Environmental Tomography: Modeling the Environment with Mobile Phones *

Stacy Patterson†, Bassam Bamieh‡, and Amr El Abbadi†
†Department of Computer Science, ‡Department of Mechanical Engineering
University of California Santa Barbara
Santa Barbara, CA 93106
{sep, amr}@cs.ucsb.edu, bamieh@engineering.ucsb.edu

Abstract

The coupling of sensors with mobile phones, which are ubiquitously available and location aware, opens the door to the creation of applications for pervasive sensing and detailed spatial modeling of environmental phenomena. In order to ensure widespread participation of mobile users, these applications must have limited per-device resource requirements, must place no expectations on individual user behaviors, and must be sensitive to user privacy concerns.

In this work, we introduce Environmental Tomography, *a novel approach to environmental sensing and spatial data modeling that meets the challenges of the mobile network. Environmental Tomography consists of two phases, a data collection phase in which the network of mobile devices computes aggregate values 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. The data collection process is robust to the dynamics of the mobile network and protects the location privacy of the participating device users. The reconstruction phase can be posed as a convex optimization problem with an objective function that takes the physics of the underlying phenomenon into account to produce accurate estimates from the limited available data. We verify the validity of our approach through extensive simulations using physically accurate models of environmental phenomena.*

1. Introduction

The coupling of sensors with mobile phones, which are ubiquitously available and location aware, presents an opportunity for the creation of large-scale sensing applications. Since modern mobile phones are GPS-enabled, it is possible to record not only the concentration of the sensed phenomenon, but also the exact location of the sensor reading. This location information opens the door for the creation of detailed spatial models of the physical data distribution.

However, the mobile computing platform also presents several challenges. 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. 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. Finally, by 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.

In this work, we introduce *Environmental Tomography*, a novel approach to environmental sensing and spatial data modeling that meets the challenges of 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) [9]. For example, in a CT scan, two dimensional X-rays are taken in multiple directions, and these two dimensional projections are combined to reconstruct a three-dimensional image. 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 patterns of the network as a whole. Specifically, during certain times of the day, namely rush hour, there is a high density

*The authors thank Stephen Boyd for a useful discussion on convex optimization.

of users, and therefore of mobile devices, along roads and sidewalks. Aggregates sensor readings can be collected by routing a query message from device to device along these roads via the local communication channels using bluetooth or 802.11. 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 our tomographic recontruction technique, formalize the tomographic reconstruction problem, and describe the solution method. In Section 5, we give experimental results. Finally, we conclude in Section 6 with a discussion of related work and future research directions.

2. System Model and 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 consider phenomena that are heavier than air. Specifically, we assume that as the substance of interest diffuses into the region, it remains concentrated close to the ground. Therefore, the distribution can be modeled by considering ground level concentration only. 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, carbon dioxide, sulfur dioxide, hydrogen cholride, and ammonia.

A high level view of the system architecture is shown in Fig. 1. The system consists 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 behavoir, we do make use of an observation about the global behavior 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.

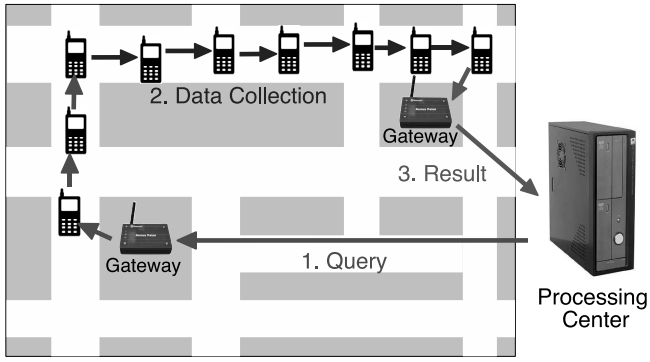

Figure 1. System Architecture

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 detail 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 of the phenomena along paths through the sensing region. We call these paths *query paths*. In medical imaging this projection is the integral of a continuous function along the path. In Environmental Tomography, the projection is the sum of sensor readings taken at specified *sampling points* along the query path.

Field	Description
start	coordinates of starting point of the query path
end	coordinates of end point of the query path
delta	distance between sensor readings along the path
sum	sum of collected sensor readings

Table 1. Query Message

For simplicity, we illustrate the data collection approach for straight line paths, though it is possible to use any arbitrary path shape so long as the processing center can determine the points at which the sensor readings in a given aggregate were taken.

The data collection process is shown by the arrows in Fig. 1. To initiate a projection, or *query*, the processing center creates a query message. This message fully specifies the query path and the location of the sampling points. For example, the query message specification for a straight line path is given in Table 1. The path is defined by *start* and *end* GPS coordinates. The *delta* field gives the distance between sensor readings that should be taken along the path. These three values completely define the location of all sampling points for the path. More complex query path trajectories can be defined using multiple line segments or a parametric representation. In addition to the path definition, the query message has a *sum* field that is updated as the readings are taken.

The processing center sends the query message to a gateway that is close to the query starting coordinates, and the gateway introduces the message into the wireless network by sending it to a nearby mobile device as shown in Step 1 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 [22, 8], 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) [11] 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 [16, 3, 24], 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.

In the data collection process, after the query has entered the mobile network, a greedy routing protocol is used to route the message to the start coordinates. The query is routed along the query path using trajectory-based forwarding, as illustrated by Step 2 in Fig. 1. 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 sum field in the message. When the message reaches the query path end coordinates, it is routed to the nearest gateway using greedy routing, as shown in Step 3. The query message containing the sum is then sent back to the processing center.

In order to participate in the data collection process, mobile devices need only to know the location of the gateways. The devices do not need any *a priori* information about the queries, and they do not need to store anything after they are done processing the query message. A participating devices can store list of gateway locations that is periodically updated by the application software, or the device can use a secure location based service [14] to retrieve the location of the nearest gateway without revealing its location to the service. Each device will mostly only participate in one query, requiring it to take only a few sensor readings. Therefore, the per device resource requirements are very low

With respect to the location privacy of the mobile devices, there are two aspects to consider, the processing center and the other devices in the mobile network. The processing center does not need to know anything about the device locations. The query results are completely decoupled from the devices that participate in the queries. However, in order 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 techniques that preserve the anonymity of the participants in location based-routing schemes [18].

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 we want to generate an estimate sensed phenomenon. Let $f : S \to \mathbb{R}$ be the underlying physical distribution, i.e. for every point $(x, y) \in S$, $f(x, y)$ is the value of the phenomenon at that point. We assume that the sensor readings are accurate, so if (x, y) is a sampling point, $f(x, y)$ is also the value recorded by the sensor at that sampling point.

Let (x_j^i, y_j^i) denote the j^{th} sampling point along the i^{th} query path. The data collection process can be expressed as a linear operator $\mathcal{A}$ on the linear space of functions $f : S \to \mathbb{R}$. Given any density distribution f, the vector of aggregate measurements along the P specified query paths is

$$\mathcal{A}(f) := \begin{bmatrix} \sum_j f(x_j^1, y_j^1) \\ \vdots \\ \sum_j f(x_j^P, y_j^P) \end{bmatrix}. \quad (1)$$

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

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

Note that m is the vector of query results whose components, m_i are each the sum of the sensor readings at the points along the i^{th} query path.

Since $\mathcal{A}$ is a linear operator, solving for the estimate $\hat{f}$ in (2) amounts to solving a system of linear equations after using some appropriate discretization scheme. In Environmental Tomography, there are severe restrictions in both the choice and number of paths due to location of roads and sidewalks. Thus, there is not enough information to yield a unique solution to Equation 2. In fact, there are infinitely many distributions that satisfy the equation. Such a system is called *underdetermined*.

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 [13], where the L^2 norm of the estimate $||\hat{f}||$ is minimized. However, in the case of a physical phenomenon such as a plume of of sulfur dioxide or a cloud of gaseous ammonia, there is no compelling argument for minimizing the L^2 norm. We instead propose to minimize a quadratic form that better captures some of the underlying physical phenomenon.

Our proposed optimization criterion is motivated by the observation that distributions of gases and pollutants typically follow diffusive dynamics. Any diffusive process in two dimensions satisfies the diffusion equation [7]

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

where $\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)$ is the two-dimensional Laplace operator which we denote by Δ. Assuming approximately steady-state conditions, *i.e.* the distribution f is quasi-static, the time variation term $\partial f / \partial t$ is then expected to be small. This assumption consequently implies that $\triangle f$ should be small. We measure the "size" of $\triangle f$ using its L^2 norm

$$|| \triangle f||_2 := < \triangle f, \triangle f > = < f, \triangle^2 f > . \quad (3)$$

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

Note that the optimization of Equation 3 subject to the constraints in Equation 2 is actually a weighted least squares problem whose solution is well understood. However, such solutions may not necessarily satisfy the physical constraint that the solution $\hat{f}$ is non-negative at every point in space. To remedy this problem, we add these non-negativty constraints explicitly into the optimization problem. Experimentally, the addition of the non-negativity constraint results in large improvements in the accuracy of the estimates. Examples of this improvement are given in Section 5.

We now summarize the above discussion by formally stating the Environmental Tomography problem as a convex optimization problem over the space of functions $f : S \to \mathbb{R}$

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

where Δ is the 2 dimensional Laplacian, and $\mathcal{A}$ is the linear measurement operator given by Equation 1.

Since Δ is a linear operator and $\|.\|$ is a norm, then $f \mapsto \|\Delta f\|$ is a convex function [2]. The linear constraints $\mathcal{A}(f) = m$, and the positive cone constraint $\forall (x, y) \in S, \;\; f(x, y) \geq 0$ 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 minimium solution [2].

There are many readily available convex optimization solvers, but in order to make use of these tools, we must convert the problem to a discrete representation. We explain this process in the next section.

4.2. Discrete Representation

To convert the convex optimization problem to a form that can be solved using a conventional tool such as MATLAB, we must define a discrete representation of the distribution f, and discrete approximations of the objective function and constraints.

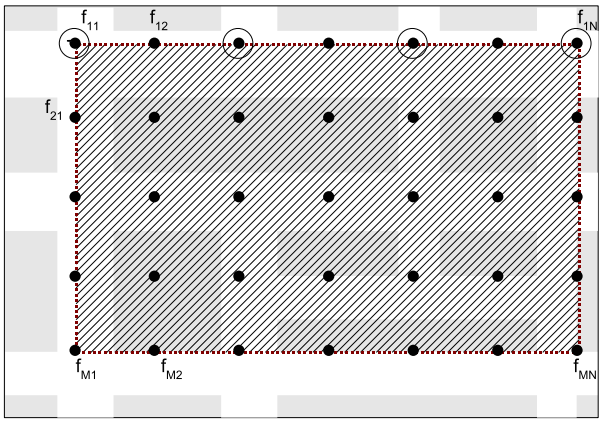

Figure 2. Discretization of Estimation Region

We consider the tomographic reconstruction problem over a region S, which is shown by the shaded rectangle in Figure 2. The region can be represented in discrete form by an $M \times N$ grid, as illustrated in the figure. Let (i, j) denote the point in S corresponding to the i^{th} row and j^{th} column of the grid. We represent the distribution f as $M \times N$ matrix $F = [\mathsf{f}_{i,j}]$, with $\mathsf{f}_{i,j}$ equal to the value of the phenomenon f at the grid point (i, j). The goal of tomographic reconstruction is then to find a matrix $\hat{F}$ that estimates F.

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 practice, if a sampling point does not correspond exactly to a grid point, we represent the sampling point by the closest grid point.

Let f be the vector form of the matrix F that is formed by concatenating the rows. We define the $(0, 1)$ matrix A that approximates the $\mathcal{A}$ operator as follows. Each row in A corresponds to a query path. Each column corresponds to a grid point. For each row, the value in the k^{th} column is 1 if the k^{th} component of f is included in the query result. Otherwise the value is 0. For example, if the circled points in Fig. 2 represent sampling points in a query path, the corresponding row in the A matrix is given by

$$\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & \cdots & 0 \end{bmatrix}.$$

The result of the query along the path is the sum of the values from components 1, 3, 5, and 7 of the f vector. Using this A matrix, the system of linear constraints is approximated by the equation

$$A\mathsf{f} = m.$$

For the objective function, we need a discretization of the Laplace operator Δ. We use a standard second order finite difference approximation which is defined over the grid as follows

$$\begin{aligned} \Delta f(i,j) \;\approx\; & (\mathsf{f}_{i,j-1} - 2\mathsf{f}_{i,j} + \mathsf{f}_{i,j+1}) \\ & + (\mathsf{f}_{i-1,j} - 2\mathsf{f}_{i,j} + \mathsf{f}_{i+1,j}). \end{aligned}$$

We denote by L be the matrix representation of this approximate operator that acts on the vectorized grid f, i.e.

$$\Delta f \approx L\mathsf{f}.$$

The L^2 norm on Δf is approximated by the standard Euclidean norm on the vector $L\mathsf{f}$.

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

$$\begin{aligned} &\text{minimize} \quad ||L\mathsf{f}|| \\ &\text{subject to} \\ &\quad A\,\mathsf{f} = \mathsf{m} \\ &\quad \mathsf{f}_k \geq 0 \quad \text{for } k = 1 \text{ to } (MN) \end{aligned}$$

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 inequaility 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.

5. Experimental Results

In this section, we present experimental results of Environmental Tomography. We use two models of environmental phenomena, a one-time emission that represents a scenario such as a chemical leak, and a continuos emission scenario that models phenomena such as factory emissions. For each model, we assume that diffusion is the sole process responsible for the change of the distribution. We also assume that the substances diffuse very slowly with respect to the data collection process, and therefore the distribution is static during throughout the query execution. As a future research direction, we plan to expand our work to include environmental models that incorporate dispersion, buoyancy, and reactivity, and to enhance the tomographic reconstruction process to account for the time variation of the distribution of the phenomena during data collection.

We explore data collection and tomographic reconstruction using a network of evenly spaced horizontal and vertical query paths so that we can study the effect of the number of paths and sampling points on the accuracy of the reconstruction. In all of these experiments, the sensing region is discretized by a 400×400 grid. We also include results

using real-world road networks from Midtown Manhattan, New York City and downtown San Francisco. In these experiments, the region is discretized by a 200×200 grid.

We calculate the error of an estimate $\hat{\mathsf{f}}$ using the following metric

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

where $||\mathsf{f}||$is the vectorized version of the discrete representation of the physical phenomenon, as described in Section 4.2.

To solve the convex optimization problems, we use CVX [5], a Matlab-based modeling system for convex programming, and we use the SDPT3 solver [19].

5.1. One-Time Emission

We first model the one-time emission of a contaminant as it diffuses over time. The ground level concentration (g/m^2) of the substance at the coordinates (x, y) at time t is given by the following equation

$$C(x, y, t) = \frac{s_0}{4\pi Dt} e^{\frac{-\left(x^2+y^2\right)}{4Dt}}.$$

s_0 is the quantity emitted at time 0 in grams, and D is the diffusion coefficient which defines how quickly the substance diffuses in a given medium [7].

For the one-time emission scenario, we consider an anhydrous ammonia leak. Ammonia is used as a refrigerant in large-scale industrial processes, and exposure to ammonia can cause eye and lung damage and even death. In the following experiments, we model a leak of 500 kg of ammonia. The diffusion coefficient of ammonia in air is

$$D = 1.96 \cdot 10^{-5}\ m^2/s.$$

We use a sensing region of 1000 square meters; the source of the emission is in the center of the region. We perform Environmental Tomography on the phenomenon at three different times, 10^8. 10^9, and 10^{10} seconds after the emission. The physical distributions are shown in Fig. 3. The diffusive behavior is evident in the change of the distribution over time.

To investigate the role that the number of paths and sampling points play in the accuracy of the tomographic reconstruction, we experiment with three different query path configurations. Each configuration consists of an equal number of horizontal and vertical query paths evenly spaced over the sensing region. In the first case, we use 10 horizontal and 10 vertical paths with 10 sampling points on each path. There are 100 sampling points in total and each sampling point belongs to two query paths. In the second configuration, we use 20 horizontal and 20 vertical paths with 20 sample points per path, totaling 400 sampling points. The third configuration uses 40 horizontal and 40 vertical paths with 40 sampling points per path for a total of 1600 sampling points.

The estimation error for each configuration at each of the three times is shown in Table 2. Overall, the error is low, and our technique appears to be most accurate for the distribution at 10^9 seconds. Interestingly, the 20 path configuration results in estimates that are nearly as accurate as the 40 path configuration for the distribution at 10^8 seconds, and the 20 path configuration is slightly more accurate than the 40 path configuration for 10^9 seconds. These results suggest an a non-obvious relationship between the amount of information available and the quality of the reconstruction estimate.

In Figure 4, we show the a graphical representation of the results of Environmental Tomography using 20 horizontal and vertical paths. For all three distribution times, the estimate closely resembles the original distribution. However, the estimate for 10^{10} seconds has large inaccuracies at the boundary of the region. Boundary conditions often play a significant part in the quality of tomographic reconstruction estimates, and it may be possible incorporate additional boundary constraints to improve the estimate accuracy.

Table 2 also gives the error estimates for tomographic reconstruction without the non-negativity constraint. For this experiment, we again used the path configuration with 20 horizontal and vertical paths. The error is significantly larger when the constraint is omitted, and this effect is seen for all three distribution times. The importance of the non-negativity constraint is even more evident in the graph of the estimate distribution. In Figure 5, we show the estimate for 10^8 seconds using the 20 horizontal and vertical paths without the non-negativity constraint. This result is clearly not an accurate estimate of the distribution in Figure 3(a).

5.2. Continuous Emission

The second model of environmental phenomena that we consider is a continuous emission scenario. This model can be used to describe a factory or processing plant that constantly emits a byproduct into the atmosphere. The ground level concentration of the substance at the point (x, y) at time t is given by

$$C(x, y, t) = \int_0^t \frac{s(\tau)}{4\pi D(t-\tau)} e^{\frac{-\left(x^2+y^2\right)}{4D(t-\tau)}} d\tau,$$

where $s(t)$ is a function that describes the amount released per unit time [7].

In the following experiments, we use an emission of sulfur dioxide (SO_2) at a rate of 1 gram per second, equivalent to approximately 35 tons per year. SO_2 is produced by the combustion of petroleum and coal and is the primary cause

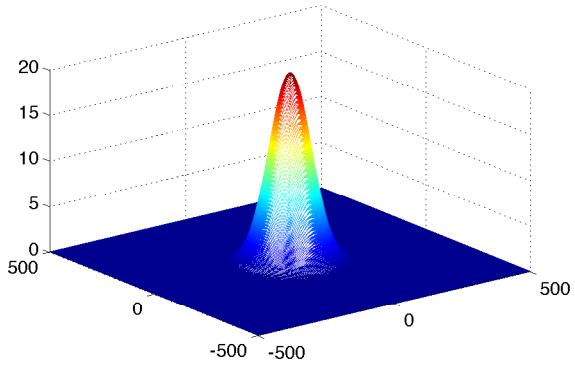

(a) 10^8 seconds

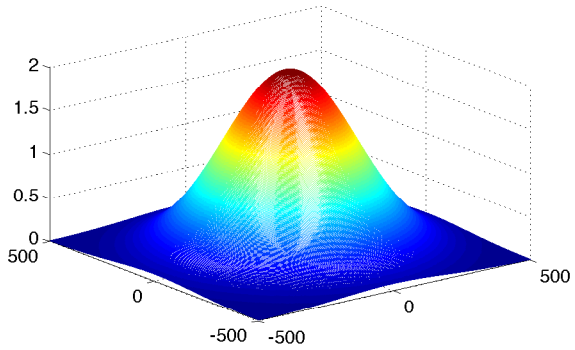

(b) 10^9 seconds

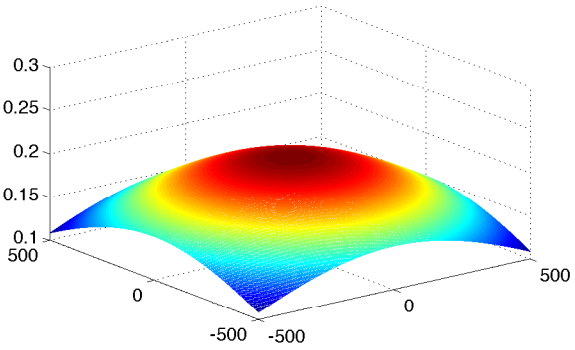

(c) 10^{10} seconds

Figure 3. Physical distribution of one-time emission of 500 kg of anhydrous ammonia

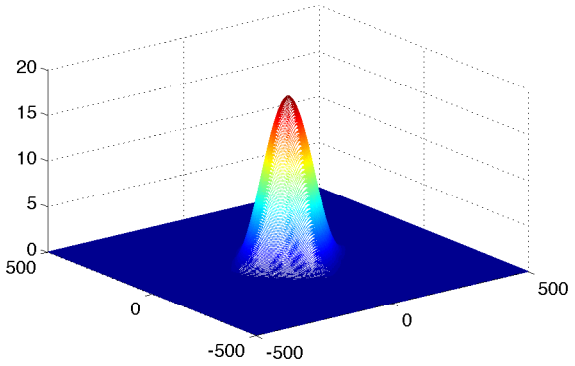

(a) 10^8 seconds

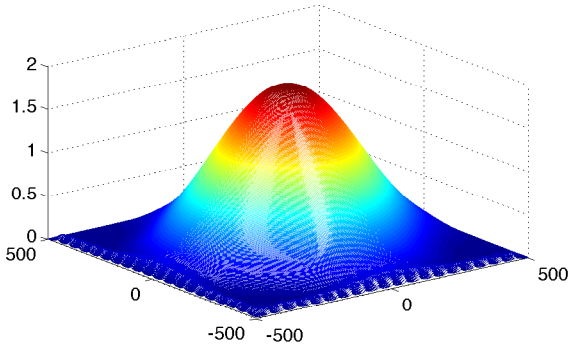

(b) 10^9 seconds

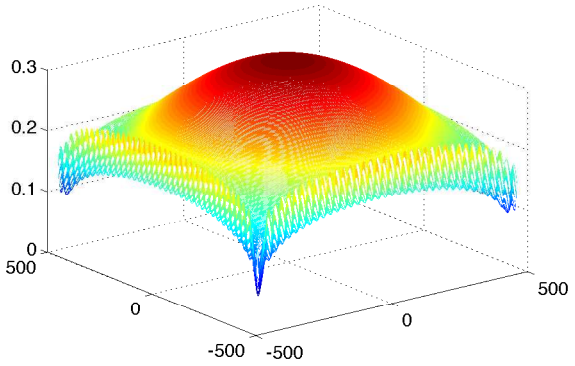

(c) 10^{10} seconds

Figure 4. Estimate of one-time emission of 500 kg of anhydrous ammonia

Table 2. Estimation error for one-time emission of 500 kg of anhydrous ammonia

Number of Horiz./Vert. Paths	10^8 seconds	10^9 seconds	10^{10} seconds
10	0.1602	0.0842	0.7686
20	0.0980	0.0777	0.3184
40	0.0946	0.0789	0.1428
20 (with negative values)	1.5062	0.2732	0.6141

of acid rain. The diffusion coefficient of SO_2 in air is

$$D = 1.22 \cdot 10^{-5}\, m^2/s.$$

We again use a sensing region of 1000 square meters, and the source of the emission is in the center of the region.

Table 3 gives the estimation error for the distribution at 10^7, 10^8, and 10^9 seconds. We use the same query path configurations as used in the one-time emission scenario. The estimation errors are larger than the errors of the estimates of the one-time emission, and unlike in the one-time emission, increasing the number of paths and sampling points improves accuracy in all cases. In Figure 6, we show the original distribution at 10^8 seconds as well as the estimate generated from the 20 horizontal and vertical query paths. While the shape of the estimate is similar to the original, the peak of the estimate distribution is not as sharp as in the original distribution. Our intuition is that we can improve the estimation process by incorporating additional constraints into the optimization problem.

Figure 7 shows the estimate of the continuous emission distribution using 40 horizontal and vertical paths, where the non-negativity constraint is not used. As in the single emission scenario, the removal of this constraint yields an obviously inaccurate estimate.

5.3. Real-World Road Networks

In this section, we give preliminary results of Enviromental Tomography over real world road networks. We use two different city maps, a map of Midtown Manhattan, New York City and a map of downtown San Francisco, California. The query paths are shown in Figures 8(a) and 9(a). For each road map, the units of measure are degrees of latitude and longitude relative to the source of the emission, which is in the center of the region. The size of each sensing region is about 5000 square meters. The roads are modeled by line segments, with 17 query paths in the Manhattan road network and 30 query paths in the San Francisco road network. For Manhattan, the sampling delta is just under one meter, which results in 1174 sampling points. For San Francisco, the sampling delta is about five meters, resulting in 222 sampling points. In order to form discrete representations of the road networks, we approximate each region by a 200×200 grid, and for each sampling point, we represent it by the closest grid point. The discretized approximation of the sampling points for each map are shown in Figures 8(b) and 9(b).

In Fig. 10, we show the physical distribution and the results of Environmental Tomography for a one-time emission of 500 kg of anhydrous ammonia at 10^9 seconds after the emmission. For both networks, the estimates closely resemble the physical distribution. Even though the Manhattan estimate was generated from a much larger number of sampling points, the estimation errors for both networks are similar. The error for Manhattan is approximately 0.3305, and for San Francisco, the error is approximately 0.3618. These results suggest that the layout of the roads plays a more significant role in the accuracy of the estimation process than the choice of sampling delta. In future work, we will further explore the relationship between the layout of the roads and the accuracy of the estimates.

6. Discussion and Future Work

In this work, we have introduced Environmental Tomography, an approach to ubiquitous sensing and environmental modeling. Our approach is unique in that it is designed to exploit the global properties of the mobile network while also being conscious of individual user requirements. 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 two realistic models of environmental phenomena.

6.1. 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, 12]. 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 reconstruc-

Table 3. Estimation error for continuos emission of 1 g/s of SO_2

Number of Horiz./Vert. Paths	10^7 seconds	10^8 seconds	10^9 seconds
10	7.0897	2.7524	4.1993
20	3.1262	0.9171	2.7422
40	1.0985	0.3553	0.3560

tion has been used for hardware [15] 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. [10] 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 [23], Participatory Urbanism [17], SenseWeb [20], and Sensor Planet [21] and the Equator Project [6]. Our work can be seen as complimentary to these projects. It can be built upon the infrastructure provided by them, and it provides novel benefits that these projects do not address. Specifically, to our knowledge, we are 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.

6.2. Future Work

We have demonstrated the feasibility of Environmental Tomography using a simple diffusion model of the physical phenomena. In future work, we will expand our approach to more complex environmental models that incorporate wind dispersion, buoyancy, atmospheric conditions, and reactivity. It is our belief that it is possible to define optimization criteria that take all of these factors into consideration in order to generate accurate estimates. We will also investigate tomographic reconstruction techniques that can be used to model dynamic distributions. Other topics of interest include optimal query path selection, the role of mobile network density, and the effects of noisy sensor readings and GPS inaccuracies.

References

[1] J. F. Bowring, A. Orso, and M. J. Harrold. Monitoring deployed software using software tomography. In *Proceedings of theACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE 2002)*, pages 2–9, Charleston, SC, November 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 2005 (ISWCS2005)*, 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] Equator. http://www.equator.ac.uk, 2007. [Online; accessed 28-October-2007].

[7] S. J. Farlow. *Partial Differential Equations for Scientists and Engineers (Dover Books on Advanced Mathematics)*. General Publishing Company, Ltd., 1993.

[8] H. Frey and I. Stojmenovic. On delivery guarantees of face and combined greedy-face routing in ad hoc and sensor networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 390–401, New York, NY, USA, 2006. ACM.

[9] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. SIAM, 2001.

[10] A. Kansal, M. Goraczko, and F. Zhao. Building a sensor network of mobile phones. In *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks (demonstration)*, pages 547–548, 2007.

[11] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *MobiCom '00*, pages 243–254, 2000.

[12] 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.

[13] W. Menke. *Geophysical Data Analysis: Discrete Inverse Theory*, volume 45 of *International Geophysics Series*. Academic Press, Inc., 1989.

[14] 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 International Conference on Very Large Data Bases, VLDB 2006*, pages 763–774, 2006.

[15] 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 (to appear)*, March 2008.

[16] B. Nath and D. Niculescu. Trajectory based forwarding and its applications. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 260–272, 2003.

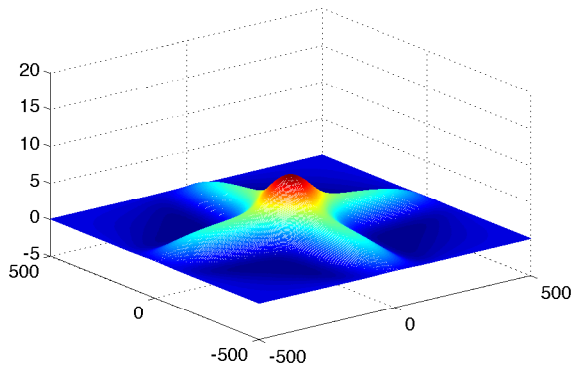

Figure 5. Estimate of ammonia leak at 10^8 seconds using 40 query paths with 20 sampling points per path, without non-negativity constraint

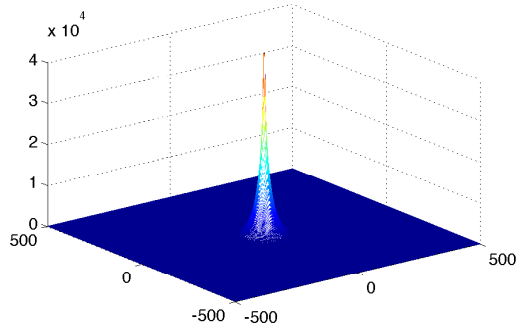

(a) Physical distribution

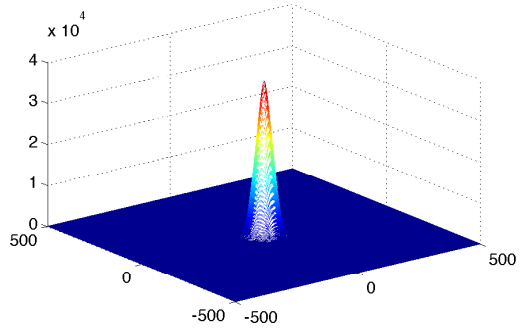

(b) Estimate with 20 horiz./vert. paths

Figure 6. Continous emission of 1 g/s of SO_2 at 10^8 seconds

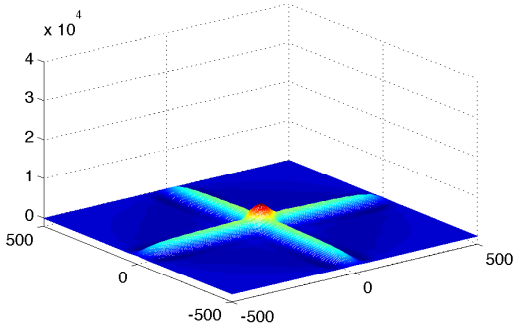

Figure 7. Continous emission of 1 g/s of SO_2 at 10^8 seconds. Estimate with 20 horiz./vert. paths, without non-negativity constraints

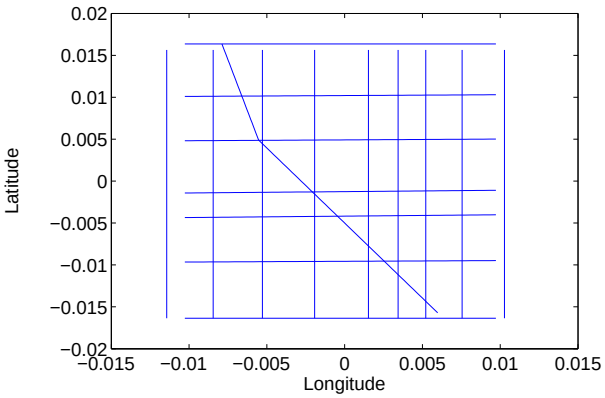

(a) Road map

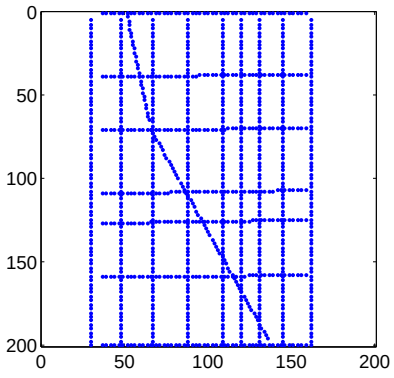

(b) Grid representation of sampling points

Figure 8. Road map and sampling points for Midtown Manhattan, New York City

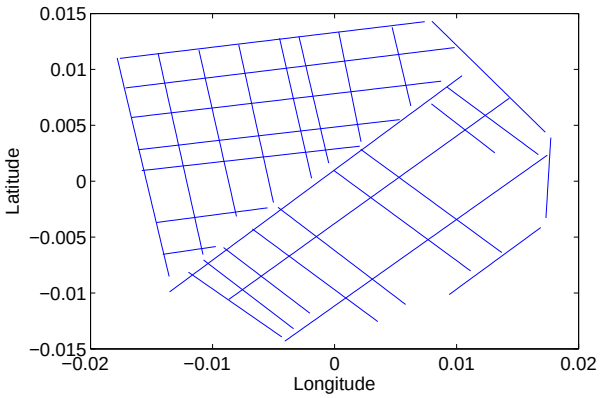

(a) Road map

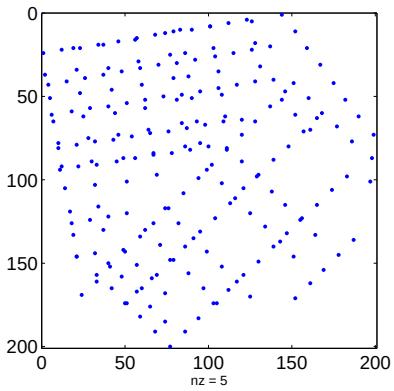

(b) Grid representation of sampling points

Figure 9. Road map and sampling points for San Francisco, California

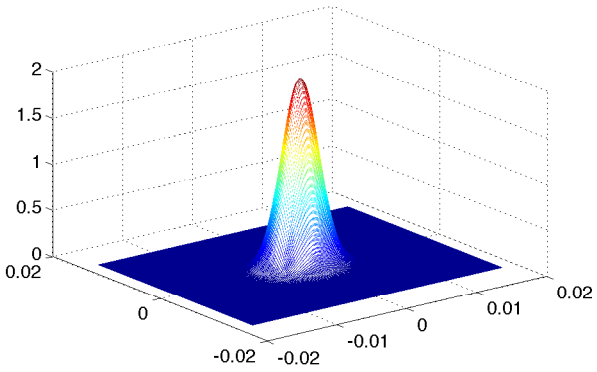

(a) Physical distribution

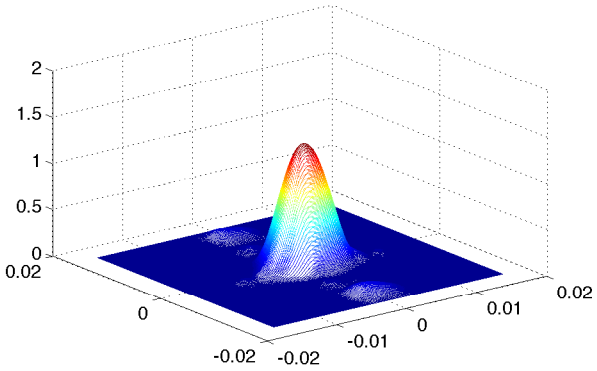

(b) Estimate generated using Manhattan road map

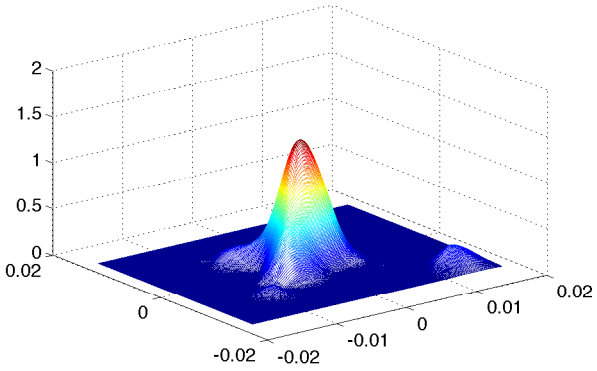

(c) Estimate generated using San Francisco road map

Figure 10. One-time emission of 500 kg of anhydrous ammonia at 10^9 seconds

[17] Participatory urbanism. http://www.urban-atmospheres.net/ParticipatoryUrbanism/index.html, 2007. [Online; accessed 28-October-2007].

[18] S. Schlott, F. Kargl, and M. Weber. Short paper: Random ids for preserving location privacy. In *SECURECOMM '05: Proceedings of the First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM'05)*, pages 415–417, Washington, DC, USA, 2005. IEEE Computer Society.

[19] 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].

[20] Senseweb. http://research.microsoft.com/nec/senseweb/, 2007. [Online; accessed 28-October-2007].

[21] Sensor planet. http://www.sensorplanet.org, 2007. [Online; accessed 28-October-2007].

[22] I. Stojmenovic. Position-based routing in ad hoc networks. *IEEE Communications Magazine*, 40(7):128–134, July 2002.

[23] Urban sensing. http://research.cens.ucla.edu/projects/2006/Systems/Urban_Sensing/default.htm, 2007. [Online; accessed 28-October-2007].

[24] 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.