Environmental Tomography: Ubiquitous Sensing with Mobile Devices

Stacy Patterson #1, Bassam Bamieh *2, Amr El Abbadi #3

#Department of Computer Science, *Department of Mechanical Engineering
University of California, Santa Barbara
Santa Barbara, California, 93106, USA
{1sep, 3amr}@cs.ucsb.edu, 2bamieh@engineering.ucsb.edu

***Abstract*—The ubiquitous nature of mobile phones, which are location-aware devices, presents a unique platform for large-scale computing applications. In particular, if mobile phones are coupled with sensors, they can be used for detection and monitoring of environmental phenomena such as pollution and radiation. In this demonstration, we present *Environmental Tomography*, a system for ubiquitous environmental sensing with mobile devices. Aggregate sensor measurements are collected by the devices along fixed paths such as roads, and these aggregates are used to reconstruct an estimate of the distribution of the underlying physical phenomenon. Our system is robust to the dynamic characteristics of mobile networks and also preserves the privacy of mobile user locations. We demonstrate a prototype that generates estimate distributions from user specified data collection paths and underlying data distributions. The accuracy of the reconstructed distributions is illustrated both numerically and graphically.**

I. Introduction

Mobile phones are ubiquitous, location-aware computing devices. If these devices are coupled with sensors, they can provide a powerful computing platform for large-scale sensing applications. Phones can be equipped with sensors that measure the quantity of harmful pollutants such as sulfur dioxide or contaminants such as radiation, and the sensor readings acquired from these devices can be used to detect and monitor the levels of these harmful substances in populated areas. Since 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 underlying 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. Therefore, any sensing application must only utilize limited resources on each device.
- Mobile device users move in independent, unpredictable patterns. 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.
- Mobile device users have privacy concerns. By reporting location information along with a sensor reading, the user is revealing his location. Users may not be willing to participate in a program that requires them to give up this private location information.

In this demonstration, we present *Environmental Tomography*, an approach for environmental sensing and spatial data modeling that overcomes these challenges. Tomography has long been used in medical imaging techniques such as Computed Tomography (CT), Magnetic Resonance Imaging (MRI), and Positron Emission Tomography (PET) [1]. 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, in Environmental Tomography, one dimensional projections, sums of sensor readings, are collected along fixed query paths across the region, and the projections are used to reconstruct an estimate of the underlying data distribution.

Our data collection approach makes use of two observations about mobile devices and users. First, modern mobile phones are equipped with local communication capabilities such as bluetooth or 802.11 radios, enabling communication between nearby devices. Second, 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 of users, and therefore of mobile devices, along roads and walkways. Aggregates sensor readings can be collected by routing a query message from device to device along a query path via the local communication channels. Devices that receive the query message take a sensor reading and add that reading to the sum of the readings stored in the query message. These sums are reported to a processing center where tomographic reconstruction is used to generate an estimate of the complete underlying data distribution. To improve the accuracy of the reconstruction, we incorporate constraints relating to the physical properties of the sensed phenomena.

Environmental Tomography provides a number of benefits. The data collection approach is designed to overcome the challenges of the mobile network. There are no requirements placed on the individual dynamics of any mobile device, and the resource requirements at each device are low. Since sensor readings are aggregated, there is no need to record the location of any participating device. Therefore, the privacy of user locations is preserved. The technique is also robust

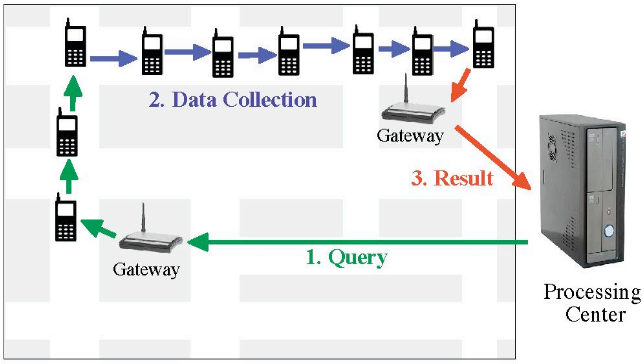

Fig. 1. System Architecture

to query failure and partial information. Accurate estimates can be generated from a small number of query results, and the estimates can be further refined as additional results are reported.

II. System Model and Architecture

A high level view of the system architecture is shown in Fig. 1. The system is made up of a collection of mobile phones that are GPS-enabled and have bluetooth or 802.11 radios for communication with devices in the local neighborhood. The mobile phone users may move about in arbitrary patterns, and devices may join (power on) and leave (power off) the system at any time. We assume that the network of mobile devices is dense along known paths, e.g. roads, so that the phones and their local communication channels along these paths form a connected network.

In addition to the mobile devices, the system also has fixed gateways that are deployed throughout the sensing region. These gateways are capable of communicating wirelessly with mobile devices that are within the local communication range. Each gateway also has a reliable connection to the processing center, which is the control center of the system. The processing center is responsible for issuing queries and processing query results.

Environmental Tomography consists of two phases. The first is the *data collection* phase, in which sums of sensor readings are collected along specified query paths. The second is *tomographic reconstruction*, in which these sums are used to generate an estimate of the underlying data distribution. These phases can occur in parallel; queries can be issued after an initial estimate is generated, and the estimate can be repeatedly refined as new query results are received. We briefly explain each phase below. We refer the reader to [2] for a more detailed development.

A. Data Collection

The entire specification of the query for an aggregate along a path is contained in the *query message*, which is shown in Table I. The mobile devices do not require any advanced

TABLE I
Query Message Specification

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	partial sum of collected sensor readings

knowledge of the queries and can determine their actions from the information in the query message. For simplicity, we assume that the query paths are straight line segments specified by the *start* and *end* coordinates, though it is possible to specify more complex path types. The *start* and *end* coordinates and the *delta* fully determine the set of coordinates at which readings should be taken along the path. We call these points *sampling points*.

The data collection process is illustrated by the arrows in Fig. 1. The processing center creates the query message and sends it to the gateway that is closest to the query start coordinates. The gateway then begins the process of forwarding the query in the mobile network. To route queries, our approach relies upon two established routing protocols for mobile ad-hoc networks: a greedy cartesian routing protocol such as Greedy Perimeter Stateless Routing (GPSR) [3] and Trajectory Based Forwarding (TBF) [4]. Both routing protocols are stateless, meaning the devices do not maintain any route information. Instead, routing decisions are made based on the location of the device and of neighboring devices and the message trajectory or destination. In GPSR, the message is routed from source to destination in a greedy manner where each device forwards the message to the neighboring device that is closest to the destination. In TBF, the goal is not to deliver the message to a destination, but to ensure that the message follows a specified trajectory. In this case, each device forwards the message to the neighbor that lies on or near that trajectory.

Each query is routed from the gateway to the start coordinates using greedy routing. Once the query reaches its starting coordinates, TBF is used to forward the query along the specified trajectory, the query path. Whenever a device with the query message moves on (or near) a specified sampling point, the device takes a sensor reading. The reading is added to the *sum* field in the query. The updated query is then forwarded to the next hop along the query path. When the query reaches the specified end coordinates, it is routed back to the nearest gateway using greedy routing, and the gateway sends the message back to the processing center. If a query cannot be successfully routed through the mobile network, it can simply be aborted, and the query message can be dropped. The processing center will detect that no result has been received and can reissue the query.

B. Tomographic Reconstruction

When the processing center receives the query results, it performs tomographic reconstruction to retrieve an estimate of the data distribution from the aggregate measurements. The reconstruction technique is described below.

We represent the results of the aggregate queries by the vector m where each component m_i corresponds to the result of the i'th query path.

We use the following notation.

- $f : S \to \mathbb{R}$ is a two-dimensional data distribution over sensing region $S \subset \mathbb{R}^2$.
- (x^i_j, y^i_j) represents the j'th sampling point along the i'th query path.
- $f(x^i_j, y^i_j)$ is the point measurement at point (x^i_j, y^i_j), e.g. the level of sulfur dioxide at the point (x^i_j, y^i_j).

Let A be the operator that generates the vector m from an underlying distribution f by computing the sums of measurements along the M query paths

$$A(f) \ := \ \begin{bmatrix} \sum_j f(x^1_j, y^1_j) \\ \vdots \\ \sum_j f(x^M_j, y^M_j) \end{bmatrix}.$$

The goal of tomographic reconstruction is to find a distribution that satisfies the equation $A(f) = m$, which is a linear system of equations in the unknown f. We call this solution the *estimate* distribution, denoted $\hat{f}$.

In medical imaging, it is possible to take projections along many paths in every direction. This provides enough measurements to obtain estimates that can be arbitrarily close to the original distribution. In Environmental Tomography, the choice of paths along which aggregates can be taken is restricted by the layout of the roads and walkways, and this limitation makes the reconstruction problem more difficult. The system is underdetermined. There are, in fact, infinitely many solutions that satisfy $A(f) = m$. In this case, one must specify criteria that define an optimal solution $\hat{f}$ from among this set of feasible solutions. A common approach is to define the optimal solution as the solution where the L^2 norm, $||\hat{f}||_2 :=< \hat{f}, \hat{f} >$, is minimized. However, in the case of a physical phenomenon such as pollution, the solution with the minimum L_2 norm is not the solution that most accurately reflects the physical data distribution. Therefore, we augment the optimization criteria to include physically meaningful constraints, specifically, constraints relating to the properties of diffusive substances.

Any diffusive process, such as diffusion of a gaseous substance, satisfies the diffusion equation [5]

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

If we also assume that the distribution f is quasi-static then the time variation term $\partial f / \partial t$ is small, implying that $\triangle f$ should be small. We quantify the size of $\triangle f$ using the L^2 norm

$$|| \triangle f||_2 :=< \triangle f, \triangle f >=< f, \triangle^2 f > .$$

Incorporating the assumptions about diffusion into our optimization, we use an optimization criterion that is a weighted combination of $||f||_2$ and $|| \triangle f||_2$.

The tomographic reconstruction problem can then be formulated as the following optimization problem: minimize $< f, Qf >$, where $Q := \left(I + \alpha \triangle^2\right)$, subject to the constraint $A(f) = m$, It can be shown that the optimal solution is given by

$$\hat{f} = Q^{-1} A^* (A Q^{-1} A^*)^{-1} m. \qquad (1)$$

The solution generated by the above approach may not satisfy the physical constraint that $\hat{f}$ is non-negative at every point in space. In this case, we can further improve the reconstruction process by introducing a non-negativity constraint. The tomographic reconstruction problem can then formulated as the following optimization problem

$$\begin{aligned} &\text{minimize} \qquad < f, Qf > \\ &\text{subject to} \\ &\qquad A(f) = m \\ &\qquad \forall (x, y) \in S, \quad f(x, y) \geq 0. \end{aligned}$$

Q is a positive semidefinite operator, therefore $< f, Qf >$ is a convex function of f. The linear constraints $\mathcal{A}(f) = m$, and the non-negative cone constraint $\forall (x, y) \in S, \quad f(x, y) \geq 0$ both yield convex constraint sets. The constraint set, which is the intersection of these two sets, is therefore also convex. The overall problem is a convex optimization problem, which implies the existence of a global minimum solution [6]. Such problems can be efficiently solved using readily available convex optimization solvers.

We quantify the accuracy of the estimate distribution using the relative error between the estimate and the original distribution $\overline{f}$.

$$\text{Err}(\hat{f}) := \frac{||\overline{f} - \hat{f}||_2}{||\overline{f}||_2} \qquad (2)$$

To calculate this error, we discretize the distributions $\overline{f}$ and $\hat{f}$ over a 2-dimensional grid and evaluate Equation 2 using the discrete representations.

III. Demonstration

Fig. 2 shows a screen shot of the prototype GUI. The user can configure various settings on the left-hand side, and the actual and estimate distributions are displayed on the right. Users can select from a number of preset underlying data distributions as well as entering a new distribution, and can also choose from a variety of real world city maps as the basis for query path selection. The query path settings are also configurable; the user can specify the sampling delta and select the query paths by clicking on roads on the displayed city map. For example, Fig. 2 shows the estimate distribution that results from 19 query paths along streets and avenues of Midtown Manhattan, NY.

The relationship between the choice of query paths and the accuracy of the estimate is indicated both graphically and numerically. As each path is selected, the estimate is updated to reflect the addition of the query result. Users also have the option of viewing an automated demonstration of the evolution of the estimate as query path results are added to the estimate at random.

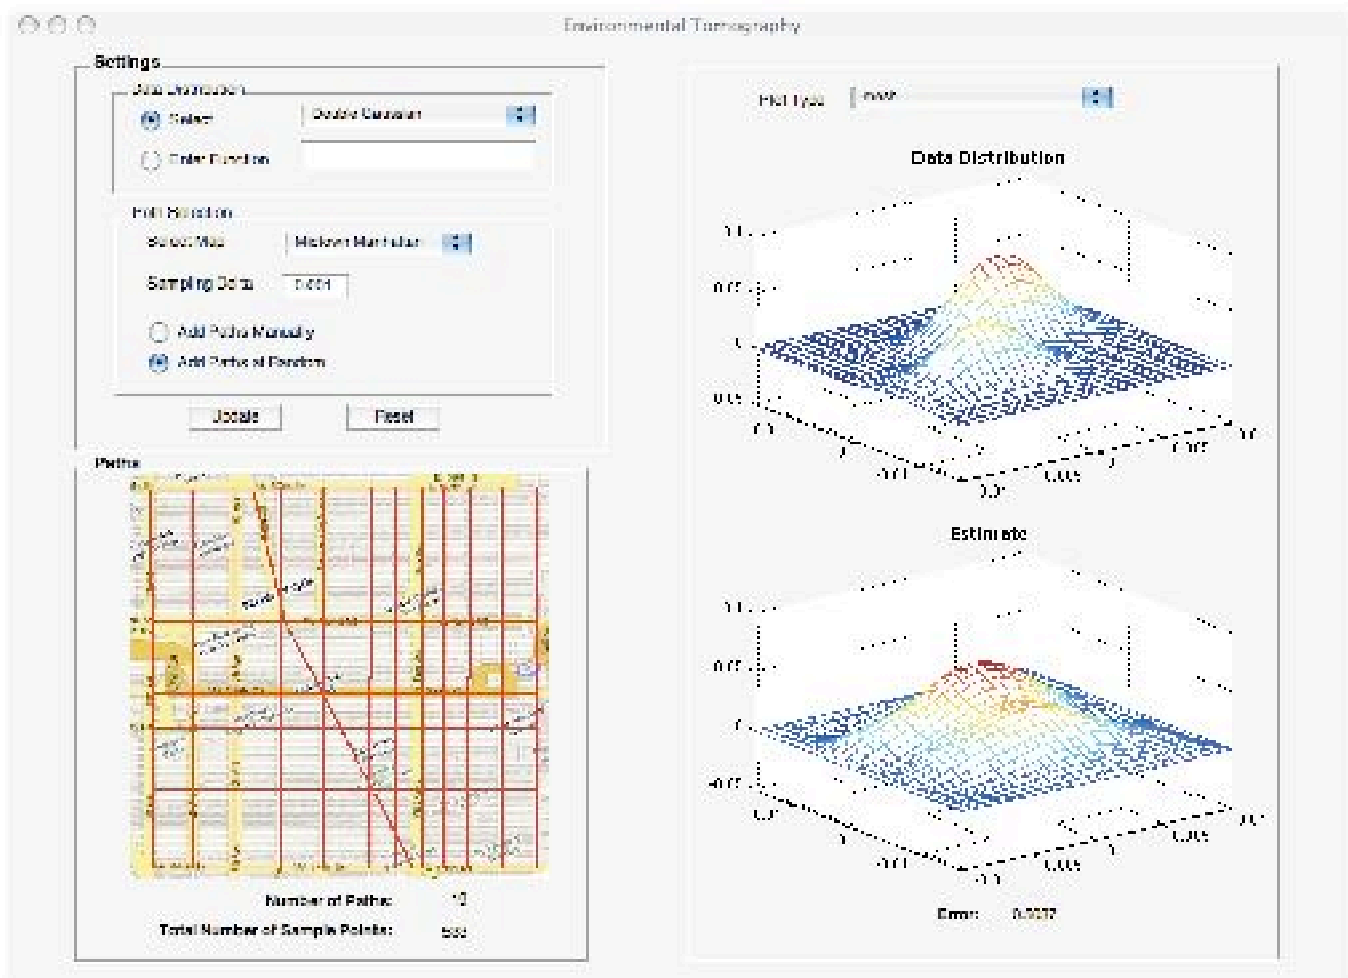

Fig. 2. The Prototype GUI

IV. Discussion

The idea of using cell phones as ubiquitous sensing devices is gaining prominance, most notably in the Department of Homeland Security's (DHS) Call-All program. The Call-All program proposes outfitting cell phones with sensors for detecting biological, radiological, and chemical contaminants. If a particular mobile device detects a contaminant, the device sends an alert of the threat. Concerns about such a program's invasion of user location privacy has led the deputy director of the DHS Advanced Research Projects Agency, Rolf Dietrich, to stress that the program would be voluntary, noting "Not all people would want to play in this game". In this demonstration, we have shown the feasibility of a ubiquitous sensing application that preserves the privacy of user location information. Our solution is also robust to the dynamics and geographical limitations of mobile networks. In future work, we plan to expand our approac to address issues such as optimal query path selection, the role of mobile network density, and the effects of noisy sensor readings and GPS inaccuracies.

Acknowledgment

This work was funded in part by NSF grant IIS 02-23022.

References

[1] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Society of Industrial and Applied Mathematics, 2001.

[2] S. Patterson, B. Bamieh, and A. El Abbadi, "Environmental tomography: Modeling the environment with mobile phones," University of California, Santa Barbara, Tech. Rep., November 2007.

[3] B. Karp and H. T. Kung, "GPSR: greedy perimeter stateless routing for wireless networks," in *MobiCom '00*, 2000, pp. 243–254.

[4] B. Nath and D. Niculescu, "Trajectory based forwarding and its applications," in *MobiCom '03*, 2003, pp. 260–272.

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

[6] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.