

CaliBree: A Self-calibration System for Mobile Sensor Networks

Emiliano Miluzzo¹, Nicholas D. Lane¹, Andrew T. Campbell¹, and Reza Olfati-Saber²

¹ Computer Science Department,
Dartmouth College, Hanover NH 03755, USA
{miluzzo,niclane,campbell}@cs.dartmouth.edu

² Thayer School of Engineering,
Dartmouth College, Hanover NH 03755, USA
olfati@dartmouth.edu

Abstract. We propose *CaliBree*, a self-calibration system for mobile wireless sensor networks. Sensors calibration is a fundamental problem in a sensor network. If sensor devices are not properly calibrated, their sensor readings are likely of little use to the application. Distributed calibration is challenging in a mobile sensor network, where sensor devices are carried by people or vehicles, mainly for three reasons: *i*) the sensing contact time, i.e., the amount of time nodes are within the sensing range of each other, can be very limited, requiring a quick and efficient calibration technique; *ii*) for scalability and ease of use, the calibration algorithm should not require manual intervention; *iii*) the computational burden must be low since some sensor platforms have limited capabilities. In this paper we propose CaliBree, a distributed, scalable, and lightweight calibration protocol that applies a discrete average consensus algorithm technique to calibrate sensor nodes. CaliBree is shown to be effective through experimental evaluation using embedded wireless sensor devices, achieving high calibration accuracy.

1 Introduction

Sensors calibration is a fundamental problem in sensor networks. Without proper calibration, sensor devices produce data that may not be useful or can even be misleading. There are many possible sources of error introduced into sensed data, including those caused by the sensing device hardware itself. Hardware error can be broken down into that caused by the sensor hardware component, the sensing device component, and the sensor drift (sensors' characteristics change by age or damage). The sensor hardware level error is corrected at the factory where a set of known stimuli is applied to the sensor to produce a map of the output. The sensing device error component is introduced when the sensor is mounted on the board itself that includes the microcontroller, the transceiver, and the circuitry that form a sensor node [1] (we call the sensor plus the supporting board a sensor device). To correct the sensing device error component calibration of the sensor device is required.

While device calibration must sometimes be done in the factory (e.g., for high precision medical sensors), a growing number of sensors are embedded in consumer devices [4] [5] and are currently used particularly in a number of popular recreational

domain [6], and emerging [7] [8] applications. This latter class of cheap sensors are generally shipped without any sensor device calibration and it is up to the user to perform the calibration procedure to make sure that the gathered sensed data is meaningful. Moreover, sensors drift from their initial calibration over time. This imparts a significant burden to the user of the sensor devices. Further, this manual method of calibration process does not scale when considering large scale people-centric deployments and applications.

We conjecture that in mobile sensor networks [3] [16] [17] there will be two classes of sensors: calibrated nodes that can be either static or mobile, and uncalibrated nodes. We refer to the nodes belonging to the former class as *ground truth* nodes. These ground truth nodes may exist as a result of factory calibration, or user manual calibration.

We propose CaliBree, a distributed, scalable, and lightweight protocol to automatically calibrate mobile sensor nodes in this environment. In CaliBree, uncalibrated nodes opportunistically interact with calibrated nodes to solve a discrete average consensus problem [9], leveraging cooperative control over their sensor readings. The average consensus algorithm measures the disagreement of sensor samples between the uncalibrated node and a series of calibrated neighbors. The algorithm eventually converges to a consensus among the nodes and leads to the discovery of the actual disagreement between the uncalibrated node's sensor and calibrated nodes' sensors. The disagreement is used by the uncalibrated node to generate (using a best fit line algorithm) the calibration curve of the newly calibrated sensor. The calibration curve is then used to locally adjust the newly calibrated node's sensor readings.

CaliBree relies on opportunistic rendezvous between uncalibrated nodes and ground truth devices because we want the calibration process to be transparent to the user. The convergence time of the algorithm depends on the density of the ground truth nodes. Still, if the density was low, the accuracy of the algorithm would not be impacted, only the convergence time would be extended. However, we expect urban sensor networks [3] [18] will have a high density of ground truth nodes. In the CitySense project, well calibrated sensor nodes are mounted on light poles in an urban area. Those sensors can be considered as ground truth nodes that could be used by mobile nodes running the CaliBree algorithm.

In order for the consensus algorithm to succeed, the uncalibrated sensor devices must compare their data when sensing the same environment as the ground truth nodes. Given the limited amount of time mobile nodes may experience the same sensing environment during a particular rendezvous, and the fact that even close proximity does not guarantee that the uncalibrated sensor and the ground truth sensor experience the same field of view, the consensus algorithm is run over time when uncalibrated nodes encounter different ground truth nodes. We experimentally determine that an uncalibrated node achieves calibration after running CaliBree with less than five different ground truth nodes.

The contribution of this paper is:

- It proposes, to the best of our knowledge, the first fully distributed approach to calibrating mobile sensor devices such as embedded sensor devices [6] [7] and sensor enabled cellphones [4] [5].
- It proves the existence of the *sensing factor* (see Section 2) which we believe is an important characteristic to be considered in the design of protocols and applications for mobile sensing systems.

- It presents a calibration technique which is efficient, scalable, and lightweight, therefore suitable to be applied to mobile sensing systems.
- It shows the experimental evaluation of the CaliBree protocol through validation using a testbed of static and mobile embedded sensor devices [1].

In the following sections, we describe the motivation, design, and evaluation of the CaliBree system. In Section 2 we motivate the need of an efficient and scalable calibration protocol for mobile sensor networks. In Section 3 we discuss the shortcomings of existing techniques proposed in the literature to achieve sensor networks calibration. The CaliBree design is illustrated in Section 4, and Section 5 describes the experimental approach we took to validate CaliBree. We summarize our work in Section 6 where we also discuss our future research direction.

2 Motivation

Curiously the issue of calibration of wireless sensor networks has received low attention in the literature despite it being recognized as a fundamental issue. Without calibration the data acquired from such networks is meaningless. This obvious fact and the difficulties in performing calibration in general is a repeated finding of real world sensor network deployments [19] [20]. Emerging mobile sensing architecture which are the focus of this paper will not be different.

To quantify the magnitude of the calibration problem we perform experiments with our own building sensor network testbed comprising both mobile and static sensors. We perform experiments to: *i*) show the variability of the individual calibration curves between multiple sensor nodes considering two different sensing modalities, and *ii*) quantify how the differences in these calibration curves would impact the actual reported sensor values from these nodes. For both of these experiments we used the Tmote Sky wireless sensor [1], a multimodal sensing platform commonly used by the experimental sensor network community.

Manual Sensor Calibration. For all the experiments performed unless otherwise noted we used a set of manually calibrated Tmote Sky sensor nodes as ground truth nodes. Given the linear response of the Tmote Sky sensors we took four calibration points for 21 different Tmote Sky on two of the available sensor suite on the node, namely the PAR (Photosynthetically Active Radiation) light sensor, which has a frequency response range approximately equivalent to that of a human eye, and the temperature Sensirion AG temperature/humidity sensor. Ground truth sensor readings were provided for temperature by the Extech SuperHeat Psychrometer RH350 [30] and for light by the Extech Dataloggin Light Meter 401036 [31]. As per the typical manual calibration process a calibration curve specific to the sensor in question was determined by taking a linear regression of measurements of the physical phenomena (provided by the ground truth sensors) and the raw output of the sensor in question. The value of the raw output used in the regression was the mean of 20 individual raw readings.


Figure 1(a) and Figure 1(b) present the larger and smaller bounds sensor specific calibration curves for, respectively, the PAR and temperature sensors. The larger bound curve is associated to a node to which we give the label “A” whereas the node associated to the lower bound curve is labeled with “B”. The x-axis represents the raw output

of the sensors while the y-axis provides actual light (in Lux) and temperature (in degrees Celsius) values that are expected to correspond to the Analog to Digital Converter (ADC) output of the sampled sensors.


We also derive the sensed data values from the factory Tmote Sky calibration curves and compare them to the values calculated with the manually derived calibration curve for same ADC outputs. The difference between the readings is plotted in Figure 1(c) and Figure 1(d) for the PAR and temperature sensors respectively and represents the calibration error of the sensor nodes coming off the factory. In Figures 1(c) and 1(d) the y-axis presents the calibration error of nodes A and B, relative to the ADC output of the sensor itself, which is reported in the x-axis. In Figure 1(d) errors of up to 55 degrees Celsius are shown depending on the temperature range. Similarly, in Figure 1(c) error as large as nearly 2,600 Lux are demonstrated.

Variations in the sensor specific calibration curves. Our experimentation demonstrated differences both in the gain and offset of the calibration curves for each sensor node and manufacturer provided generic calibration curve (see Figures 1(a) and 1(b)). Not only were there differences between the individual sensor nodes but there were patterns in the type of calibration error depending on the modality considered. The light sensor had very small difference in the offset of the calibration curves while it had substantial difference in the gain between each curve. For instance the error when the light sensor is exposed to bright environments, where the ADC's output is near 4000 units, is nearly 2600 Lux, whereas in darker environments (low ADC output) the error can become small. In contrast was temperature sensor for which the calibration curves had both gain and offset variation but with offset differences being more substantial. This suggests any general approach to calibration must be able to both determine the unique gain and offset values for each sensor device.


Existing calibration approaches within wireless sensor networks do not assume the presence of well calibrated sensor nodes in the network (i.e., [21] [23] [22]). This is in line with the typical assumption about the use of cheap, lower quality (i.e., radio interface and sensors), resource constrained sensor nodes that collectively form a dense network over the sensor field. However, these assumptions do not hold in an emerging class of sensor network architectures [16] [3] [14] [5] deployed in urban areas, where such restrictions are not longer motivated, and comprising new advanced forms of embedded sensor platforms such as sensor enabled cellphones [4]. In particular, this impacts how calibration should be performed within such networks. Specifically, the potential exists for a subset of the devices to be capable of acting as ground truth nodes, i.e., sources of reliable calibrated sensor data. Such nodes are able to support the rest of the network comprised by uncalibrated sensors. Due to the availability of ground truth sensor data a traditional approach to calibration becomes possible. This being the approach of determining the calibration curve for a particular sensor based upon a collection of sensor values that can be compared to the ground truth sensor data. This comparison is opportunistic in the sense that due to the uncontrolled mobility patterns of either or both the ground truth and uncalibrated nodes, the uncalibrated nodes probabilistically encounter ground truth nodes. We envision a scenario where urban-scale sensor network deployments provide a number of well calibrated sensors [3] [18] with which uncalibrated nodes could rendezvous in order to perform calibration.




(a) Plot of the upper and lower bounds of the PAR calibration curves obtained by manually calibrating 21 Tmote Sky nodes.




(b) Plot of the upper and lower bounds of the temperature calibration curves obtained by manually calibrating 21 Tmote Sky nodes.




(c) Plot of the light calibration error measured by comparing the sensed data obtained from the factory calibration curves and the manually generated calibration curves.



(d) Plot of the temperature calibration error measured by comparing the sensed data obtained from the factory calibration curves and the manually generated calibration curves.



(e) Contour plot of the mean light intensity measured with a set of 21 Tmote Sky nodes arranged in a 3 by 7 grid with a 0.5 meters spacing within an indoor office space environment.



(f) Contour plot of the mean temperature gradient measured with a set of 21 Tmote Sky nodes arranged in a 3 by 7 grid with a 0.5 meters spacing within an indoor office space environment.

Fig. 1.

Superficial consideration of this approach to calibration suggests that performing calibration in such networks is trivial. However, several factors make the calibration challenging. Firstly, the calibration can only occur when the ground truth node and the uncalibrated node are experiencing identical sensing environment. This is necessary

because the comparison between calibrated and uncalibrated data is only meaningful when the same input to the sensors is applied. Secondly, the calibration rendezvous is complicated by the existence of the *sensing factor*. The sensing factor is identified by the tendency of a physical phenomenon to be localized to a small region around the entity taking the measurement. If for example we consider the light sensing modality, given the high directional nature of light, the light readings reported by a light sensor are relative to the proximity region of the sensor. In contrast, for the temperature modality, the temperature gradient around a temperature sensor presents a much smaller variation. In general, the existence of the sensing factor is largely independent of the specific modality in question and can be related to a broad class of sensors (e.g, light, dust/pollen, CO_2 , sound, etc.). The variability of the sensor data relative to an originating location of sampling increases rapidly as the distance from this origin increases (for example consider the exponential decay of various phenomena such as light and heat). We note that the sensing factor, since it is based upon invariant physical laws, will remain the same regardless of the components of the sensor devices themselves. Unlike discussions of short communication rendezvous durations found in DTNs (Delay Tolerant Networks) (such as discussed in [13]) which could be an artifact of the short range low power radios, the sensing factor will be present regardless the radio technology and more or less dependent on the sensing modality.

Characterizing the Sensing Factor. To quantify the sensing factor we performed an experiment where both light and temperature were sampled from a 3 by 7 static grid of 21 calibrated Tmote Skys separated by a 0.5 meters distance. The nodes were placed in an indoor environment during daylight hours. Figure 1(e) is a contour plot of light readings and Figure 1(f) is the contour plot for the temperature readings. Both Figures 1(e) and 1(f) clearly demonstrate the variability of light and temperature over relatively short distances. Gradients exist with the sampled phenomena and the variability of these gradients increases with distance. It is evident that the variation of the light intensity is larger than for the temperature (light drop is of about 500 Lux in just 0.5 meters whereas one Celsius degree variation in temperature is obtained over more than 2 meters). This implies that if for example a light ground truth node was positioned in the $(x,y)=(1,0)$ location in Figure 1(e), an uncalibrated light sensor node needs to move close within 0.5 meters distance from the ground truth node in order to sample a sensing environment similar to the one of the ground truth node and perform accurate calibration. For the temperature sensor, the distance between ground truth node and uncalibrated nodes within which the calibration can be performed becomes larger.

In general, mobility combined with the sensing factor reduces the time interval in which nodes experience the same environment which is a requirement to perform accurate calibration. CaliBree is designed to operate quickly when uncalibrated nodes enter the same sensing environments of ground truth nodes. It also allows uncalibrated nodes to exploit distance information between themselves and ground truth nodes and make decisions about whether to rendezvous with them or not.

3 Related Work

A significant amount of work spanning many decades addresses the general problem of sensor calibration. However, relatively few solutions are developed for the more

recently formed conception of wireless sensor networks [10] [12] [11]. The bulk of research in this more focused area deals primarily with energy-efficient networking and distributed computing, rather than with accurate sensing. In fact, the payload of packets in these networks is often treated as a black box that is ignored or abstracted away. The work that does exist in calibrating wireless sensor networks assumes a dense network of static and highly resource-constrained nodes [28], and is not directly applicable to sensor networks with uncontrolled mobility (e.g., [3]), the environment assumed in this paper. These networks comprise loose federations of heterogeneous nodes with variable mobility patterns (i.e., static and mobile) that lead to variable and often sparse nodes density.

Motivated by the unscalability of manual calibration techniques with a known standard input signal, the authors of [24] propose a technique called “macro-calibration” for use in networks of thousands of nodes. The technique builds an optimization problem from trends and relationships observed in the aggregate sensor data provided by the network to generate calibration equations. However, the design and evaluation of [24] focuses on the accuracy of range estimates between nodes to support localization. Others have also contributed solutions limited to the needs of localization [21] [22] [23] [29] and time synchronization rather than the sensor modality agnostic type of calibration that is our focus. More general and less modality and application specific calibration is considered in [25]. This work also adopts the aforementioned “macro-calibration” approach in densely deployed networks. More recent work [26] does not require the same levels of density, but assumes that sampled sensor data is band limited. By sampling this data above the Nyquist rate, the actual sensor values will exist in a lower dimensional subspace of the higher dimensional space comprised by the uncalibrated readings.

A calibration approach involving robotic network elements is presented in [27], whereby a robot with calibrated sensors gathers samples within the sensor field to allow already deployed static sensors to be calibrated/re-calibrated as required. Our work differs in that we do not depend on controlled robotic mobility but rather we exploit opportunistic rendezvous between mobile uncalibrated and calibrated sensor nodes carried by humans and their vehicles. Further, our solution does not require the introduction of costly and complex robotic hardware.

4 CaliBree Design

In this section we present the design of the CaliBree protocol. Recall the definition of sensing contact time as the time window in which mobile nodes experience approximately the same sensing environment. Similarly, we term the spatial region where nodes experience the same sensing environment as the *common sensing range*. The sensing contact time depends on several factors including mobile node’s speed, sensor orientation, obstacles to a sensor’s field of view, and the physical sensing range limit. The common sensing range varies with sensor type. For example, as shown in Section 2, the common sensing range for light tends to be small due to its highly directional nature, whereas for temperature the common sensing range is larger since the temperature gradient is typically small in the proximity of a human-carried sensor device. Given that an uncalibrated node should experience the same sensing context as the ground truth node during the calibration process, if the common sensing range is small the sensing

contact time could be very short. Moreover, if either or both the uncalibrated or ground truth nodes are moving the sensing contact time may be further shortened. We show in Section 5 that in the case of a location-dependent sensing modality like the light sensor, the sensing contact time is in the order of few seconds under human mobility patterns. There is then a need to design a calibration protocol that is fast, completing during the short sensing contact time.

To this end, CaliBree is designed to solve a distributed average consensus problem. Equation (1) shows the formulation of the discrete consensus problem we use:

$$\bar{d}_i(k+1) = \begin{cases} (1-\epsilon) \cdot \bar{d}_i(k) + \epsilon \cdot \sum_{j=1}^{N_i} \frac{s_i(k)-s_j(k)}{N_i}, & k > 0, \\ d_i^{uncal}, & k = 0. \end{cases} \quad (1)$$

$\bar{d}_i(k)$ is the average disagreement measured by node i up to round k ; $0 < \epsilon < 1$ sets the weight given to the current round's disagreement consensus; N_i is the set of ground truth nodes in i 's neighborhood at round k ; $\frac{s_i(k)-s_j(k)}{N_i}$ is the average disagreement between i 's sample and those of the ground truth nodes in range at round k ; d_i^{uncal} is the difference between the uncalibrated data and one of the ground truth node's data when the calibration starts.

The consensus algorithm formulated in (1) works as follows. Each ground truth node periodically transmits a beacon advertising its availability to participate in the calibration routine for at least one sensor type. If an uncalibrated node wishes to calibrate its sensor of an advertised type, it replies to this advertisement, triggering the CaliBree protocol (a distance-based energy optimization to this trigger is described in Section 4.3). Upon receiving a reply to its advertisement, the ground truth node starts broadcasting a series of packets containing its instantaneous sensed data value for the sensor type under calibration. We call these broadcasts packets sent by the ground truth nodes *calibration beacons*. As the uncalibrated node starts receiving the calibration beacons it begins running the consensus algorithm. The uncalibrated node calculates the difference between its own sensed data and the sensed value from the ground truth node and feeds this difference into Equation 1 as $s_i - s_j$. Equation 1 outputs the current estimate of the average 'disagreement' between the uncalibrated and ground truth nodes. The average disagreement \bar{d} from Equation 1 decreases as the uncalibrated node physically approaches the ground truth node(s) and their sensing ranges begin to overlap. Considering a particular node pair (i, j) , the minimum \bar{d} occurs at the time of maximum sensing range overlap between the uncalibrated node and the ground truth node (i.e., they both are sampling a similar environment). This \bar{d}^{min} estimates the uncalibrated sensor device's true offset from the ground truth sensor.

In the following, we demonstrate the ability of the consensus algorithm to converge to the minimum disagreement. We investigate the calibration of the light sensor for this experiment and throughout the paper given the challenge implied by the highly directional nature of light. By showing that CaliBree is able to work for light, which potentially leads to small sensing contact times due to its sensitivity to sensor orientation and obstructions, we gain confidence that it works well for other sensor types, under less demanding constraints as well. We implement the consensus algorithm on Tmote Sky [1] wireless sensor nodes, and start by manually calibrating a single ground truth node. The ground truth node is placed on a shelf next to the window in our lab.

Figure 2(a) shows the evolution of the average disagreement \bar{d} between the uncalibrated light sensor and the single ground truth node as a human carries the uncalibrated sensor periodically towards and then away from the ground truth sensor. The minima in Figure 2(a) represent the minimum disagreements and are achieved every time the uncalibrated node arrives within the common sensing range of the ground truth node. In this experiment the minimum disagreement is approximately 700 Lux. In Section 5 we show that the common sensing range, i.e., the spatial region where nodes experience the same sensing environment, depends on the relative context of the nodes (e.g., light sensor orientation).


At the moment the minimum disagreement is achieved, the uncalibrated sensed data plus the average minimum disagreement \bar{d}^{min} gives the actual ground truth sensor readings. For uncalibrated node i and ground truth sensor j , we call the value of i 's sensor at the moment of minimum disagreement s_i^{min} and the value of j 's sensor at the same moment s_j^{min} . Then we have that $s_i^{min,j} + \bar{d}^{min,j} = s_j^{min}$, where $s_i^{min,j}$ is the value of i 's sensor at the moment of minimum disagreement during the rendezvous with ground truth sensor j . Moreover, we use $ADC_i^{min,j}$ to refer to the output of the analog to digital converter (ADC) that samples the sensor on the sensing device i at the time of the minimum disagreement during the rendezvous with the ground truth node j . The $s_i^{min,j}$ value is a function of this ADC value, where the function is defined by the sensor manufacturer. The uncalibrated node stores the following bundle in an internal buffer for later reference when generating its calibration curve: $\{ADC_i^{min,j}, s_i^{min,j}, \bar{d}^{min,j}\}$.

In case ground truth and/or uncalibrated nodes follow a sleep schedule to reduce their power consumption, CaliBree is not triggered if either the ground truth nodes or the uncalibrated nodes are in sleep mode (use of a radio wake-up mechanism [33] is possible, but is outside the scope of this paper). If a rendezvous occurs when both ground truth and uncalibrated nodes are awake, the node(s) stay awake until the calibration rendezvous completes. Due to CaliBree's reliance on the uncalibrated node and the ground truth node sharing a common sensing context, which may quickly change due to human mobility, for best performance we recommend prioritizing the CaliBree service over other platform services during the ground truth rendezvous period. We conjecture this is not excessively disruptive for the current sensing device running application since as we show in Section 5.2 the calibration rendezvous lasts for two seconds at most.

4.1 Best Fit Line Algorithm

The best fit line algorithm takes as an input the collected bundles and generates a sensor device's calibration curve in the form $y = f(x)$. Here x represents the ADC input and y the calibrated sensed data output. The best fit line algorithm takes as input the two finite sets $\{x_i = ADC_i^{min}\}$ and $\{y_i = s_i^{min,j} + \bar{d}^{min,j}\}$, for $j \in [1, N]$ where N is the number of encountered ground truth nodes, and produces the calibration equation. In Section 5 we show that fewer than five ground truth nodes are needed by the best fit line algorithm to compute the calibration curve.

By running the calibration algorithm every time an uncalibrated node encounters ground truth nodes we reduce the calibration error (e.g., due to slightly differing



(a) Plot of the disagreement over time between an uncalibrated node and one ground truth node when the uncalibrated node rendezvous multiple times with the ground truth node. (b) Plot of the RSSI measured at the mobile node from packets received by the ground truth node as a function of the distance from the ground truth node.

Fig. 2.

sensing context during the calibration rendezvous) that might be introduced by performing calibration by relying on one or few ground truth nodes.

4.2 Epsilon Adaptation

To make CaliBree more responsive to the dynamics of a mobile sensor network CaliBree adapts the ϵ value in Equation 1 according to the sensing context of the uncalibrated and ground truth nodes. Recall that as ϵ increases the weight given to the newly calculated disagreement value increases relative to the average historical value. CaliBree adaptively changes the value of ϵ when some ground truth readings might introduce large errors and increase the time to convergence of the minimum disagreement estimation. In particular, the value of ϵ is reduced (less weight to the current sample) when:

- the distance between the uncalibrated node and the ground truth node is large;
- the orientation of the uncalibrated node is different than the orientation of the ground truth node(s);
- the time since the ground truth node was last calibrated is large;
- there are hardware differences between the uncalibrated sensing device and the ground truth node(s) (e.g., different ADC scaling), such that comparison of the sensor readings is not possible.

To allow the uncalibrated node to make these determinations, the ground truth sensor includes its current orientation, calibration time stamp, and hardware specifications in each calibration beacon. Location is also stamped in the calibration beacon if a localization system is in place in the network. If this is not the case, we describe in the next section how well distance between nodes can be estimated using the RSSI values of exchanged packets. In our current implementation we do not yet make use of this information to adapt ϵ , but instead we experimentally find that a fixed value of ϵ is 0.05 balances the sensitivity and convergence time of the average consensus algorithm.

4.3 Distance Estimation

CaliBree is triggered by an uncalibrated node when it determines via advertisement packet reception that it is approaching a ground truth node. CaliBree is most efficient, in terms of number of calibration beacons sent, if the protocol is triggered as close as possible to a given ground truth node. This is true since the average consensus algorithm will not converge to the minimum disagreement until the uncalibrated node and the ground truth node share a common sensing range. Calculating incremental disagreements far outside this common sensing range provides no benefit and is therefore wasteful. To determine the relative distance between itself and a ground truth node, the uncalibrated node can leverage a localization system when available, or can perform estimates based on the Received Signal Strength Indicator (RSSI) measurements taken from advertisement packets transmitted by the ground truth node(s).

We run an experiment to verify whether the RSSI can be used as a satisfactory means to infer distance. The ground truth node is placed approximately in the middle of a long hallway. The ground truth node sends packets periodically from which a mobile node extracts and records RSSI information. The mobile mote is carried in one case in a necklace, and in the other case inside a pocket. Figure 2(b) shows a plot of RSSI at the single mobile node versus actual distance from the ground truth node. Error bars indicate the 95% confidence intervals. The ground truth node is placed about 38 meters from one end of an office building hallway ($x=0$). It can be seen that for both the necklace and pocket cases the measured RSSI increases as the mobile node approaches the ground truth node and decreases when the mobile node goes away from the ground truth node. While neither the rising edge or the falling edge of either curve in Figure 2(b) is monotonic (perhaps RSSI can not be used for accurate ranging) we conjecture that RSSI can be effective as a coarse proximity indicator. In future work we will report on determining an appropriate RSSI threshold to serve as a boundary between “close” and “not close”, inasmuch as these labels pertain to triggering Calibree. In our current implementation of CaliBree, ground truth nodes send advertisement packets at the lowest possible transmit power setting, which on the Tmote Sky platform and in our experimental field gives a radio range of approximately 8 meters, to minimize the number of wasted calibration beacons. Approaches such as [15] can be used at close range to refine the ranging estimate within an error of few tens of centimeters.

Another way to determine the uncalibrated-ground truth nodes distance derives from more accurate forms of localization, e.g., on board GPS or hardware location engines [32], which will be largely used by future mobile sensing devices.

5 Experimental Evaluation

In this section we evaluate the performance of the CaliBree protocol using Tmote Sky [1] wireless embedded devices. We implement CaliBree in TinyOS [2] which is currently the de-facto open source standard operating system for embedded experimental wireless sensor systems. We use a testbed of 20 static Tmote Sky nodes, calibrated to provide ground truth sensed data, deployed across the three floors of an office building. The experiments characterize the performance of CaliBree in calibrating the light

sensor of a single¹ Tmote Sky sensing device as it is carried by a human moving at walking speed around the building. The mobility of the human brings the uncalibrated sensor through the sensing ranges of various arrangements of ground truth nodes, as described in the following.

Although we evaluate CaliBree with Tmote Sky platforms, CaliBree could be used equally well with any sensing platform that requires post-factory calibration or calibrated nodes that experience calibration drift over time. We leave a survey of such devices to future work.

5.1 Sensing Contact Time


In the following, we quantify the sensing contact time, i.e., the amount of time two nodes experience the same sensing environment, between a mobile sensor and a static ground truth node. An uncalibrated sensor is carried at human walking speed through the sensing ranges of statically deployed ground truth nodes. In Figure 3(a) the output of the consensus algorithm, the average disagreement, is shown over time as the mobile uncalibrated node rendezvous with three ground truth nodes. The minima in the graph occur when the mobile node is within the common sensing range of each respective ground truth node, where the difference between uncalibrated and calibrated data is minimized. The plot inset in Figure 3(a) is a zoom in of the leftmost minimum. It shows that the amount of time the difference between the uncalibrated and ground truth data is minimum is on average about two seconds. This time interval, the sensing contact time, is relatively short even at human walking speeds when the ground truth node is static, underscoring the importance of efficient messaging and fast consensus convergence in the calibration protocol. Beyond this time interval the uncalibrated node no longer experiences the same sensing environment as the ground truth node and the accuracy of the calibration output decreases.

5.2 Ground Truth Nodes Beacon Rate

As described in Section 4, once an uncalibrated node triggers the calibration routine on the ground truth node, a series of calibration beacons are broadcast by the ground truth node. The broadcast nature of the calibration beacons allow other uncalibrated nodes to trigger their own calibration processes. These beacons allow the uncalibrated node to run the CaliBree consensus algorithm. In our implementation the calibration beacon packet is 18 bytes in size. In determining the best rate at which ground truth nodes should send calibration beacons, one must consider how the resultant consensus update rate (consensus round interval) of the average consensus algorithm impacts the ability of the mobile sensing device to detect when it has left the common sensing range of the ground truth node.

In our implementation, presence in the common sensing range is inferred by detecting the difference from the minimum in a moving window average of consecutive values of \bar{d} from Equation 1. The output of Equation 1 is updated on the reception of every calibration beacon. Therefore, the speed at which a mobile node can detect it has

¹ CaliBree calibration accuracy is designed to be independent of the number of uncalibrated nodes that operate the calibration procedure concurrently.



(a) Plot of the sensing contact time. An uncalibrated node experiences the same sensing environment as the ground truth node for a short time, even under favorable operating conditions (low speed). (b) Plot of the amount of time an uncalibrated node overestimates its presence in the common sensing range of the ground truth node as a function of the calibration beacon transmission rate. The larger the transmission rate, the faster the consensus algorithm reacts to changed conditions and the lower the error.


Fig. 3.

left the common sensing range is proportional to the calibration beacon rate. Figure 3(b) shows a plot of the amount of time an uncalibrated node overestimates its presence in the common sensing range of the ground truth node as a function of the calibration beacon transmission rate. The overestimation presented here is calculated with respect to the common sensing range dwell time inferred when the beacon rate was 65 Hz, the maximum of our tested rates². In order for the mobile node to immediately detect it has left the common sensing range, the calibration beacon rate should be infinite. Under more practical beacon rates and practical conditions (e.g., occasional packet loss), Figure 3(b) shows that when the calibration beacon rate is smaller than 25 beacons per second the consensus algorithm overestimates the dwell time in the common sensing area by 2.5 to 5.5 seconds, persisting in the minimum disagreement state even after leaving the common sensing range, and leading to inaccurate calibration results. For beacon rates larger than 25 pps the consensus algorithm updates the state faster and the estimated common sensing range dwell time is close to that given by the highest tested beacon rate. In our implementation we use a calibration beacon rate of 65 packets/sec. While this high data rate seems to be incurring high cost in terms of bandwidth and energy, we show that the calibration rendezvous lasts for few seconds (less than two) and it would be triggered only at the first time of usage of the sensing device and after long time scale (months or years) if needed due to sensor drift.

5.3 Node Calibration

In this section we show the performance of CaliBree when: *i*) the uncalibrated node comes across multiple co-located ground truth nodes, and *ii*) the uncalibrated node

² 65 packets/sec approaches the maximum possible packet transmission rate of the TinyOS networking stack on the Tmote Sky platform.



(a) Plot of the calibration performance when a set of five ground truth nodes sit within a $1m^2$ area. (b) Plot of the calibration performance when the uncalibrated node sequentially encounters five different ground truth nodes.

Fig. 4.

encounters one ground truth node a time while being carried around the three floor office building.

Co-located ground truth nodes. In this experiment five ground truth nodes are co-located within a $1m^2$ area and are turned on sequentially at intervals of several seconds. The purpose of the experiment is to show the convergence of CaliBree when an uncalibrated node rendezvous with co-located ground truth nodes. The mobility of the uncalibrated sensor node carries it first towards the ground truth cluster, then away until five ground truth rendezvous have completed. After the fifth rendezvous the node remains co-located with the cluster for the duration of the experiment. The result is shown in Figure 4(a) where a comparison between the static ground truth data, the uncalibrated data, and calibrated sensor readings is reported. The values with the flat pattern on the upper part of the figure are the ground truth readings. As the number of active ground truth nodes increases, the calibration curve obtained with the best fitting line algorithm becomes more accurate and the final result when the fifth ground truth node is activated is that the calibrated data lays somewhere in between the highest and lowest ground truth data values. It is possible to see that precise calibration results can be obtained already after the first two rendezvous.

Sparse ground truth nodes. In the case of sparsely placed ground truth nodes across the building, the mobility of the uncalibrated node brings it through the sensing ranges of five different ground truth nodes. After the fifth rendezvous, the uncalibrated node remains near the fifth ground truth node to show that, after having computed the calibration curve from a data set of five data points from different locations, the uncalibrated node has achieved accurate calibration. In Figure 4(b), where again the flat segments in the upper part of the Figure represent the sensor readings of the ground truth nodes, we see that after the first few rendezvous the uncalibrated node starts producing accurate sensor readings. After the fifth rendezvous the CaliBree protocol is not run anymore and by placing the uncalibrated node near the last ground truth node for 200 seconds we

observe that the uncalibrated node is accurately calibrated. In fact, as Figure 4(b) shows, the calibrated data curve overlaps the ground truth curve after the fifth rendezvous.

5.4 Sensor Nodes Orientation

In Section 4 we mention the need for adapting the ϵ value in Equation 1 according to the relationship between the orientation of the nodes. In support of this argument we show in this section how light sensing is impacted by the sensor orientation. We run an experiment with one ground truth node and one uncalibrated node in two different settings: *i*) a “horizontal” configuration where both the sensors face upwards, and *ii*) a “vertical” configuration where the uncalibrated node is tilted by 90 degrees and faces the ground truth node. In both the scenarios the distance between the uncalibrated and ground truth nodes is increased over time by 30 cm each measurement. The difference between the light sensed data of the two nodes for the horizontal and vertical configuration is plotted in Figure 5. It is shown that the vertical configuration doubles the common sensing range (from 30 cm to 60 cm). The experiment confirms the directional nature of the light sensor and shows that when an uncalibrated node is near a ground truth node the mutual orientation of the nodes matters. In the case where the sensors on the nodes have different orientations, a mechanism that reduces the weight of the ground truth sensed readings, like for example reducing the value of ϵ , could improve the performance of the calibration system. The mutual orientation between nodes could be inferred by compass/magnetometer readings for example. Assuming that the future generations of mobile sensor platforms will be equipped with compass sensors is reasonable considering the continuing advances in the embedded sensing technology and the increasing interest in providing smarter sensing devices for people-centric sensing applications [7].




Fig. 5. Plot of the impact of the uncalibrated - ground truth nodes mutual orientation

6 Conclusions

We presented CaliBree, a distributed self-calibration protocol for mobile wireless sensor networks. CaliBree is a very promising technique and, to the best of our knowledge, it is an important first step towards the introduction of calibration algorithms for mobile sensing systems. CaliBree is scalable, robust, and self-adaptive to the dynamics of a mobile sensor network. We demonstrate through experimentation with real sensor devices the existence of the sensing factor which we believe will be one of the drivers in the design and implementations of protocols and applications in the mobile sensing systems domain. We also demonstrate that calibration can be achieved after rendezvous with less than five calibrated nodes. Thus, it can be considered a suitable technique to calibrate mobile sensor nodes in a scalable, lightweight and efficient way.

As part of the future research direction we intend to implement mechanisms to let the uncalibrated nodes infer the context of the ground truth nodes to make clear decisions about the suitability of the ground truth nodes to provide ground truth data. For context we mean those conditions, like for example mutual sensor orientations, that impact the overall calibration performance.

We also intend to implement and validate CaliBree on sensor enabled cellphones and finally, we plan to test the performance of CaliBree when the uncalibrated nodes move at a higher speed than pedestrian to verify the suitability of the protocol to vehicular mobility patterns as well.

Acknowledgments

This work is supported in part by Intel Corp., Nokia, NSF NCS-0631289, and the Institute for Security Technology Studies (ISTS) at Dartmouth College. ISTS support is provided by the U.S. Department of Homeland Security under award 2006-CS-001-000001, and by award 0NANB6D6130 from the U.S. Department of Commerce. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of any funding body.

The authors would like to thank Shane Eisenman for his valuable technical feedback and great help in editing this manuscript.

References

1. Polastre, J., Szewczyk, R., Culler, D.: Telos: Enabling Ultra-Low Power Wireless Research. In: Proc. of IPSN/SPOT 2005, April 25-27 (2005)
2. TinyOS, <http://tinycos.net>
3. Campbell, A.T., Eisenman, S.B., Lane, N.D., Miluzzo, E., Peterson, R.A.: People-Centric Urban Sensing (Invited Paper). In: Proc. of 2nd ACM/IEEE Int'l Conf. on Wireless Internet, Boston (August 2006)
4. Nokia 5500 Sport Phone, <http://www.nokia.com>
5. SensorPlanet, <http://www.sensorplanet.com>
6. Nike + iPod, <http://www.apple.com/ipod/nike>

7. Eisenman, S.B., Miluzzo, E., Lane, N.D., Peterson, R.A., Ahn, G.A., Campbell, A.T.: The BikeNet Mobile Sensing System for Cyclist Experience Mapping. In: Proc. of SenSys 2007, Sydney, Australia, November 6-9 (2007)
8. Eisenman, S.B., Campbell, A.T.: SkiScape Sensing. In: Proc. of 4th ACM Conference on Embedded Networked Sensor Systems (SenSys 2006), Boulder, Colorado, November 1-3 (2006) (Poster abstract)
9. Olfati-Saber, R., Murray, R.: Consensus problems in networks of agents with switching topology and time-delays. *IEEE Trans. Autom. Control* 49(9), 1520–1533 (2004)
10. Culler, D., Estrin, D., Srivastava, M.: Guest Editors' Introduction: Overview of Sensor Networks. *Computer* 34 (2004)
11. Estrin, D., Girod, L., Pottie, G., Srivastava, M.: Instrumenting the world with wireless sensor networks. In: International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2001), Salt Lake City, Utah (May 2001)
12. Estrin, D., Govindan, R., Heidemann, J., Kumar, S.: Next century challenges: scalable coordination in sensor networks. In: Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking, Seattle, USA (1999)
13. Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., Scott, J.: Impact of Human Mobility on Opportunistic Forwarding Algorithms. *IEEE Transactions on Mobile Computing* 6(6), 606–620 (2007)
14. Abdelzaher, T., et al.: Mobiscopes for Human Spaces. *IEEE Pervasive Computing - Mobile and Ubiquitous Systems* 6(2) (April-June 2007)
15. Lowton, M., Brown, J., Finney, J.: Finding NEMO: On the Accuracy of Inferring Location in IEEE 802.15.4 Networks. In: Proc. of 2nd ACM Workshop on Real-World Wireless Sensor Networks (REALWSN 2006), Uppsala, Sweden, June 19 (2006)
16. Hull, B., et al.: CarTel: A Distributed Mobile Sensor Computing System. In: 4th ACM SenSys., Boulder, CO, USA (November 2006)
17. Ganti, R.K., Jayachandran, P., Abdelzaher, T.F., Stankovic, J.A.: SATIRE: a software architecture for smart AtTIRE. In: Proc. of MobiSys 2006, Uppsala, Sweden (2006)
18. CitySense. An Open, Urban-Scale Sensor Network Testbed, <http://www.citysense.net>
19. Tolle, G., et al.: A macroscope in the redwoods. In: Proc. of SenSys 2005, San Diego, California, USA (November 2005)
20. Szewczyk, R., Polastre, J., Mainwaring, A., Culler, D.: Lessons From A Sensor Network Expedition. In: Proc. of the First European Workshop on Sensor Networks (EWSN) (January 2004)
21. Ihler, A.T., et al.: Nonparametric belief propagation for self-calibration in sensor networks. In: Proc. of IPSN 2004, Berkeley, California, USA (2004)
22. Girod, L., Lukac, M., Trifa, V., Estrin, D.: The design and implementation of a self-calibrating distributed acoustic sensing platform. In: Proc. of SenSys 2006, Boulder, CO, USA (November 2006)
23. Taylor, C., et al.: Simultaneous localization, calibration, and tracking in an ad hoc sensor network. In: Proc. of IPSN 2006, Nashville, Tennessee, USA (2006)
24. Whitehouse, K., Culler, D.: Calibration as parameter estimation in sensor networks. In: WSNA 2002: Proc. of the 1st ACM international workshop on Wireless sensor networks and applications, Atlanta, Georgia, USA (2002)
25. Bychkovskiy, V., Megerian, S., Estrin, D., Potkonjak, M.: A Collaborative Approach to In-Place Sensor Calibration. In: Zhao, F., Guibas, L.J. (eds.) *IPSN 2003*. LNCS, vol. 2634, pp. 301–316. Springer, Heidelberg (2003)
26. Balzano, L., Nowak, R.: Blind calibration of sensor networks. In: In Proc. of IPSN 2007, Cambridge, Massachusetts, USA (April 2007)
27. LaMarca, A., et al.: Making Sensor Networks Practical with Robots. In: *Pervasive 2002*, Zurich, Switzerland, August 26-28 (2002)

28. Bychkovskiy, V., et al.: A Collaborative Approach to In-Place Sensor Calibration. In: Zhao, F., Guibas, L.J. (eds.) IPSN 2003. LNCS, vol. 2634, pp. 301–316. Springer, Heidelberg (2003)
29. Moses, R.L., Patterson, R.: Self-calibration of sensor networks. In: The Society of Photo-Optical Instrumentation Engineers (SPIE) Conference (2002)
30. Extech. SuperHeat Psychrometer RH350,
<http://www.extech.com/instrument/products/alpha/datasheets/RH350.pdf>.
31. Extech. Dataloggin Light Meter 401036,
http://www.extech.com/instrument/products/400_500/datasheets/401036.pdf.
32. Chipcon CC2431, <http://www.ti.com/corp/docs/landing/cc2431/index.htm>.
33. Shih, E., Bahl, P., Sinclair, M.J.: Wake on wireless: an event driven energy saving strategy for battery operated devices. In: Proc. of MobiCom 2002, Atlanta, Georgia, USA, September 23-28 (2002)