

Collaborative Calibration and Sensor Placement for Mobile Sensor Networks

Yun Xiang[†], Lan S. Bai[‡], Ricardo Piedrahita*,
Robert P. Dick[†], Qin Lv[◊], Michael Hannigan*, Li Shang[△]

[†]EECS Department, University of Michigan, Ann Arbor, U.S.A.

[‡]EMC, Pleasanton, CA 94566, U.S.A.

[◊]Dept. of Computer Science, University of Colorado Boulder, CO 80309, U.S.A.

[△]Dept. of ECEE, University of Colorado Boulder, CO 80309, U.S.A.

*Dept. of Mechanical Engineering, University of Colorado Boulder, CO 80309, U.S.A.

[†]xiangyun@umich.edu, [‡]lanceybai@hotmail.com, [†]dickrp@eecs.umich.edu,
{*ricardo.piedrahita, [◊]qin.lv, *hannigan, [△]li.shang}@colorado.edu

ABSTRACT

Mobile sensing systems carried by individuals or machines make it possible to measure position- and time-dependent environmental conditions, such as air quality and radiation. The low-cost, miniature sensors commonly used in these systems are prone to measurement drift, requiring occasional re-calibration to provide accurate data. Requiring end users to periodically do manual calibration work would make many mobile sensing systems impractical. We therefore argue for the use of collaborative, automatic calibration among nearby mobile sensors, and provide solutions to the drift estimation and placement problems posed by such a system.

Collaborative calibration opportunistically uses interactions among sensors to adjust their calibration functions and error estimates. We use measured sensor drift data to determine properties of time-varying drift error. We then develop (1) both optimal and heuristic algorithms that use information from multiple collaborative calibration events for error compensation and (2) algorithms for *stationary sensor placement*, which can further decrease system-wide drift error in a mobile, personal sensing system. We evaluated the proposed techniques using real-world and synthesized human motion traces. The most advanced existing work has 23.2% average sensing error, while our collaborative calibration technique reduces the error to 2.2%. The appropriate placement of accurate stationary sensors can further reduce this error.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design

General Terms

Algorithms, design

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.

IPSN'12, April 16–20, 2012, Beijing, China.

Copyright 2012 ACM 978-1-4503-1227-1/12/04 ...\$10.00.

Keywords

Mobile sensing, Collaborative, Calibration, MILP

1. INTRODUCTION

Mobile sensing applications are increasingly popular. The fast development of smartphones and sensor technology makes many such applications possible, e.g., mobile noise pollution sensing networks [14] and mobile personalized air quality sensor networks [10]. Compact, light, and energy-efficient sensors are now becoming available at prices that permit widespread use by non-scientists (and scientists). In the future, individuals will carry multiple unobtrusive sensors with them, within or networked with their smartphones, forming dense and interconnected sensor networks. Mobile sensing applications will soon become mainstream.

Mobile sensing systems have many advantages over conventional systems composed of a few accurate, low-drift, stationary, and expensive sensing stations. For example, in the personal air quality sensing application, many pollutants have nonuniform spatial distributions [22]. As a result, personal exposure is poorly estimated by using sparsely distributed stationary sensors. If each participant in a sensing system were to carry a sensor, we would be able to better understand human exposure and provide more relevant information to users.

Temporal drift is generally not a concern for expensive stationary sensors, since they are regularly calibrated by experts. However, low cost and compact sensors can accumulate substantial errors in short time spans [8, 17]. Our own measurements of 15 volatile organic compound (VOC) sensors kept in a controlled environment showed up to 20% drift in a single day. Erroneous measurements caused by sensor drift can result in incorrect scientific conclusions, false alarms, and bad decisions. Therefore, low cost sensors require frequent re-calibration.

Manually calibrating sensors to compensate for drift is time-consuming and burdensome; it can annoy users and limit their desire to use the sensors, which will result in an ineffective system. Automatic calibration (which requires no explicit user intervention) has the potential to solve these problems, thereby increasing mobile sensing opportunities.

We propose a system supporting automatic, opportunistic, and collaborative calibration among mobile sensors. Our solution takes into account the gradual increase in sensor drift error with time, and appropriately weights different calibration events based on the

time-dependent estimated errors of the other sensors, i.e., we consider the temporal and spatial properties of the graph formed by (transitive) calibration events. Although we do not require the presence of stationary sensors, we support their inclusion in the system, and also provide algorithms for determining their best locations. Our evaluation makes use of controlled sensor drift studies as well as measured human motion patterns.

The proposed collaborative calibration approach is appropriate for applications with the following characteristics.

1. Spatial variation of sensor readings are low within certain physical distance.
2. Sensor nodes are able to communicate with each other and detect when they are within calibration distance, e.g., either by tracking their own locations or by measuring signal attenuation between nodes.
3. Sensor drift can be compensated for using a drift predictor. The residual error of this predictor has a Gaussian distribution with variance that increases as a function of time, as explained in Section 4.2 and demonstrated in Section 6.1.

Our technique can potentially be used in many mobile sensing applications, such as radiation sensing applications in which sensors are carried by individuals and unmanned aerial vehicles, remote sensing applications in which detailed data are available from in-field sensors and sparse data are available from satellites, and personal environmental sensing. Although the concepts we develop apply to a broader range of mobile sensing systems susceptible to drift error, in the rest of paper, we focus our discussion on a personal air quality sensing application.

It should be noted that collaborative calibration minimizes the increase in the rate of uncompensable drift error, but does not eliminate error. Without the stationary accurate sensors, the mobile sensor network’s overall accuracy degrades over time. The use of a few stationary accurate sensors to augment mobile collaborative calibration is beneficial; it allows the drift error to be bounded.

Our work makes the following main contributions.

1. We formulate and solve the opportunistic collaborative mobile sensor calibration problem.
2. We formulate and solve the mobility aware stationary sensor placement problem to augment collaborative calibration.
3. We propose a sensor drift model built using experimental data from 15 VOC sensors.

To better understand and characterize the effects of real-world human motion on calibration, we also carried out an indoor human motion pattern study on a university campus. Compared with our collaborative calibration scheme, the most advanced existing auto-calibration technique has an average error of 23.2%, while our efficient heuristic has an error of 2.2%. We also present two algorithms for placing stationary sensors to further improve mobile collaborative calibration. The use of well-placed stationary sensors within the collaborative calibration system techniques reduces sensing error significantly, e.g., by about 40% for a density of 1 stationary sensors per 25 mobile sensors. The approximation algorithm based placement technique results in only 6.2% more error than an MILP based technique.

The rest of this paper is organized as follows. Section 2 gives a motivating example. Section 3 summarizes the related work on collaborative calibration and stationary sensor placement. Section 4

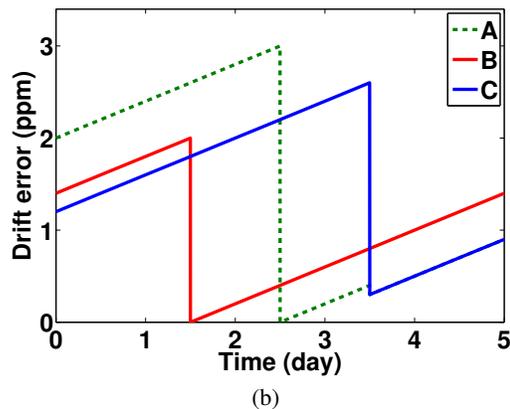
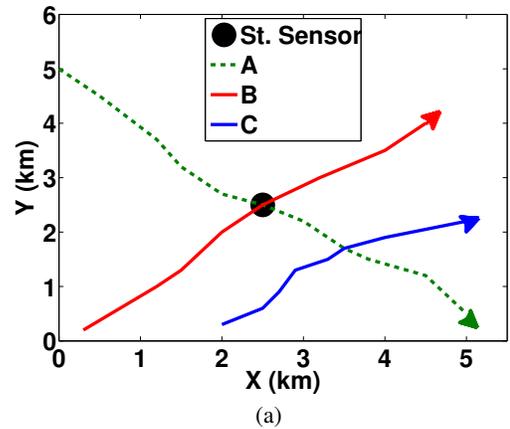


Figure 1: (a) Human motion traces and calibration events and (b) drift errors for three sensors.

describes the sensor random drift model and our collaborative calibration method. Section 5 generalizes the human mobility model, and provides an MILP based solution for the human motion aware stationary sensor placement problem as well as an approximation algorithm. Section 6 describes our controlled-environment experiments for sensor drift and the data analysis results. It also evaluates the performance of our techniques using simulations based on real-world and synthesized human motion traces. Section 7 concludes the paper.

2. MOTIVATING EXAMPLE

Consider a mobile sensor network formed by sensing devices carried by individuals to monitor their air pollution exposures. Each device houses small, energy efficient, and inexpensive metal oxide gas sensors that measure various air pollutants. The sensor measurements gradually drift over time. Drift rates can vary greatly; to minimize error, the sensors must be re-calibrated frequently. In many cases, accurate stationary sensors are not readily accessible for users, and the occasional calibration opportunities they provide are insufficient to cover all the participants in the sensing system. By using collaborative calibration together with optimized placement of stationary sensors, accuracy can be significantly improved.

Figure 1 illustrates an example of our mobile sensor network calibration technique. Figure 1(a) shows the trajectories of three mobile sensors (A, B, and C). Figure 1(b) shows their uncompensable drift errors over time. Each vertical drop in Figure 1(b) corresponds to one calibration event. Between calibration events, the drift error

increases with time as a result of reduced drift prediction accuracy. Given the mobile sensor motion traces, our sensor placement approach decides where to put accurate stationary sensors to maximize the probabilities of mobile sensors being calibrated against the stationary sensor. In this example, the stationary sensor is located at a position both sensor A and B visit, thus providing ground truth calibration for two sensors. When sensor A and B get close to the accurate stationary sensor, their errors drop due to calibration (refer to Figure 1(b)). Our problem formulation and solution also consider a realistic human mobility model that considers individual motion traces able to represent day-to-day variation. With our collaborative calibration technique, even though sensor C never directly calibrates with any (accurate) stationary sensor, its drift error still reduces in the third day by calibrating with sensor A, which has a smaller error due to recent calibration with an accurate stationary sensor.

3. RELATED WORK

This section summarizes prior work on auto-calibration and placement for distributed sensor networks.

Bychkovskiy et al. [3] proposed a two-phase post-deployment calibration technique for dense stationary sensor networks. In the first phase, linear relative calibration relations are derived for pairs of co-located sensors. In the second phase, the consistency of the pair-wise calibration functions among groups of sensor nodes is maximized. Their technique requires a dense deployment of stationary sensors. In contrast, our work focuses on mobile sensor networks.

Miluzzo et al. [15] proposed an auto-calibration algorithm for mobile sensor networks, called CaliBree. In their approach, uncalibrated mobile nodes opportunistically calibrate themselves when interacting with stationary sensors. In their work, calibration events always involve stationary sensors. Our work supports calibration with stationary sensors, but in contrast also supports calibration among mobile sensors, allowing either higher accuracy or a reduction in the number (and therefore cost) of stationary sensors.

Tsujita et al. [21, 22] studied calibration for air pollution monitoring networks. They [22] observed that at a certain time of day, the nitrous oxide pollutant concentration becomes low and uniform in certain areas. They use these opportunities to calibrate mobile sensors using the pollutant concentration reported from nearby environment monitoring stations. In their other work [21], when multiple sensors are close to each other, the average of their readings is used as ground truth to estimate sensor drift. In contrast, we account for the gradual increase in drift error as a function of time, allowing an optimal weighting for each of the many calibration events used to determine drift compensation parameters. Our experimental results show that the technique proposed by Tsujita et al. technique has 23.2% error relative to the optimal result; our proposed heuristic only has 2.2% error.

Berry et al. [2] used an MILP based method to solve the \mathcal{NP} -hard problem of placing sensors in water networks for optimal contamination detection. Chakrabarty et al. [4] tried to find an optimal sensor placement scheme to minimize the cost of sensors while meeting coverage constraints. Our problem formulation differs in that mobile sensors are carried by individuals. A realistic human mobility model is therefore necessary to solve our placement problem. We build our human mobility model based on previous research and our indoor human motion study, and solve the stationary sensor placement problem using a high quality but potentially slow MILP method and an efficient approximation algorithm based technique.

4. COLLABORATIVE CALIBRATION

This section describes our collaborative calibration technique. We present the problem definition, mathematical analysis, and our algorithm to solve this problem optimally.

4.1 Overview

Our collaborative calibration technique uses drift modeling and sensor fusion to reduce drift-related sensor measurement error. Sensor drift models, or drift predictors, are built based on past measured or estimated drift errors. They are used to estimate sensor drift at any point of time and (partially) compensate for drift errors in sensor measurements. In addition, the drift model allows the *residual error* of the drift predictor to be predicted as a function of time. Sensor fusion uses measurements from co-located sensors to improve the accuracy of the combined results. The fusion algorithm determines how to combine multiple sensor measurements based on their residual errors in order to maximize the combined accuracy. In implicit mobile calibration, sensor fusion happens whenever sensors happen to be close to each other; our calibration technique is opportunistic and collaborative.

Actual sensor drifts are only known when sensor measurements are compared with a highly accurate sensor (ground truth). Such opportunities are rare in most personal sensing applications. At most times, drifts must be estimated using models based on prior calibration events with other sensors having varying accuracies. The temporal correlation in sensor drift decreases with temporal distance. Therefore, all other things being equal, the uncertainty of the model (i.e., the residual error of the predictor) increases with the elapsed time since the last drift measurement.

Since nearby sensors are exposed to similar physical conditions, readings from co-located sensors can be combined to statistically improve accuracy. As mentioned before, each sensor has a residual error associated with its post-drift-compensation measurement. Each calibration event allows this error to be reevaluated and potentially reduced. If the two residual errors are independent, the measurement with the smaller residual error should be given more weight during combination. Calibration relationships introduce correlations in sensors' residual errors that the calibration algorithm must account for. Section 4.3 describes our correlation-aware fusion algorithm in detail.

4.2 Collaborative Calibration Problem Definition

Our analytical framework can handle classes of mobile and stationary sensors with arbitrary drift rates. However, we will focus our discussions on systems composed of inexpensive, high drift rate mobile sensors, and expensive but accurate stationary sensors with low drift rates. Although our collaborative calibration technique does not rely on accurate (and most likely stationary) sensors, the use of a few stationary sensors will allow network-wide drift error to be reduced. We assume that these stationary sensors provide accurate readings, either because they are inherently resistant to error or because they are maintained by experts.

Mobile sensors are generally calibrated before deployment but they drift over time. Drift is a function of various factors such as sensing material, exposure to sulfur compounds or acids, aging, or condensate on the sensor surface [1, 8]. It is reported that short-term sensor drift can be modeled accurately with simple models but long-term drift is less predictable [5]. We assume only that (1) there exists an unbiased drift predictor whose residual error has Gaussian distribution and that (2) we have knowledge of how its variance increases over elapsed time since the most recent calibration event.

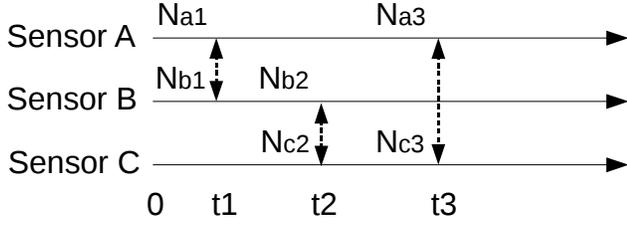


Figure 2: An example of sensor error correlation as a result of previous calibration events.

As explained in Section 6.1, we observed that high-quality predictors for our sensors have this property.

Our goal is to develop a distributed technique that automatically compensates for sensor drift error; there is no notion of a central controller that has access to data from all sensors. Avoiding dependence on a central controller can reduce sensing system energy consumption, cost, and security problems.

We now present the formal problem definition. Given N mobile sensors and M accurate stationary sensors, the location of a mobile sensor i at time t is $L_i(t)$, $i \in N$. The location of accurate stationary sensor j is L_j , $j \in M$. Sensor i 's raw reading (including drift error) at time t is $r_i(t)$. Its drift prediction function is $f_i(t, k_1, k_2, \dots, k_n)$. The parameters of this function may be different for each sensor and may change over time. The error associated with the drift predictor $e(t)$ changes over time. The drift-compensated sensor reading is $R_i(t) = r_i(t) - f_i(t)$. The accurate value of the monitored parameter at location l and time t is G_l^t . Let $C_i(t)$ be the post-calibration sensor reading. In other words, $C_i(t)$ is the sensor reading after drift compensation and sensor fusion. The goal is to determine k_1, k_2, \dots, k_n for each sensor to minimize its total mean squared error, i.e., $\sum_t (G_l^t - C_i(t))^2$. Each sensor i at time t , only has access to $R_j(t)$ of sensor j when $|L_i(t) - L_j(t)| < D_c$ (D_c is the calibration range).

Our measurements in several rooms suggest that in well-ventilated rooms with no obvious pollution sources, the pollutant mixture is spatially homogeneous within 2 m distance. We will use this distance as calibration range D_c in simulations. Note that the spatial distributions of air pollutant concentrations vary based on nearby pollution sources and ventilation conditions, thus the calibration range depends on circumstances.

4.3 Error Estimation and Error Propagation

As we mentioned before, each sensor has a residual error that is adjusted after each calibration event. In this section, we describe how this residual drift error is calculated and minimized via calibration and prediction. We address the problem of predictor design for one particular type of sensor in this paper. In general, the predictor should be provided by the sensor manufacturer or determined by pre-deployment lab calibration.

We start with a simple scenario where errors of two sensors are independent. Assume two co-located sensors A and B. Sensor A's current error estimate is n_a and sensor B's current error estimate is n_b , where n_a and n_b are random numbers with Gaussian distributions N_a and N_b and standard deviations E_a and E_b (in the rest of the paper, we use N to represent a Gaussian distribution, n to represent a random number following distribution N , and E to represent its standard deviation). Assume this is the first time sensors A and B calibrate with other sensors. N_a and N_b are independent and their standard deviations, E_a and E_b , are determined by how long the sensors remain uncalibrated. Let G be the ground truth

value of the physical condition measured by the sensors. Readings from these two sensors can be represented as $R_a = G + n_a$ and $R_b = G + n_b$. The weighted sum of R_a and R_b is $R_{ab} = \alpha \cdot R_a + (1 - \alpha) \cdot R_b = G + N(0, \sqrt{\alpha^2 \cdot E_a^2 + (1 - \alpha)^2 \cdot E_b^2})$. It is easy to prove that when

$$\alpha = E_b^2 / (E_a^2 + E_b^2), \quad (1)$$

the weighted sum has minimal standard deviation for both calibrated sensors, i.e., $G + N(0, E_a E_b / \sqrt{E_a^2 + E_b^2})$. A reading from the sensor with smaller error is given more weight. After calibration, both sensors should adjust their readings to R_{ab} and use R_{ab} to estimate their current ground truth readings as well as to predict future drifts.

Now we consider the scenario in which N_a and N_b are correlated. This may happen as a result of both sensors directly or transitively calibrating with the same mobile sensor prior to their calibration with each other. In this case, we need to know the correlation between N_a and N_b to compute the optimal combination of their readings. Let us consider the example shown in Figure 2. Assume three sensors A, B, and C all start operating at time 0. At time t_1 , sensors A and B calibrate. Their calibration parameters are independent of each other at that time and thus the analysis in the previous paragraph for independent errors can be applied. Assume weights of 0.2 and 0.8 are used, thus the error after calibration is $0.2n_{a1} + 0.8n_{b1}$. At time t_2 , sensors B and C calibrate. Assume sensor B's drift prediction error increased by n_{b12} from time t_1 to t_2 . The errors of B and C are still independent. Assume the optimal weight is 0.5 in this case. After calibration, B's and C's errors are $0.1n_{a1} + 0.4n_{b1} + 0.5n_{b12} + 0.5n_{c2}$. At time t_3 , sensors A and C calibrate. A's error is now $n_{a3} = 0.2n_{a1} + 0.8n_{b1} + n_{a13}$ and C's error is $n_{c3} = 0.1n_{a1} + 0.4n_{b1} + 0.5n_{b12} + 0.5n_{c2} + n_{c23}$. Note that at that moment, these two sensors contain the same errors generated from the previous calibration, which are n_{a1} and n_{b1} . Now N_a and N_c are correlated and Equation 1 cannot be directly applied. However, it is still possible to use the weight assignment technique to find an optimal solution. To do that, we can remember all the independent distributions and weight assignments from previous calibration events.

Now we present the general approach that accounts for correlation introduced by transient calibration events among sensors. Each sensor's error distribution is represented as a weighted sum of multiple independent error distributions. Each independent distribution is from the other sensor's or its own increased prediction error over the uncalibrated time interval. Label the two calibrating sensors as sensor 1 and 2. Let S_1 and S_2 be the sets of independent error distributions for sensors 1 and 2. Let C be the intersection of S_1 and S_2 , i.e., $C = S_1 \cap S_2$. Let C_1 and C_2 be S_1 and S_2 's non-overlapping regions, i.e., $C_1 = S_1 - C$, $C_2 = S_2 - C$. Let W_{1i} and W_{2i} be the weights associated with the error distributions for sensors 1 and 2, δ_i be the standard deviation of each distribution, and G be the ground truth value of measured object. Sensor 1's reading after drift compensation is

$$R_1 = G + \sum_{i \in C_1} W_{1i} N(0, \delta_i) + \sum_{j \in C_2} W_{1j} N(0, \delta_j). \quad (2)$$

Sensor 2's reading is

$$R_2 = G + \sum_{i \in C_1} W_{2i} N(0, \delta_i) + \sum_{k \in C_2} W_{2k} N(0, \delta_k). \quad (3)$$

In order to generate more accurate results by combining the readings of sensor 1 and 2, we use a linear weighted sum function to combine their drift-compensated measurements. Assuming the

weights are α and $1 - \alpha$ for sensor 1 and 2 respectively, the combined result is

$$\begin{aligned} R_{12} &= \alpha R_1 + (1 - \alpha) R_2 \\ &= G + \sum_{i \in C} [\alpha W_{1i} + (1 - \alpha) W_{2i}] N(0, \delta_i) \\ &+ \sum_{j \in C_1} \alpha W_{1j} N(0, \delta_j) + \sum_{k \in C_2} (1 - \alpha) W_{2k} N(0, \delta_k). \end{aligned} \quad (4)$$

The variance of the error for the combined reading is

$$\begin{aligned} Var &= \sum_{i \in C} [\alpha W_{1i} + (1 - \alpha) W_{2i}]^2 \delta_i^2 + \sum_{j \in C_1} W_{1j}^2 \alpha^2 \delta_j^2 \\ &+ \sum_{k \in C_2} W_{2k}^2 (1 - \alpha)^2 \delta_k^2. \end{aligned} \quad (5)$$

The derivative of the variance is

$$\begin{aligned} \frac{dVar}{d\alpha} &= 2\alpha \sum_{i \in C} (W_{1i} - W_{2i})^2 \delta_i^2 + 2 \sum_{i \in C} W_{2i} (W_{1i} - W_{2i}) \delta_i^2 \\ &+ 2\alpha \sum_{j \in C_1} W_{1j}^2 \delta_j^2 + 2\alpha \sum_{k \in C_2} W_{2k}^2 \delta_k^2 - 2 \sum_{k \in C_2} W_{2k}^2 \delta_k^2. \end{aligned} \quad (6)$$

To minimize the variance, we have $\frac{dVar}{d\alpha} = 0$, therefore

$$\alpha = \frac{\sum_{i \in C} W_{2i} (W_{2i} - W_{1i}) \delta_i^2 + \sum_{k \in C_2} W_{2k}^2 \delta_k^2}{\sum_{i \in C} (W_{1i} - W_{2i})^2 \delta_i^2 + \sum_{j \in C_1} W_{1j}^2 \delta_j^2 + \sum_{k \in C_2} W_{2k}^2 \delta_k^2}. \quad (7)$$

Equation 7 gives the general expression for weight assignment. In the case of two independent sensors (C is empty), we have

$$\alpha = \frac{\sum_{k \in C_2} W_{2k}^2 \delta_k^2}{\sum_{j \in C_1} W_{1j}^2 \delta_j^2 + \sum_{k \in C_2} W_{2k}^2 \delta_k^2} = \frac{E_2^2}{E_1^2 + E_2^2}, \quad (8)$$

which is consistent with Equation 1.

Note that the above analysis applies only to the scenario where collaborative calibration involves two sensors. It is possible to extend the evaluation to an arbitrary number of co-located sensors, although this would increase the complexity of the weight assignment expression.

4.4 Collaborative Calibration Algorithm

We have presented the key concept allowing the optimal calibration algorithm to combine readings from co-located sensors. Now we present the complete algorithm for collaborative calibration, which includes drift compensation, weight assignment, and drift reevaluation. Note that calibration opportunity detection is not part of our algorithm. There are multiple existing approaches to discover calibration opportunities, including radio communication (e.g., Bluetooth), ultrasound, and passive audio environment based proximity detection schemes [10, 16, 20].

The key data structure used is a table that stores all the independent error distributions and their corresponding weight assignments for each sensor. Each entry is a tuple of name, weight, and standard deviation. The names are used to distinguish independent error distributions. The calibration algorithm for a mobile sensor labeled i that calibrates with sensor j is shown in Algorithm 1.

Mobile sensors participating in the collaborative calibration system carry out three actions every time a calibration event happens: (1) estimate its current drift with its drift predictor and use the result to compensate its raw reading, (2) estimate the ground truth

value and update its error table, and (3) use the estimated ground truth value to recompute its drift, residual error, and drift predictor. The type of co-located sensor determines the details of step (2). If the co-located sensor is an accurate stationary sensor, its reading can be directly used as ground truth to estimate the mobile sensor's drift. The mobile sensor ignores its own reading and directly overwrites its own reading with the reading from the stationary sensor and its current error immediately drops to zero. As a consequence, it can forget all previous calibration errors as they become irrelevant (clear the table). Otherwise, if the co-located sensor is also a mobile sensor with a non-zero error, its drift-compensated reading is combined with the mobile sensor's drift-compensated reading according to Equation 7 to generate an estimate of ground truth and the error distribution table will be updated accordingly.

Algorithm 1 Collaborative calibration algorithm for mobile sensor i

Require: r_i // i 's raw reading
Require: R_j // j 's calibrated reading
Require: T_i // i 's error table
Require: T_j // j 's error table
Require: t // current time

if j is accurate stationary sensor **then**
 $R_i \leftarrow R_j$
 $D_i'(t) \leftarrow r_i - R_i$
Update drift model
 $T_i.clear()$

else
Predict current drift D_i
 $R_i \leftarrow r_i - D_i$
 $T_i.insert(i, t, g(t - last_cali_t), 1)$
 $C \leftarrow T_i \cap T_j$
 $C_1 \leftarrow T_i - C$
 $C_2 \leftarrow T_j - C$
Compute α using Equation 7
 $R_{ij} \leftarrow \alpha R_i + (1 - \alpha) R_j$
Update current drift $D_i'(t) \leftarrow r_i - R_{ij}$
Update drift model
for $k \in C$ **do**
 $T_i[k].weight \leftarrow T_i[k].weight \times \alpha + T_j[k].weight \times (1 - \alpha)$
end for
for $k \in C_1$ **do**
 $T_i[k].weight \leftarrow T_i[k].weight \times \alpha$
end for
for $k \in C_2$ **do**
 $T_i[k] \leftarrow (T_j[k].name, T_j[k].var, T_j[k].weight \times (1 - \alpha))$
end for
end if
 $last_cali_t \leftarrow t$

5. STATIONARY SENSOR PLACEMENT

In this section, we consider placement of stationary sensors to further assist the collaborative calibration of mobile sensors. Our discussion will focus on human-carried sensors.

5.1 Overview

Adding stationary sensors to a system composed of collaboratively calibrating mobile sensors can further improve accuracy. The number of stationary sensors is constrained by cost; they must be carefully positioned to enable frequent calibration opportunities with mobile sensors. Fortunately, humans move with patterns that can

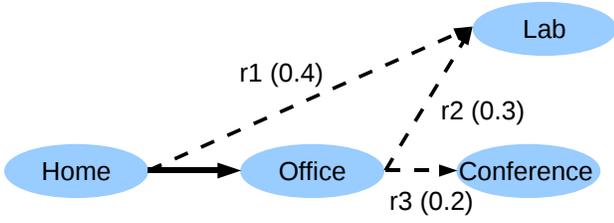


Figure 3: Example human motion trace with 3 patterns.

be used to our benefit; some locations are more frequently visited than others [13].

Recent research has shown that most people’s daily motion patterns are predictable [7, 18, 19]. We present a stochastic human mobility model capable of capturing the most relevant motion patterns for the stationary sensor placement problem. The field for stationary sensor deployment is modeled as a grid in which implicit calibration may occur among sensors in the same grid element. It is possible to eliminate discretization problems by making grid elements arbitrarily small and permitting calibration between nodes in multiple grid elements within the calibration distance. We define a motion pattern as a set of locations (grid elements) that a person is likely to visit on a particular day. An individual’s mobility model is a probability-weighted collection of possible motion patterns. Extreme sensor drift typically occurs on a timescale of days, not hours, enabling a simplified model that neglects the order of visited locations within a single day. In our evaluation, these models are extracted from measured motion traces as well as those generated by software provided by human motion pattern researchers [13].

Daily motion patterns are weighted with probabilities. For example, as shown in Figure 3, there are three distinct patterns: r_1, r_2 , and r_3 . A value ranging from 0 to 1 is associated with each pattern to indicate its probability. It is possible for multiple stationary sensors to be encountered by a person in a day. However, encountering one is sufficient for calibration.

5.2 Sensor Placement Problem Definition and MILP-Based Solution

We now define the problem of stationary sensor placement to assist calibration of mobile sensors.

Problem Definition: The field for stationary sensor deployment can be represented by a grid G . A set of people S move within the grid. Each person $s \in S$ carries a mobile sensor. A person’s motion pattern for a particular day, r_s , is a set of locations. R is the set of all motion patterns, and the motion patterns associated with a particular person s are represented with R_s . Each motion pattern r is associated with a value p_{sr} , which is the probability of person s having pattern r . The sum of the calibration probabilities of all patterns of person s is P_s . A total number of k sensors are deployed in the field. The optimization objective is to find a set of grid elements in which stationary sensors should be placed to maximize the average daily probability of mobile sensor calibration, i.e., $\frac{\sum_{s \in S} P_s}{k}$.

This problem is \mathcal{NP} -hard. Let each pattern be represented by an element associated with a probability weight and each possible stationary sensor placement location be represented by a subset. An element belongs to a subset if and only if the corresponding pattern contains the placement location. Given a resource constraint, k , the original problem can be stated as selecting at most k subsets such that the covered elements have maximum total weight. This is the

weighted maximum coverage problem [11]. We will now describe an MILP formulation for the problem.

Maximize

$$\frac{\sum P_s}{k}, \forall s \in S,$$

subject to

$$\sum_{(i,j) \in G} x_{ij} \leq k, \quad (9)$$

$$\forall r \in R, \sum_{(i,j) \in r} x_{ij} - M d_r \leq 0, \quad (10)$$

$$\forall r \in R, \sum_{(i,j) \in r} x_{ij} - m d_r \geq 0, \quad (11)$$

$$P_s - \sum_{r \in R_s} d_r * p_{sr} = 0, \quad (12)$$

$$1 \geq x_{ij}, \text{ and } d_r \geq 0. \quad (13)$$

x_{ij}, d_r are integers. M and m are constants and are set to $k + 1$ and 0.5. The probabilities p_{sr} are known. The properties of binary indicators x_{ij} and d_r are described below.

$$x_{ij} = \begin{cases} 1 & \text{if a sensor is placed at grid element } (i, j) \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

and

$$d_r = \begin{cases} 1 & \text{if pattern } r \text{ is covered by at least one sensor} \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

M is greater than the largest possible value of $\sum_{(i,j) \in r} x_{ij}$ (which is satisfied by setting M to be $k + 1$) and m is less than the smallest possible non-zero value of $\sum_{(i,j) \in r} x_{ij}$ (which is satisfied by setting m to be 0.5).

5.3 Approximation Algorithm Based Placement Technique

Algorithm 2 Approximation based placement technique

Require: G // deployment field grid

Require: R // set of all patterns

Require: P // probabilities

Require: k // stationary sensor count constraint

$C \leftarrow \{\}$ // output set

while $\text{size}(C) \leq k$ **do**

Select $g \in G$ s.t. $\sum_{r \in g} P_r$ is maximized

Remove the covered patterns from R

$C \leftarrow C \cup g$

end while

Normally MILP-based solutions are not tractable for large instances of hard problems. Fortunately, the number of patterns per person is limited: it is possible to directly use the MILP formulation for substantial problem instances. The solver performance is further improved because human motion traces tend to be spatially clustered [13]. We will show in Section 6.3 that our algorithm can be applied to deployment cases with up to 840 km² area or 200 patterns. It is conceivable that some problem instances will exceed



Figure 4: Calibration chamber used for sensor drift experiments.

the size tractable for MILP solvers. Therefore, we also present an approximation algorithm based polynomial time heuristic.

The maximum coverage problem can be solved with the polynomial time $(1 - \frac{1}{e})$ -approximation algorithm shown in Algorithm 2. This is minimum achievable bound [11]. However, the $(1 - \frac{1}{e})$ -approximation bound only applies for the average calibration probability between stationary and mobile sensors. There are many other factors influencing the network sensing accuracies, such as collaborative calibration events, calibration time, and calibration order. Section 6.3 evaluates the approximation algorithm based technique in detail.

6. EXPERIMENTAL RESULTS

In this section, we first describe our controlled drift experiments (Section 6.1), which support the hypothesis in Section 4.2. Section 6.2 presents simulation results for our optimal and efficient collaborative calibration techniques and compares them with two existing works that are most related. Section 6.3 reports on the performance of our MILP based stationary sensor placement algorithm and compares it with the efficient approximation algorithm we propose.

6.1 Calibration Procedure and Drift Experiments

Section 4.2 describes our sensor drift model. We assume that drift can be (partially) compensated for by an unbiased predictor, and the residual error can be modeled using a Gaussian distribution with a variance that predictably increases with time. To test this hypothesis, we have conducted a drift experiment in our controlled chamber.

Before the drift experiment, we manually calibrated all sensors. Calibrations were performed using de-humidified zero grade air (i.e., air with less than 1 ppm total hydrocarbons) and controlled-concentration iso-butylene (a VOC unlikely to damage graduate students when used at low concentration). The purpose of this calibration is to compensate for initial measurement offsets, possibly due to variation in the manufacturing process. During calibration and drift experiments, sensors are mounted on a custom printed circuit board enclosed in the 250 cm³ polycarbonate chamber as shown in Figure 4. A fan is mounted inside the chamber to improve mixing and make convection heat loss from the sensors uniform. The temperature and humidity inside the chamber are stabilized at

43.8 ± 1.3 °C, and $7.8 \pm 1.7\%$ respectively. A LabVIEW interface controls the gas mixture using mass flow controllers. During calibration runs, the sensors are held at concentrations of 0, 0.25, and 1.0 ppm (parts per million by volume) of iso-butylene in a total volume flow of 4 liters per minute, for 20 minutes each. The sensors are powered continuously throughout the experiment period, and were warmed up for two weeks prior to starting the experiments to allow the sensors to reach an initial equilibrium, as recommended by the manufacturer.

During the drift experiment, 15 pre-calibrated Figaro TGS 2602 VOC sensors are placed in the controlled gas chamber and exposed to 4 liters per minute air. These exposure tests last 120 minutes and are performed daily. Since the sensors are powered continuously, they should drift constantly during the experiment. The drift data are calculated by averaging the last 30 minutes of readings from each test to avoid any warm-up effects from changes in the air flow rate.

We use the analog to digital converter on Labjack U3 data acquisition modules to measure the voltage output of the TGS sensors, at a sampling frequency of 0.5 Hz. We use log-based transfer function to convert the voltages to VOC concentrations, based on calibrations performed before the experiment. The concentration readings after conversion are shown in Figure 5. Since the ground truth reading should be 0 ppm, the readings after the conversion already represent drift. Seven of the 48 measurements were discarded due to inconsistent air flow rate or relative humidity levels due to transient problems with the testing chamber air supply.

We now evaluate a simple drift predictor based on linear extrapolation of two consecutive drift errors to predict future errors. The difference between the predicted drift value and the measured drift is the portion of the drift error that is not captured by the drift model. We have also evaluated higher-order non-linear predictors but they did not have higher prediction accuracies than the linear predictor. The linear predictor compensated for 94.1% and 87.7% of the drift on average when predicting one day and two days ahead. We therefore consider it to be a good predictor for this kind of sensor. Note that for different sensor types, the forms of the predictor function may be different. In some cases, a higher order non-linear fitting function might be necessary.

We applied the Lillie normality test to the residual error of the linear predictor. The residual error has a Gaussian distribution, with an exception for predictions eight days in advance. For most cases, the linear predictor meets Gaussian residual requirement posed in Section 4.2. For specific sensors and time offsets passing the normality test, we perform t-tests to assess whether the distributions have means of 0 ppm. The significance levels used in the Lillie test and t-test are both 0.05 and the test results are shown in Figure 6(a). Figure 6(b) shows the standard deviation of the remaining drift error after applying the linear predictor for up to 10 days in the future. The results clearly show an increasing trend for all the sensors, consistent with our hypothesis in Section 4.2 that the variance increases over time. The standard deviations of the short-term drift errors can be well predicted using simple linear functions.

With one possible anomaly at an eight-day offset, the drift experiment results confirm our hypothesis that the residual error after drift prediction has a Gaussian distribution with mean 0 and predictable variance that increases over time.

6.2 Evaluation of Collaborative Calibration

To evaluate our collaborative calibration algorithm, we compare it with two other approaches proposed in relevant and recent work. In the first approach, Calibree [15], all mobile sensors calibrate with stationary accurate sensors. In contrast, our calibration tech-

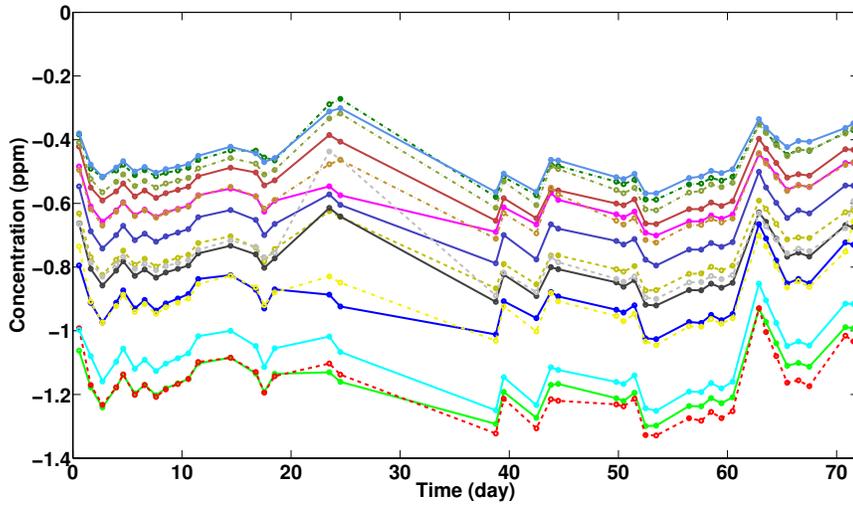


Figure 5: Measured drift error as a function of time for Figaro TGS2602 VOC sensors.

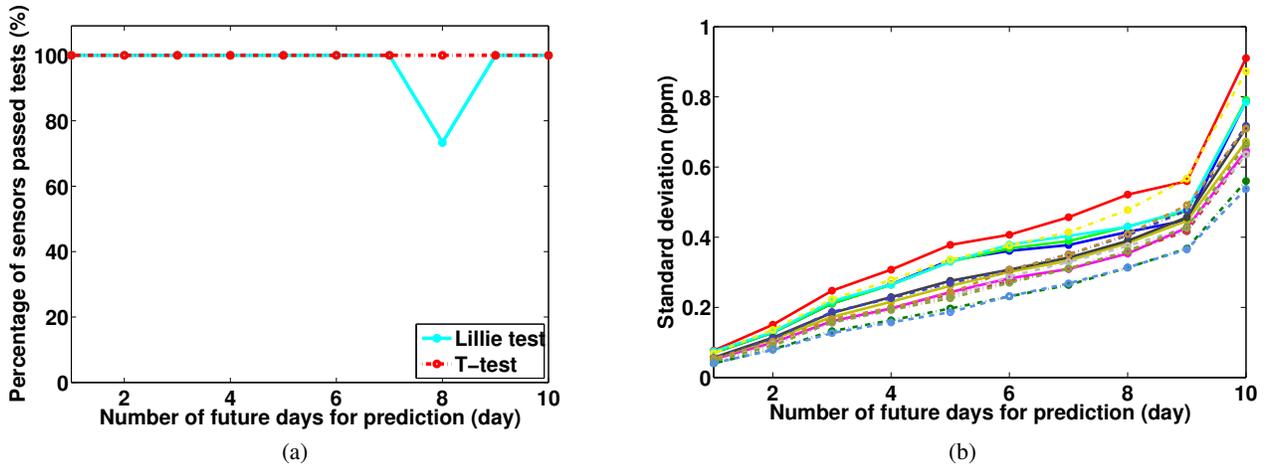


Figure 6: (a) The normality test results and (b) the standard deviations of prediction errors using the 2-day linear predictor to compensate for 1 to 10 days of future drift.

Table 1: Aggregated Sensor Error with Synthesized Human Motion Traces

Trace	Num. of cali. events			Total aggregated mean squared error			
	Total	Uncorrelated	Stationary	CaliBree	Averaging	Heuristic	Optimal
1	44,290	5,072	21,818	964.6	393.6	321.9	312.1
2	43,378	3,368	20,144	1,716.6	559.0	454.9	434.8
3	9,701	1,722	4,429	3,059.0	1,461.1	1,244.3	1,229.8
4	5,659	1,048	2,589	6,805.8	2,359.6	1,984.0	1,966.3
5	14,308	2,496	4,398	8,610.6	3,234.7	2,681.8	2,643.6
Average overhead (%)				224.8	23.2	2.2	0

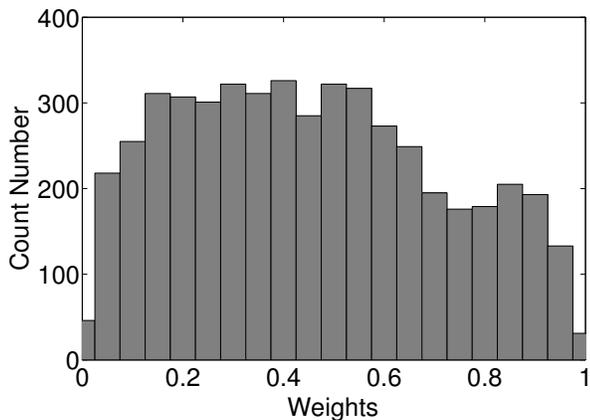


Figure 7: Histogram of assigned weights for an example trace using the optimal collaborative calibration scheme.

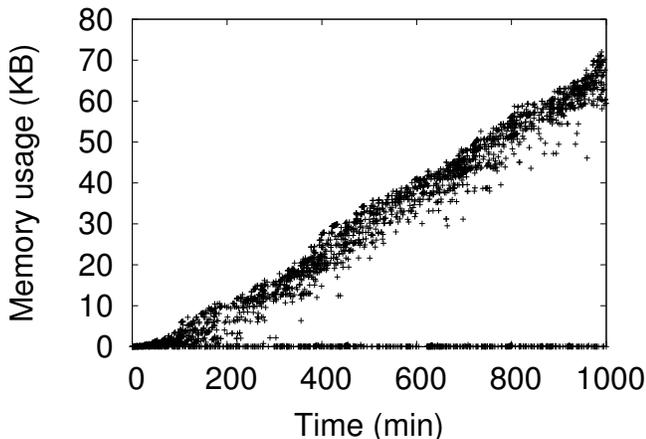


Figure 8: Memory use of the optimal collaborative calibration scheme.

nique allows sensors to calibrate with each other as well as stationary sensors. In the second approach [21], readings from co-located sensors are averaged to estimate the ground truth value. In contrast, our technique enables more accurate drift compensation by considering the differing drift prediction errors of calibration events, i.e., sensors. We also propose and evaluate a *calibration heuristic* that reduces computation complexity and memory use at the cost of a very slight reduction in calibration accuracy. This heuristic ignores correlations between prediction errors. Instead tracking independent error distributions from previous calibration events and temporal error growth, this algorithm only stores an aggregated error for each sensor. During calibration, it uses Equation 1 to assign weights to readings from co-located sensors. We evaluate the four approaches with the same set of motion traces and sensor placements, and compare the resulting accumulated mean squared error. For this experiment, we use 10 stationary accurate sensors placed at the most frequently visited locations and use a random walk model for sensor drift.

Table 1 shows the results for the four approaches with five synthesized motion traces generated using the SLAW human mobility model [13]. The second to the fourth columns present statistics for calibration events for the optimal algorithm. The second column shows the total number of calibration events. A pair-wise calibration between two sensors is considered to be two calibration events. The third column shows the number of calibration events in which

Table 2: Statistics for Human Mobility Case Study

Participant	Duration (days)	On campus prob. (%)	# of patterns	# of locations
1	30	90.0	12	11
2	30	86.7	5	5
3	22	77.3	4	4
4	23	100.0	5	4
5	21	76.2	7	6
Average	25.3	85.2	6.6	6

the errors from two sensors are independent. The fourth column shows the number of calibrations with stationary accurate sensors. The last four columns show the aggregated mean squared errors of all sensors during the entire experiment.

On average, CaliBree [15] has 224.8% more error than optimal. This is because it only considers calibration events between stationary and mobile sensors, and thus misses opportunities for calibration between mobile sensors. 43.6% of calibration events occur between mobile and stationary sensors; the rest occur between pairs of mobile sensors.

Tsujita’s technique (averaging) has 23.2% more error than optimal result. Figure 7 shows the distribution of the weights generated with the optimal algorithm for Trace 5. The weights are widely distributed from 0 to 1. Only 25.4% are in the range from 0.4 to 0.6. The structure of this histogram has implications for the effectiveness of Tsujita’s approach: the closer weights are to 0.5, the more effective Tsujita’s approach.

Our heuristic produces results with accuracy that deviates from optimal by only 2.2%. Even though the percentage of correlated events is fairly large (41.8%), ignoring the correlation does not significantly degrade accuracy. However, this algorithm greatly reduces required memory compared with the optimal algorithm. With the optimal algorithm, the memory use increases linearly with time for most sensors. Figure 8 shows the memory use over time for all sensor nodes in our experiment with trace 1. Each point corresponds to a sensor node involved in a calibration event. We therefore conclude that the heuristic is more efficient and likely to be appropriate for most practical applications.

The optimal algorithm allows us to evaluate the quality of various calibration approaches. In summary, utilizing the interactions among mobile sensors improves the accuracy by 224.8% compared to only permitting mobile sensors to calibrate with stationary sensors. The accuracy is improved by 23.2% by considering the heterogeneity of drift estimation parameters among different sensors. Considering correlations among sensors due to calibration imposes large computation complexity and memory use with a relatively small gain (2.2%). In summary, a technique using collaborative calibration among mobile sensors that considers heterogeneity in drift estimation parameters but ignores calibration event induced inter-sensor correlations represents a good trade off between accuracy and run-time overhead/complexity.

6.3 Evaluation of Stationary Sensor Placement

This section introduces our human motion pattern case study and evaluates our stationary sensor placement algorithms with both measured and synthesized human mobility traces.

6.3.1 Measured Human Mobility Case Study

Much human mobility modeling research is based on outdoor GPS data [7, 13, 19]. However, GPS is inaccurate indoors, where humans spend 90% of their time [6]. According to a survey-based model, office worker indoor activities can be modeled using a few

Table 3: Statistics for Measured and Synthesized Human Motion Traces and Solver Performance

Trace	Area (km ²)	Total pat.	Sensor no.	Cand. loc.	Runtime (s)
Case study	N/A	33	5	17	0.01
KAIST	840.1	92	92	41,270	1.2
NCSU	142.3	35	35	10,691	0.13
New York	618.8	39	39	12,180	0.05
Orlando	122.0	41	41	26,662	0.07
State fair	1.2	19	19	4,422	0.03
1	0.01	200	50	1,225	0.13
2	0.01	200	50	1,001	0.24
3	1.0	200	50	26,448	2.44
4	1.0	200	50	39,695	5816.10
5	4.0	400	100	101,891	> 6 h

patterns [12]. In our evaluation, we use mobility traces generated using algorithms proposed by other researchers as well as data gathered in our real-world human mobility study, which was conducted on the campus of University of Colorado Boulder.

In our study, five graduate students, undergraduate students, and professors used their mobile phones to record their daily motion patterns. Participants manually entered locations and times into their smart phones as they moved and these data were sent to a server via the Internet. Locations in which users spent fewer than five minutes were omitted from the motion patterns. The study was conducted between August 3rd, 2011 and September 12th, 2011. Statistics from the study are shown in Table 2. Motion patterns contain 1.94 locations on average, which implies that the indoor activities of the participants were spatially concentrated, which is consistent with the findings of other human motion studies [12, 19].

6.3.2 Experiment on Measured and Synthesized Human Motion Traces

To solve the MILP problem, we use the CPLEX v.12.2 solver [9] on an Intel 4-core Xeon E3 1230 CPU running at 3.2 GHz with 8 GB of memory. The evaluation is performed on both real-world and mobility model generated [13] human motion traces.

The statistics of the real-world and synthesized human motion traces [13], as well as our case study trace, and their MILP solver performances are shown in Table 3. The case study trace does not contain detailed location information, but lasts for multiple days. The rest of the real-world traces contain detailed location information, but are finished within a day each, i.e., each person has one motion pattern. The duration for each trace is 4 days, i.e., each person has 4 patterns. According to our real-world case study, the average probabilities of the top 4 patterns are 0.48, 0.2, 0.1, and 0.08. The same probability values are used in the synthesized traces. The fourth column of the table shows the total number of mobile sensors in each trace. The fifth column shows the total number of candidate locations where stationary sensors may be placed. Grid elements visited by one or more person are considered as placement location candidates. The total number of the candidate locations is equal to the number of variables x_{ij} in Equation 9.

The MILP placement algorithm quickly solves all the problem instances, except for synthesized trace 5. For this trace, the solver terminated after six hours without producing a solution. This trace contains 400 patterns and 101,891 candidate placement locations. We conclude that the MILP solution is suitable for many useful-scale problem instances, but there may be some real-world cases for which a more efficient solution is required, e.g., that in Section 5.3.

The results of the MILP placement algorithm are shown in Figure 9. For most of the solutions, the number of sensors is far less than the number of patterns. This is consistent with the hypothesis that people’s motion traces tend to be clustered, repetitive, and frequently overlap each other. The synthesized human motion traces typically required fewer sensors despite having more motion patterns because a relatively small geographical area was considered in these traces. In summary, although personal mobile sensors are needed to monitor the conditions experienced by many individuals, the accuracy of these sensors can be improved substantially by using a few accurate stationary sensors to assist a collaborative calibration technique.

The results of evaluating the algorithms on both real-world and synthesized human motion traces are shown in Table 4. We assume that repeated calibration with a stationary sensor during the same day does not further reduce error. The aggregated network error (the sum of mean square errors of all the sensors in the network for readings taken every 30 seconds) is measured when both placement algorithms are permitted to use the number of stationary sensor listed in the second column of Table 4. For the synthesized traces, we assume that all the patterns occur with the same probability. The fifth column of Table 4 shows the aggregated network error using our optimal collaborative calibration technique, assuming there are no stationary sensors. The results show that the approximation algorithm based technique increases aggregated network error by 6.2% compared to the MILP placement algorithm. Note that for Trace 4, the approximation algorithm based technique outperforms the MILP solution. In that case, the approximation algorithm had already reached 99% average calibration probability, making its solution essentially equivalent to the MILP solution. Note that in our placement problem formulation, the error caused by calibration order is neglected. However, since the uncompensable drift error within a day is small (less than 0.1 ppm as shown in 6(b)), this simplification has very little impact on solution quality.

7. CONCLUSIONS

We have presented a collaborative calibration and sensor placement framework for mobile sensor networks. We developed a random sensor drift model based on controlled experiments and developed a collaborative calibration technique to compensate for drift error. We also describes placement techniques for stationary sensors used to augment collaborative calibration among mobile sensors. We conducted a human motion study on a university campus to build these models and evaluate our placement algorithms. Experimental results indicate that, compared with our collaborative calibration algorithm, the most advanced existing work has an average sensor error of 23.2%. Our stationary sensor placement algorithms further reduce the effects of drift error.

8. ACKNOWLEDGEMENTS

We would like to thank Yifei Jiang and Xin Pan for help with the human mobility case study. This work was supported in part by the National Science Foundation under awards CNS-0910995, CNS-0910816, and CNS-1059372.

9. REFERENCES

- [1] K. Arshak, E. Moore, G. M. Lyons, J. Harris, and S. Clifford. A review of gas sensors employed in electronic nose applications. *Sensor Review*, 24(2):181–198, 2004.
- [2] J. Berry, L. Fleischer, W. Hart, C. Phillips, and J. Watson. Sensor placement in municipal water networks. *J. of water resources planning and management*, 131(3):237–243, 2005.

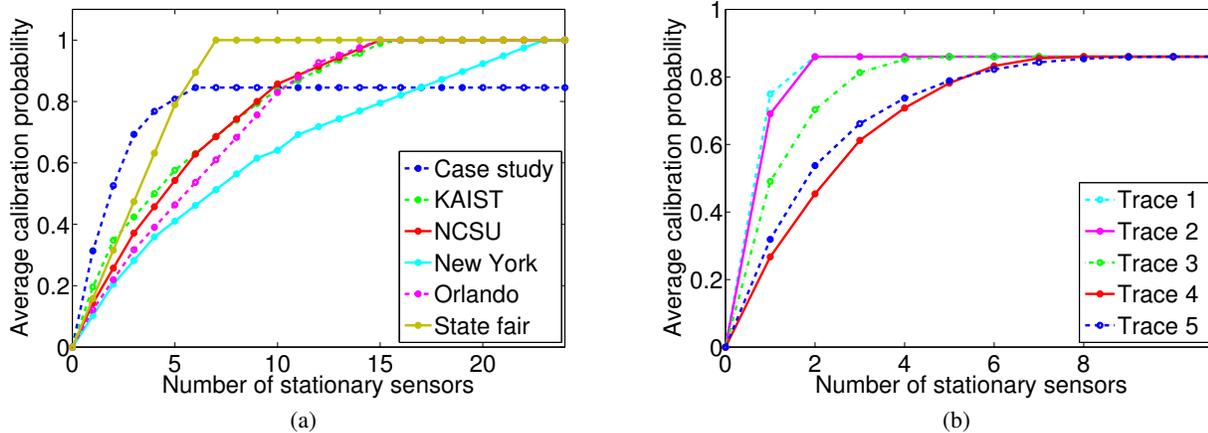


Figure 9: The MILP stationary sensor placement results for (a) measured human motion traces and (b) synthesized human motion traces.

Table 4: Aggregated Sensor Errors for Different Placement Algorithms

Trace	Sensor number			Aggregated error			
	MILP	Approx. Algo.	Improvement	All Mobile	MILP	Approx. Algo.	Improvement
KAIST	16	19	18.8%	9,880	7,875	8,465	7.5%
NCSU	15	15	0.0%	6,075	3,095	3,333	7.7%
New York	23	26	13.0%	4,720	2,076	2,504	20.6%
Orlando	15	16	6.7%	7,208	3,683	3,954	7.4%
State fair	7	7	0.0%	5,303	2,649	2,786	5.2%
1	2	2	0.0%	910	523	551	5.4%
2	2	3	50.0%	1,083	701	738	5.3%
3	5	5	0.0%	2,326	1,783	1,831	2.7%
4	8	9	12.5%	3,370	2,522	2,511	-0.4%
5*	10	11	10.0%	3,924	3,195	3,205	0.3%

*The MILP solution is derived by setting the relative tolerance of the MILP solver to be 0.3%.

- [3] V. Bychkovskiy, S. Megerian, D. Estrin, and M. Potkonjak. A collaborative approach to in-place sensor calibration. In *Proc. Int. Conf. Information Processing in Sensor Networks*, pages 301–316, Apr. 2003.
- [4] K. Chakraborty, S. Iyengar, H. Qi, and E. Cho. Grid coverage for surveillance and target location in distributed sensor networks. *IEEE Trans. Computers*, 51(22):1448–1453, Dec. 2002.
- [5] A. Emami-Naeini, M. Akhter, and S. Rock. Effect of model uncertainty on failure detection: the threshold selector. *IEEE Trans. Automatic Control*, 33(12):1106–1115, Dec. 1988.
- [6] EPA. Buildings and their impact on the environment: a statistical summary, 2009.
- [7] M. Gonzalez, C. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.
- [8] J.-E. Haugen, O. Tomic, and K. Kvaal. A calibration method for handling the temporal drift of solid state gas-sensors. *Analytica Chimica Acta*, 407(1–2):23–39, Feb. 2000.
- [9] IBM ILOG CPLEX Division. IBM ILOG CPLEX 12.0 user manual, 2008.
- [10] Y. Jiang, K. Li, L. Tian, R. Piedrahita, Y. Xiang, O. Mansata, Q. Lv, R. P. Dick, M. Hannigan, and L. Shang. MAQS: A personalized mobile sensing system for indoor air quality monitoring. In *Proc. Int. Conf. Ubiquitous Computing*, pages 271–280, Sept. 2011.
- [11] S. Khuller, A. Moss, and J. Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.
- [12] J. Kim, V. Sridhara, and S. Bohacek. Realistic mobility simulation of urban mesh networks. *Ad Hoc Networks*, 7(2):411–430, 2009.
- [13] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong. SLAW: a mobility model for human walks. In *Proc. Int. Conf. Computer Communications*, pages 855–863, Apr. 2009.
- [14] N. Maisonneuve, M. Stevens, M. Niessen, P. Hanappe, and L. Steels. Citizen noise pollution monitoring. In *Proc. Int. Conf. Digital Government Research*, pages 96–103, 2009.
- [15] E. Miluzzo, N. D. Lane, A. T. Campbell, and R. Olfati-saber. CaliBree: a self-calibration system for mobile sensor networks. In *Proc. Int. Conf. Distributed Computing in Sensor Networks*, pages 11–14, June 2008.
- [16] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proc. MOBICOM*, pages 32–43, 2000.
- [17] A. Romain and J. Nicolas. Long term stability of metal oxide-based as sensors for e-nose environmental applications: an overview. *Sensors and actuators, B, Chemical*, 146(2):502–506, Apr. 2010.
- [18] U. Schlink, K. Strebel, M. Loos, R. Tuchscherer, M. Richter, T. Lange, J. Wernicke, and A. Ragas. Evaluation of human mobility models, for exposure to air pollutants. *Science of The Total Environment*, 408(18):3918–3930, Aug. 2010.
- [19] C. Song, Z. Qu, N. Blumm, and A.-L. Barabasi. Limits of predictability in human motion. *Science*, 327:1018–2021, Feb. 2010.
- [20] S. P. Tarzia, P. A. Dinda, R. P. Dick, and G. Memik. Indoor localization without Infrastructure using the acoustic background spectrum. In *Proc. Int. Conf. on Mobile Systems, Applications, and Services*, pages 155–168, June 2011.
- [21] W. Tsujita, H. Ishida, and T. Moriizumi. Dynamic gas sensor network for air pollution monitoring and its auto-calibration. In *IEEE Proc. Sensors*, pages 56–59, Oct. 2004.
- [22] W. Tsujita, A. Yoshino, H. Ishida, and T. Moriizumi. Gas sensor network for air-pollution monitoring. *Sensors and Actuators, B, Chemical*, 110(2):304–311, Oct. 2005.