate and that sampling points correspond exactly to grid points in the discretized sensing region. In this section, we show how our solution can be slightly reformulated to account for errors introduced by inaccurate sensors, as well as location errors caused by inaccurate GPS information and those created by the discretization of the problem.

In the original formulation, the measurement taken by the sensor at position  $(x_i^j, y_i^j)$  is exactly equal to the value of the distribution at that point,

$$m_i^j = C(x_i^j, y_i^j).$$

If the sensor is not accurate, the measurement will differ from the actual value by some perturbation. We model this inaccuracy by a zero mean random variable with variance  $\sigma_d$ , call it  $d_i^j$ ,

$$m_i^j = C(x_i^j, y_i^j) + d_i^j.$$

Additionally, if the GPS service does not report the location of the device accurately, the sensor reading will be taken at a location that differs slightly from the expected location for the reading. Similarly, if the sampling point does not correspond exactly to a point in the discretized sensing region, the reading will be attributed to a grid point with a location that differs from the location of the actual reading. In both cases, the processing center assumes that a reading was taken at a point  $(x, y)$ . However, the actual location of the reading is  $(x + \delta_x, y + \delta_y)$ . The relationship between the value of the distribution at this expected point and the actual point is given by the first order expansion,

$$C(x + \delta_x, y + \delta_y) \approx C(x, y) + \delta_x \frac{\partial}{\partial x} C(x, y) + \delta_y \frac{\partial}{\partial y} C(x, y).$$

If we assume that the first derivative of the distribution is bounded, as is expected with diffusive substances, then the error introduced by inaccurate location information can also be modeled by additive noise. Therefore, the errors introduced by both the sensor and location inaccuracies can be modeled by a single additive noise term, yielding,

$$\bar{m}_i^j = C(x_i^j, y_i^j) + u_i^j,$$

where  $u_i^j$  is also a zero mean random variable with variance  $\sigma_u^j$ .

Since each individual reading is inaccurate, the sums collected along the query paths are also inaccurate. Therefore, we can no longer expect the underlying distribution to be consistent with the measurement vector  $\mathbf{m}$ . Instead, we select the solution that best fits the measurement vector while also still optimizing for the original criterion. The new optimization problem is then to find a solution that minimizes a combination of the original criterion  $\|\mathbf{L}\mathbf{c}\|_2$  and the weighted sum of squared residuals. Let  $J$  be the diagonal matrix where the diagonal entry in row  $i$  is the number of measurements taken along query path  $i$ . The convex optimization problem can then be stated as

$$\begin{aligned} & \text{minimize} && \alpha \|\mathbf{L}\mathbf{c}\|_2 + \beta \|J^{-1} (A\mathbf{c} - \bar{\mathbf{m}})\|_2 \\ & \text{subject to} && \\ & && \mathbf{c}_k \geq 0 \quad \text{for } k = 1 \text{ to } (MN). \end{aligned}$$

By selecting different values for  $\alpha$  and  $\beta$ , it is possible to tune the optimization problem based on the dynamics of the phenomenon and the size of the errors. For error variances that are comparable to those produced using currently available GPS and sensor technologies, we have found that setting  $\alpha = \beta = 1$  gives the best results for our experimental settings.

## 5. POLLUTION MODELS

One major challenge to the development of any application for ubiquitous environmental sensing is the evaluation of the accuracy of estimates generated by the application. Ideally, this evaluation can be done by testing the application using real world pollution data. However, there are not yet any readily-available data sets that give concentrations of pollutants at the level of spatial granularity comparable to the granularity of the estimates produced by Environmental Tomography. Therefore, we evaluate our approach on synthetic pollution data generated using physically accurate models.

As stated earlier, diffusive pollutants move through the atmosphere according to the advection-diffusion equation (3). Different solutions to the advection-diffusion equation can be used to model different types of pollution. In this work, we consider two such models. The first is a one-time emission from a point source, also called a *puff*, which models phenomena such as chemical spill from an industrial plant. The second is a continuous emission from a point source, also called a *plume*. This model can be used to simulate emissions from a factory smoke stack. We describe both models in detail below.

### Puff Model

Our first model is a simple one-time emission model where the effects of wind are negligible. For an emission of  $Q$  grams of pollutant at the point  $(x_0, y_0)$  at time  $t = 0$  the solution to the advection-diffusion equation is

$$C(x, y, t) = \frac{S}{4\pi Dt} e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{4Dt}}. \quad (5)$$

$C(x, y, t)$  gives the concentration in  $g/m^3$  at the point  $(x, y)$  at  $t$  seconds after the release of the pollutant.  $D$  is the diffusion coefficient which defines how quickly the substance diffuses in a given medium [8]. While this model is a simplification of the behavior of point source pollutants, it can be used as a baseline test of the validity of the optimization criteria and the effects of choosing different query paths.

### Gaussian Plume Model

The second pollution scenario we consider is a steady-state solution to the advection-diffusion equation, the Gaussian Plume Model [26, 19]

$$C(x, y) = \frac{S}{\pi u \sigma_y \sigma_z} \left( e^{-\frac{(y-h)^2}{2\sigma_y^2}} + e^{-\frac{(y+h)^2}{2\sigma_y^2}} \right) e^{-\frac{y^2}{2\sigma_z^2}}. \quad (6)$$

This solution models the downwind concentration of pollutant generated by a continuous point source such as a factory smoke stack. The model is parameterized by the height of the stack  $h$  (m), the wind speed  $u$  (m/s), and source strength  $S$  (g/s). The plume spread is characterized by the plume coefficients  $\sigma_y$  and  $\sigma_z$ , which are based on observed data and depend on atmospheric conditions, wind speed, and distance from source. The Gaussian Plume Model is one of the air quality models employed by the U.S. Environmental Protection Agency (EPA).

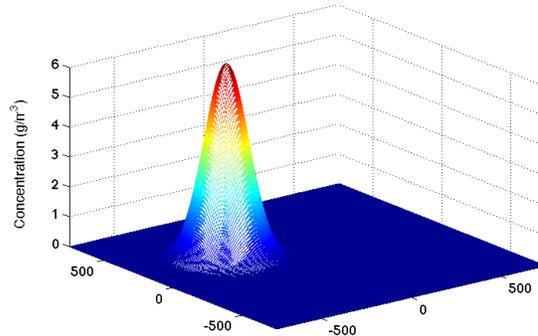
## 6. SIMULATIONS

In this section, we give results for simulations of Environmental Tomography using several representative pollution models. Our simulation system is implemented in Matlab. We use CVX [5], a Matlab-based modeling tool for convex programming, and we use the SDPT3 solver [22]. Each sensing region is discretized by a  $300 \times 300$  grid. This system is a full realization of the exploratory demonstration proposed in [20].

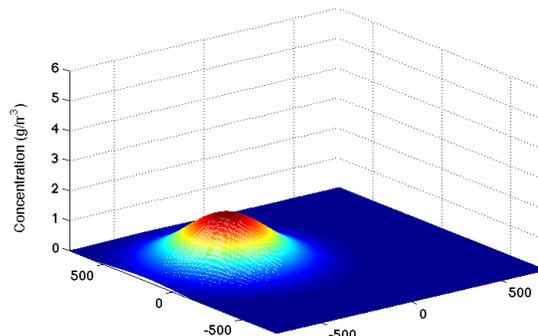
We first present the pollution distribution models and then describe our evaluation strategy in detail. Finally, we present some summary results.

### 6.1 Simulation Scenarios

For our simulations, we study one puff distribution generated by a 400 kg emission of anhydrous ammonia, which



(a) Puff at  $2.5 \times 10^8$  seconds after release.



(b) Puff at  $1 \times 10^9$  seconds after release.

**Figure 3: Puff Distributions**

has a diffusion constant of  $1.96 \times 10^{-5}$ . We use distribution examples from two different time periods,  $2.5 \times 10^8$  seconds and  $1 \times 10^9$  seconds after the release of the chemical. The sensing region is a 1500 square meter region centered at  $(0,0)$ , and the source of the emission is the northwest corner of the region, at the coordinates  $(-300, 200)$ . The distributions, as generated by Equation 5, are shown in Figures 3(a) and 3(b).

We also present results for one plume distribution. We again use a 1500 square meter sensing region. The source of the plume is a smoke stack in the west of the region, located at  $(0,0)$ . The height of the stack is 10 meters, and the rate of emissions is 50 g/s. We model the plume using conditions on a cloudy evening with a westerly wind of 2 m/s. We use the plume coefficients defined by the EPA's Industrial Source Complex Air Quality Model [6]. The resulting plume distribution is shown in Fig. 4.

In order to investigate the effects of road network layout on the performance of Environmental Tomography, we use two different road network scenarios. The first is a random road network where road start and end points are selected uniformly at random from within the sensing region. Each road is a line segment that connects the start and end points. The second network type is a biased network where start and points are selected according to gaussian distributions centered at  $(-375, 375)$  and  $(375, -375)$ , respectively, with a standard deviation of 200 meters. For each road network, we select  $\delta$  so that a network with  $R$  roads has  $10R$  total

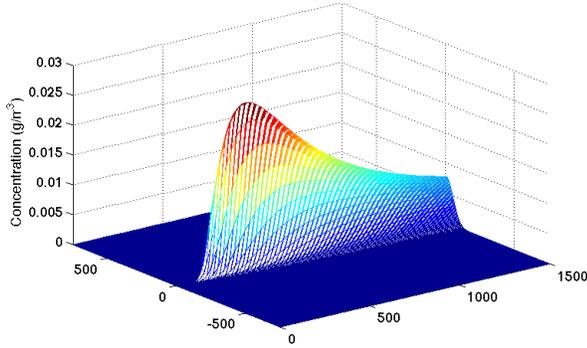


Figure 4: Plume

sampling points, i.e a road network with 10 roads has 100 sampling points, a road network with 20 roads has 200 sampling points, etc. By choosing  $\delta$  in this way, the number of sampling points assigned to each road in a particular network is proportional to the road length.

For both types of road networks, we evaluate Environmental Tomography using simulations in which sensor readings and locations are correct and simulations that include both sensor reading and location errors. To model location errors, we use sampling points that deviate from the expected sampling points by a gaussian random perturbation with a standard deviation of 3 meters. To model sensor errors, we then perturb reading at the erroneous location by a random value that is selected according to a gaussian distribution with a standard deviation of 5% of the mean value of the distribution. We present results for these various scenarios in the next section.

## 6.2 Simulation Results

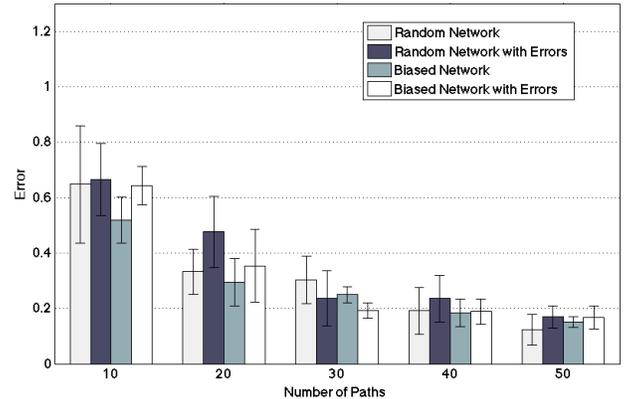
To quantify the performance of Environmental Tomography, we measure the difference between the reconstructed estimate  $\hat{C}$  and the discretization of the original underlying distribution  $C$ , using the following metric,

$$\text{Err}(\hat{c}) := \frac{\|c - \hat{c}\|_2}{\|c\|_2},$$

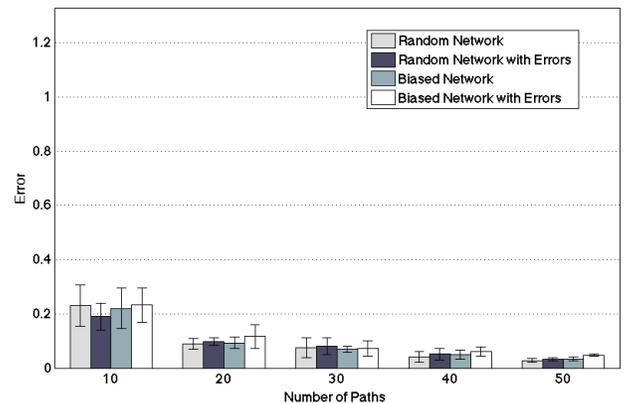
where  $c$  and  $\hat{c}$  are the vectorizations of  $C$  and  $\hat{C}$ .

In Fig. 6.2, we present the errors for estimates generated for each of the three pollution models. We show results for random and biased networks both with and without sensor and location error. Each bar gives the mean error and standard deviation of 10 simulations, using 10 different randomly generated road networks.

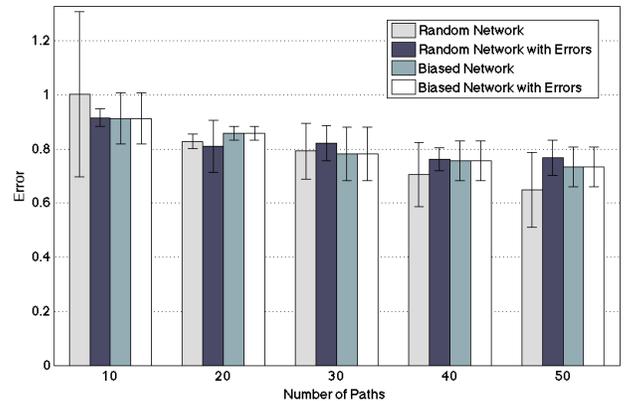
Figure 5(a) shows the estimation error for the puff distribution at  $2.5 \times 10^8$  seconds after the pollutant is released. As expected, increasing the number of roads along which aggregates are collected increases the amount of information available in the reconstruction and therefore reduces the error in the estimation process. It is interesting to note that as the number of paths increases, the addition of more paths provides less benefit to the estimation accuracy. Much greater improvement is seen when the number of paths is increased from 10 to 20 than is seen in an increase from 40 to 50 paths. We also note that the performance in random and biased networks is comparable. The simulations that



(a) Puff at  $2.5 \times 10^8$  seconds after release.

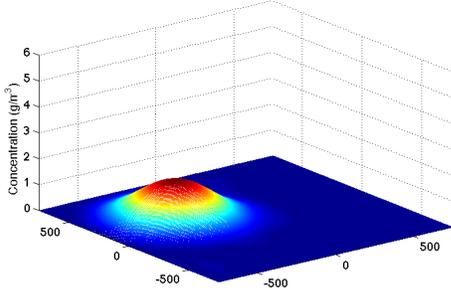


(b) Puff at  $1 \times 10^9$  seconds after release.

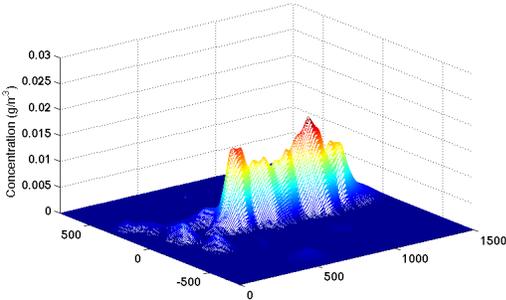


(c) Plume.

Figure 5: Mean estimation error and standard deviation for different road network types and number of roads.



(a) Estimate of puff distribution at  $1 \times 10^9$  seconds after release.



(b) Estimate of plume distribution.

**Figure 6: Estimates generated from a 50 road random network.**

include errors also perform comparably in most cases.

For the puff distribution at  $1 \times 10^9$  seconds after release, our technique performs very well, as shown by the results in Fig. 5(b). The mean error is significantly lower than for the previous distribution for both random and biased networks, with and without errors. The standard deviation is also small for all networks, which suggests that Environmental Tomography can be used to create accurate estimates of this type of distribution under a large variety of road networks. In Fig. 6(a), we show the three dimensional plot of an estimate of the puff distribution at  $1 \times 10^9$  seconds after release that was generated using a 50 road random network without errors. This distribution is almost identical to the original distribution in Fig. 3(b), giving further evidence of the accuracy of our approach in modeling this type of distribution. Our intuition is that, because this distribution is relatively smooth and spread out over the sensing region, the data collection process gathers information about a large portion of the distribution. This increased information may yield more accurate estimates.

The estimation errors for the plume distribution are shown in Fig 5(c). The errors for this distribution are larger than for both puff distributions. Additionally, the estimation process does not appear to be sensitive to road network layout, as the standard deviation is small for all network sizes, nor to errors, as the performance between the simulations with and without errors is similar. We also present an example estimate of the plume distribution in Fig. 6(b). This estimate is generated from a 50 road random network without errors. While there are obvious differences between the es-

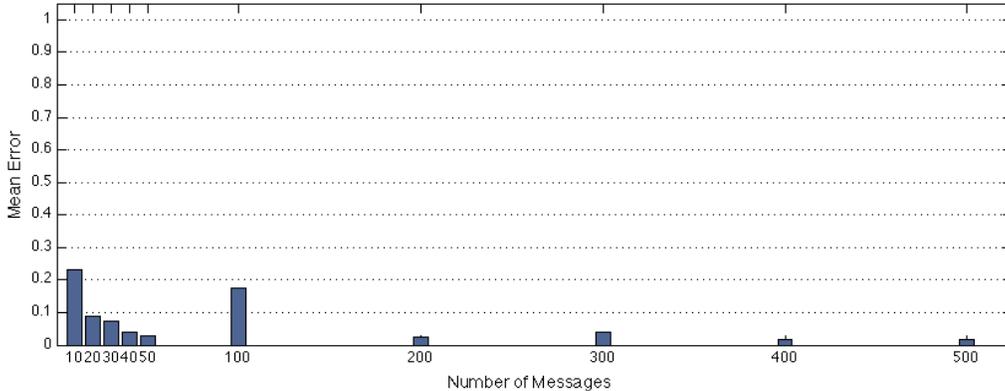
timate distribution and the original distribution in Fig. 4, the estimate distribution is concentrated in the correct location and is of similar magnitude to the original. It appears that even if the estimates have a larger error, Environmental Tomography is still successful in identifying areas of high pollutant concentration from only 50 path aggregates over a region of 1500 square meters.

In plume distributions, the pollutant is highly concentrated in smaller areas of the regions, and so, for most road networks, the data collection process will “miss” a large portion of the distribution. Unless a sampling point happens to intersect the this small region of highest concentration, a large quantity of the distribution will not be included in the aggregation process. The tomographic reconstruction process may not be able to completely compensate for this missing information, and thus less accurate estimates are produced. However, any data collection process that relies on voluntary user participation will most likely suffer from the same data collection restrictions regardless of whether that technique uses aggregation to protect user privacy. Therefore, it is important to not only examine absolute errors, but also to compare the estimates with other estimates generated under similar information restrictions. We present this type of comparison in our next set of results.

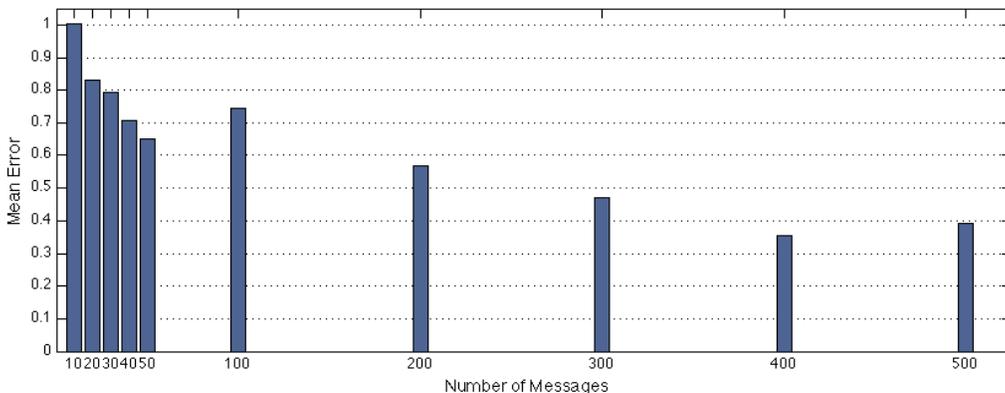
In order to compare our technique to an approach with similar information availability, we consider estimates that are generated from *all* of the available sensor readings as opposed to the aggregates used in Environmental Tomography. In this technique, each individual sensor reading is reported in its own message to the processing center, and the processing center uses the same reconstruction technique on this set of individual data points. For example, in a road network of 10 roads with 100 sampling points, each of the 100 readings is reported to the processing center. Hence, the processing center performs a reconstruction using all 100 samples instead of the 10 aggregate values used in Environmental Tomography.

In Figures 7(a) and 7(b), we present a comparison of errors for the puff distribution at  $1 \times 10^9$  seconds after release and the steady-state plume distribution. Data collection was done over random road networks for both distributions, and each bar presents the mean error of simulations on 10 different networks. These simulations do not include location and sensor errors. Results for simulations over biased networks and simulations with errors are similar.

In the figure, we show the error for estimates that are generated for each specified numbers of messages. For the cases of 10, 20, 30, 40, and 50 messages, the messages contain aggregate values over the 10, 20, 30, 40, and 50 paths. For the cases of 100, 200, 300, 400, and 500 messages, the messages correspond to each of the individual readings over the same 10, 20, 30, 40, and 50 paths. We note that, for the puff distribution, estimates generated from just 20 path aggregates are significantly more accurate than those generated from 100 individual sensor readings. Estimates generated from 40 and 50 path aggregates are comparable to those those generated from 200 individual sensor readings for both distributions. Most importantly, for the plume distribution, the errors are large for estimates generated from both the aggregates and the individual readings. These results indicate that by using aggregates instead of individual readings, we do not lose much information about the distribution, and we can retain the benefits of scalability and location privacy.



(a) Puff at  $1 \times 10^9$  seconds after release.



(b) Plume.

**Figure 7: Mean estimation errors for various numbers of messages in a random road network. Errors for 10, 20, 30, 40, and 50 messages correspond to estimates generated from path aggregates using Environmental Tomography. Errors for 100, 200, 300, 400, and 500 messages correspond to estimates generated from individual sensor readings.**

## 7. DISCUSSION AND FUTURE WORK

In this work, we have introduced Environmental Tomography, a novel approach to ubiquitous sensing and environmental modeling. Our technique is unique in that it is designed to exploit the global properties of the mobile network while also being conscious of individual user requirements, most significantly protecting the location privacy of participants. We have shown that tomographic reconstruction can be formulated as a convex optimization problem with an objective and constraints that are based on the physical properties of the underlying phenomenon. Finally, we have demonstrated the feasibility of our approach through experiments using various road networks and realistic models of environmental phenomena.

### 7.1 Related Work

Tomography and tomographic reconstruction have been applied in several computing disciplines. Network tomography has been proposed to estimate individual network link delays from end-to-end delay measurements [4, 13]. Since

the end-to-end measurement of a path varies, the system is overdetermined. The authors therefore employ tomographic reconstruction techniques that minimize the error due to the measurement variation. Similarly, tomographic reconstruction has been used for hardware [16] and software [1] analysis from end-to-end measurements

The notion of using cell phones to build large-scale sensor network has been the subject of much attention. The recent work by Kansal et al. [11] suggests using cell phone microphones and cameras as sensors and proposes an infrastructure for collecting this sensor data. There are also several ongoing research projects that focus on urban sensing and participatory sensing by using ubiquitous entities such as mobile phones and vehicles to develop pervasive sensor networks. These projects include Urban Sensing [27], Participatory Urbanism [18], SenseWeb [23], and Sensor Planet [24] and the Equator Project [7]. Our work can be seen as complimentary to these projects. It can build upon the infrastructure provided by them, and it provides novel benefits that these projects do not address. Specifically, to our

knowledge, this is the first work to propose a data collection and modeling approach that is sensitive to user privacy concerns and specifically designed for urban mobile networks.

## 7.2 Future Work

In future work, we will investigate the extension of Environmental Tomography to include additional pollution models such as multiple source and line source distributions. We also plan to extend our work to accommodate data collection over large time scales and the generation of dynamic pollution models. Finally, we will investigate the benefits of utilizing stationary sensors to supplement the data collection process and improve the distribution estimation.

## 8. REFERENCES

- [1] J. F. Bowring, A. Orso, and M. J. Harrold. Monitoring deployed software using software tomography. In *Proceedings of the ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE)*, pages 2–9, 2002.
- [2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [3] A. Capone, L. Pizziniaco, I. Filippini, and M. de la Fuente. A sift: an efficient method for trajectory based forwarding. In *2nd International Symposium on Wireless Communication Systems (ISWCS)*, pages 135–139, 2005.
- [4] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu. Network tomography: Recent developments. *Statistical Science*, 19(3):499–517, 2004.
- [5] cvx: Matlab software for disciplined convex programming. <http://www.stanford.edu/boyd/cvx/>, 2007. [Online; accessed 28-October-2007].
- [6] EPA technology transfer network, support center for regulatory atmospheric modeling. <http://www.epa.gov/scram001/dispersionindex.htm>. [Online; accessed 01-June-2008].
- [7] Equator. <http://www.equator.ac.uk>. [Online; accessed 01-June-2008].
- [8] S. J. Farlow. *Partial Differential Equations for Scientists and Engineers (Dover Books on Advanced Mathematics)*. General Publishing Company, Ltd., 1993.
- [9] H. Frey and I. Stojmenovic. On delivery guarantees of face and combined greedy-face routing in ad hoc and sensor networks. In *Proceedings of the 12th annual international conference on Mobile computing and networking (MobiCom)*, pages 390–401, 2006.
- [10] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. SIAM, 2001.
- [11] A. Kansal, M. Goraczko, and F. Zhao. Building a sensor network of mobile phones. In *Proceedings of the 6th international conference on Information processing in sensor networks (demonstration) (IPSN)*, pages 547–548, 2007.
- [12] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proceedings of the 6th annual international conference on mobile computing and networking (MobiCom)*, pages 243–254, 2000.
- [13] E. Lawrence, G. Michailidis, V. N. Nair, and B. Xi. *Frontiers in Statistics*, chapter Network Tomography: A Review and Recent Developments. Imperial College Press, July 2005.
- [14] W. Menke. *Geophysical Data Analysis: Discrete Inverse Theory*, volume 45 of *International Geophysics Series*. Academic Press, Inc., 1989.
- [15] M. F. Mokbel, C.-Y. Chow, and W. G. Aref. The new casper: Query processing for location services without compromising privacy. In *Proceedings of the 32nd International Conference on Very Large Data Bases (VLDB)*, pages 763–774, 2006.
- [16] S. Mysore, B. Mazloom, B. Agrawal, and T. Sherwood. Understanding and visualizing full systems with data flow tomography. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 211–221, March 2008.
- [17] B. Nath and D. Niculescu. Trajectory based forwarding and its applications. In *Proceedings of the 9th annual international conference on Mobile computing and networking (MobiCom)*, pages 260–272, 2003.
- [18] Participatory urbanism. <http://www.urban-atmospheres.net/ParticipatoryUrbanism/index.html>. [Online; accessed 01-June-2008].
- [19] F. Pasquill. *Atmospheric Diffusion*. Wiley, New York, 1974.
- [20] S. Patterson, B. Bamieh, and A. El Abbadi. Environmental tomography: Ubiquitous sensing with mobile devices. In *Proceedings of the IEEE 24th International Conference on Data Engineering (ICDE) (demonstration)*, pages 1560–1563, 2008.
- [21] S. Schlott, F. Kargl, and M. Weber. Short paper: Random ids for preserving location privacy. In *Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM)*, pages 415–417, 2005.
- [22] SDPT3 version 4.0 (beta) – a matlab software for semidefinite-quadratic-linear programming. <http://www.math.nus.edu.sg/mattohkc/sdpt3.html/>, 2007. [Online; accessed 28-October-2007].
- [23] Senseweb. <http://research.microsoft.com/nec/senseweb/>. [Online; accessed 01-June-2008].
- [24] Sensor planet. <http://www.sensorplanet.org>. [Online; accessed 01-June-2008].
- [25] I. Stojmenovic. Position-based routing in ad hoc networks. *IEEE Communications Magazine*, 40(7):128–134, July 2002.
- [26] O. Sutton. *Micrometeorology: A Study of Physical Processes in the Lowest Layers of the Earth's Atmosphere*. McGraw-Hill, New York, 1953.
- [27] Urban sensing. [http://research.cens.ucla.edu/projects/2006/Systems/Urban\\_Sensing/default.htm](http://research.cens.ucla.edu/projects/2006/Systems/Urban_Sensing/default.htm). [Online; accessed 01-June-2008].
- [28] M. Yuksel, R. Pradhan, and S. Kalyanaraman. An implementation framework for trajectory-based routing in ad hoc networks. *Ad Hoc Networks*, 4(1):125–137, January 2006.