

Privacy-Preserving Collaborative Blind Macro-Calibration of Environmental Sensors in Participatory Sensing

Jan-Frederic Markert¹, Matthias Budde^{1,*}, Gregor Schindler¹, Markus Klug¹, Michael Beigl¹

¹Karlsruhe Institute of Technology (KIT), TECO / Chair for Pervasive Computing Systems, Vincenz-Prießnitz-Straße, 176131 Karlsruhe, Germany

Abstract

The ubiquity of ever-connected smartphones has led to new sensing paradigms that promise environmental monitoring in unprecedented temporal and spatial resolution. Everyday people may use low-cost sensors to collect environmental data. However, measurement errors increase over time, especially with low-cost air quality sensors. Therefore, regular calibration is important. On a larger scale and in participatory sensing, this needs to be done in-situ. Since for this step, personal sensor data, time and location need to be exchanged, privacy implications arise.

This paper presents a novel privacy-preserving multi-hop sensor calibration scheme, that combines Private Proximity Testing and an anonymizing MIX network with cross-sensor calibration based on rendezvous. Our evaluation with simulated ozone measurements and real-world taxicab mobility traces shows that our scheme provides privacy protection while maintaining competitive overall data quality in dense participatory sensing networks.

Received on 30 June 2017; accepted on 8 August 2017; published on 15 January 2018

Keywords: Participatory Sensing, Location Privacy, Sensor Calibration, Mobile Sensing, Environmental Monitoring, Calibration Rendezvous, Citizen Science, Air Pollution

Copyright © 2018 J.-F. Markert, M. Budde *et al.*, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi:10.4108/eai.15-1-2018.153564

1. Introduction

Air pollutants like ozone or particulate matter pose a danger to people, as they have been proven to cause diseases such as asthma and lung cancer in high concentrations, as well as damage to the environment [1]. The awareness for these problems has increased in societies around the world in the last years. Traditional air pollution monitoring by governmental authorities is mostly done with large stationary equipment and characterized by a high accuracy, reliability and cost and a low temporal and spatial resolution. As an alternative, low-cost gas and particle sensors and platforms have emerged [2, 3], that can be used in different sensing scenarios, in which mobile wireless sensor networks are formed [4]. As these sensors are

not as accurate, reliable and/or stable as classical stationary measurement units, regular calibration is a possible strategy to compensate for systematic error and to prevent quality loss. Among the different calibration techniques that exist, rendezvous-based blind calibration [5] is most suitable for so-called Participatory Sensing scenarios, since low-cost sensors are carried by a large number of people and calibration against reliable, high-quality reference stations is infeasible. However, the proximity-based data exchange approach generally entails privacy issues: Partial traces might be identified based on location information, such as frequently visited places or velocity, network characteristics (e.g. latency), or others.

This paper presents a novel privacy-protecting calibration scheme for participatory environmental sensing. Collaborative blind macro-calibration is combined

*Corresponding author. Email: budde@teco.edu

with several privacy preserving measures, such as private proximity testing, personalized exclusion zones, spatial generalization, pseudonymization and MIX networks.

2. Background and Related Work

A lot of research has been done related either to mobile participatory sensing and its privacy implications, or to the calibration of dynamic sensor systems. However, there is only very limited work combining the two. To the best of our knowledge, *PPCS* [6] is the only privacy-preserving calibration mechanism presented so far. *PPCS* is a MIX-network-based pseudonymization scheme for mobile sensor systems with server-client architecture. It uses so-called *non-blind* calibration, i.e., relies on high quality ground truth reference data. Therefore, *PPCS* can not easily be applied to multi-hop settings with rendezvous-based calibrations. It thus is not so well-suited for end-user participatory sensing. The same applies to a slightly different version of *PPCS*, which was published under the name *PRICAPS* [7]. Another system, *TAPAS* [8], presents approaches to privately select participants for collection tasks – a technique that is not applicable to calibration.

In calibration, so-called *blind* methods achieve calibration gain without ground truth reference data [9]. HASENFRATZ ET AL. proposed a multi-hop calibration scheme for mobile sensors, in which the sensors utilize each others measurements from rendezvous in order to improve the calibration “on-the-fly” [5]. This work incorporates a similar approach.

In privacy preservation, *Proximity Testing* can be used to privately and “continuously report all events of mobile users being within the distance of each other” [10]. This can be used for the task of finding sensors that have measured the same phenomenon at approximately the same time and location. While private one-to-one matchings can reliably be done with pairwise exchanged keys, one-to-many matchings with proximity tests against an unknown number of strange users fail due to the bad scalability, especially regarding key exchange and pairwise distance calculation. Instead, spatial generalization as proposed in [11] can be applied. A *MIX network* [12] is a way to ensure that the network traffic and the corresponding devices are unlinkable.

Our approach combines blind multi-hop calibration with Private Proximity Testing and a MIX network to build a privacy preserving rendezvous-based calibration scheme for participatory sensing scenarios.

3. Preliminary Assumptions

The definition of privacy in this work is to “guarantee that participants maintain control over the release of their sensitive information” [13].

Attacker Model Attackers can be administrators, participants as well as external entities. The attackers’ role is either passive or active: passive attackers may eavesdrop on communication, while active attackers might also compromise servers and communication. Their motivation is assumed to be either malicious or honest-but-curious. Attackers’ objectives can be rather general, e.g., desiring the traces themselves, or more specific, e.g., being interested in the location of a certain person at a specific time. Furthermore, the attackers can enhance their capabilities by utilizing additional information, e.g. publicly available address information from yellow pages or frequently visited places found on social media.

Trust Model The participants trust the devices’ soft- and hardware to correctly implement the scheme. Moreover, they trust the system administrator for choosing reasonable privacy-affecting parameters. The network provider is also trusted, as it already knows the nodes’ approximate location.

Further, the server is not taken as honest or benevolent. Positioning services such as GPS are assumed to utilize passive client applications and thus need not to be trusted.

4. Approach

Our privacy-protecting collaborative blind macro-calibration method can be decomposed into the following separate parts:

1. Sensing
2. Proximity Testing
3. Calibration
4. Upload

Step 1: Sensing The first step naturally is the process of the sensors taking measurements. Each reading consists of essentially three data entries: location, time and the measurement itself.

Measurements with low-cost sensors typically deviate to a certain degree from the ground truth. This measurement error is composed of two parts: (1) The statistical error, caused by random hardware noise or inaccuracies in the measurement apparatus, as well as the statistical nature of the measurement process; (2) the systematic error, depending on multiple factors such as the sensed phenomena and the environment. With low-cost sensors, the systematic error may increase with time, e.g. due to *sensor aging* or other causes [14]. Some sensors, e.g. electro-chemical gas sensors, are more susceptible to this kind of sensor drift than others.

In order to represent the reliability of a sensor’s measurement, we introduce the *validity* (v) as a

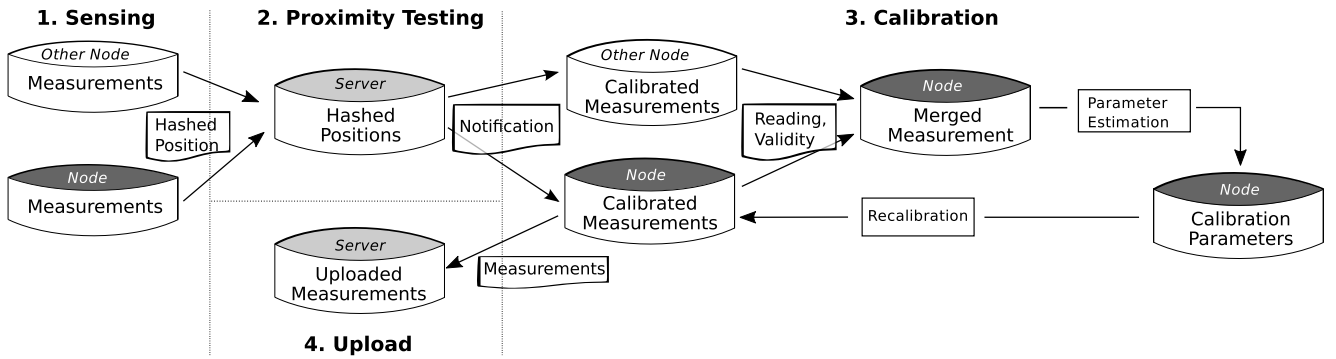


Figure 1. Calibration pipeline: Rendezvous are matched by the server and the sensor nodes' respective measurements and validities are merged to calibration tuples. The nodes estimate the calibration parameters by linear regression and recalibrate the measurements accordingly. Finally, the measurements are uploaded to the server.

meta attribute. The validity ranges between 1 and 0, with an initial value of 1 representing a status of perfect calibration. As the sensors' systematic error increases continuously due to sensor aging, the validity decreases monotonically. The daily decrease depends on a global parameter *dailyValidityLoss* and is calculated as follows:

$$v(t+1) = v(t) * (1 - \text{dailyValidityLoss}) \quad (1)$$

Accordingly, the half-life (*hl*) of the validity can be calculated as

$$\text{hl}(\text{dailyValidityLoss}) = \ln(2) / \text{dailyValidityLoss} \quad (2)$$

For an exemplary half-life of five days, a daily validity loss of 0.138 would be suitable.

While the statistical error is random, calibration is required to estimate the systematic measurement error and subsequently minimize the measurements' deviation from the actual values. The systematic error approximation is based on rendezvous and uses the fact that two spatially and temporally close sensors should measure the same value for a phenomenon. Depending on the homogeneity of the phenomenon different values for the temporal and spatial closeness are required. These so called rendezvous are determined through Private Proximity Testing via a server.

Step 2: Proximity Testing While pairwise distance comparison against a proximity threshold suffices for proximity testing [15], a private implementation requires the exchange of private keys between each pair of nodes and pairwise operations hamper scalability. This complexity is not manageable in a large-scale network of mutually strange nodes. Instead, we utilize a reduction of proximity testing to equality testing via spatial generalization similar to [11]. The positions are mapped to cells via a globally deterministic

function and the resulting cells are compared for equality. Additionally, from the privacy perspective, this coarsens the location and reduces the detail of the released personal information.

The grid characteristics have an impact on the quality of the proximity detection. The basic grid form is a composition of distinct rectangular cells. In order to better approximate a circular neighborhood, multiple mutually offset hexagonal grids can be utilized [15]. Furthermore, the size of the grid cells impact the neighborhood relation: the larger the cells, the greater the rendezvous neighborhood, and the less detailed the released personal information.

The temporal and geographic sampling position is first discretized to the corresponding cell in the grid. As the discretization also involves the temporal dimension, the same geographic position will be in a different cell regularly, preventing frequency analysis attacks to infer population density of certain locations. The distinct cell identifier is then mapped with a cryptographic hash function, making it impossible to recognize the original cell. Depending on the grid, an appropriate hash length needs to be chosen in order to prevent conflicting hashes. Finally, the hash value is uploaded to the server along with a pseudonym in order to query for rendezvous.

The rendezvous detection is done centrally on the server. For every newly uploaded query, the server checks for matches in the set of already uploaded queries. If a match is found, a data exchange between the co-located nodes represented by the pseudonyms is established.

The exchanged data includes the measurements and the respective validity. As the validity is sensor- and thus person-specific, this can have privacy implications. The probability of validities to be relatively distinct is rising with the decrease of the measurement density in the corresponding cells. As a countermeasure,

the validity is discretized according to a global discretization step before exchange. In order to protect the participants privacy during the exchange, a secure communication channel is established via asymmetric encryption. The discretized validity and the calibrated measurement are then sent to the respective rendezvous partner.

The presence of ground truth sensors (i.e. reference stations) in the system is not required, but can improve the calibration performance. Reference stations act as regular nodes, except that they do not move and exhibit no measurement error (constant validity of 1). Their measurements along with the respective cell hashes are also accessed by the server. In case of rendezvous, the server performs the data exchange on behalf of the stations.

In order to be able to protect the participants' privacy also in low density areas, we added tailored sensing in the form of personalized exclusion zones to our scheme. The participants can set up so called sensitive locations, for instance their home or workplace. Subsequently, entries that are located within a given radius of such a sensitive location are discarded. In such areas, subsequent measurements might otherwise be linkable to a trace utilizing prior knowledge on mobility patterns, such as speed and frequent whereabouts.

Step 3: Calibration The computation of the calibration parameters and the calibration application is done locally on the nodes to reduce possible privacy implications as the server could link successive characteristic parameters to re-identify participants.

To perform a calibration, two prerequisites have to be met: (1) At least one of the participating sensors possesses a sufficient validity (*rendezvousValidityThreshold*). This ensures that a calibration actually results in an accuracy improvement. (2) There was no recalibration based on a rendezvous that happened later. As a result of network latencies, the server e.g. might recognize a rendezvous before the data on a different rendezvous that actually happened before that one is processed. If a preceding rendezvous is recognized later, it is therefore discarded as outdated.

The two calibrated measurements retrieved from the rendezvous are merged in order to estimate the unknown ground truth. The validities representing the measurement's reliability are utilized as weights. Thus, the estimation is calculated as a validity-weighted arithmetic mean of the two measurements. The estimation and the rendezvous time constitute the so-called calibration tuple: $\{estimate, time\}$.

After that, the most recent calibration tuples are merged with the sensor's respective uncalibrated measurements $\{r\}$. The number of chosen calibration tuples depends on a global parameter

(*calibrationWindowSize*), and its choice has great impact on the calibration performance: while a higher value yields a more solid calculation basis for regression, the chance of considering already outdated measurements increases.

For the calibration parameters' calculation, different regression methods can be applied depending on the characteristics of the systematic error. For a systematic error best described as polynomial of first order depending on ground truth, linear regression with the method of least squares is utilized for error approximation. However, when the data range is below a threshold (*minimumDataRange*), linear regression can lead to poor results. In this case, we model the systematic error as constant and disregard the present dependency on the ground truth. In both cases, the calibration parameters are updated after error approximation and the following measurements are calibrated accordingly.

In order to account for the calibration gain, the validity is updated after the calibration. While the sensor with the higher validity keeps his validity, the other rendezvous partner adopts the higher value.

Additionally, a so called validity boost is applied, slightly increasing the validity for both. The boost accounts for the calibration gain that not results from calibrating with more valid sensors, but from the fact that rendezvous among uncalibrated nodes still yield positive effects when accumulated for many sensors with different errors. The validity boost, parameterized by a global parameter *validityBoost*, is applied by the following function:

$$v' = \frac{v + \text{validityBoost}}{1 + \text{validityBoost}} \quad (3)$$

This function ensures that the validity never exceeds 1.

Step 4: Upload We implemented different measures to ensure privacy in the data upload step: The participants' privacy with respect to network communication is protected through the use of a MIX network and dynamic pseudonyms. There are different types of MIX networks, that exhibit e.g. different latency characteristics. We assume that a suitable implementation as in [16] is realized.

Pseudonyms are freshly chosen for every communication, in order to prevent any linking. The pseudonym length and the decentral generation mechanism are chosen in a way to prevent pseudonym collusion, which depends on the size of the area to be monitored, the number of nodes and the communication frequency.

Additionally, to prevent attacks based on the upload time, uploads are globally limited to certain points in time defined by a periodic interval. Finally, the uploaded data consists of the calibrated measurement

and the respective time and location. There is no need for an identifier, as the calibration is finished with the upload.

5. Evaluation

We evaluated our scheme by combining simulated ozone measurements with real-world taxicab mobility traces: For the location traces of the simulated mobile nodes we use data from the *epfl/mobility* dataset at CRAWDAD [17]. The dataset contains real-world GPS traces of 537 taxicabs tracked while serving in San Francisco, USA. As the dataset is limited to 22 days, so is our simulation time. The measurement frequency results from the respective GPS logging frequency and amounts to once per minute on average.

For the simulation of the ground truth ozone distribution, we use data from a noise generator based on a free implementation of the *OpenSimplex* noise generation algorithm [18]. The three-dimensional noise has a continuous gradient in all dimensions and nearly no artifacts. We assume ozone to be homogeneous in the order of 30 minutes in time respectively 100 meters in space, in line with [19]. Its amplitude ranges between 0 and 140 ppb (parts per billion) as common ozone concentrations range between 0 ppb and 70 ppb [5] and EU regulations state 90 ppb as information threshold and 120 ppb as alert threshold [20].

In line with e.g. [21], we model the measurement error as the sum of two separate components: The statistical error is modeled with a Gaussian distribution: $n \sim N(0, \sigma^2)$ ppb. Its variance is chosen at the beginning of each day individually for each sensor: $\sigma \sim N(1, 3)$ ppb.

The systematic error b is modeled as a function of the measured value as well as the sensor age. Based on the literature [5, 22] the systematic error linearly depends on the ground truth and increases with time:

$$b(gt, t) = b_0(t) + b_1 * gt \quad (4)$$

where the coefficients are determined by uniform distributions:

$$b_0 \sim U(-9 - \frac{d}{5}, 9 + \frac{d}{5}) \text{ ppb} \quad (5)$$

$$b_1 \sim U(-0.2, 0.2) \text{ ppb} \quad (6)$$

By introducing a temporal dependency for b_1 , sensor aging is incorporated. The coefficients are updated on the beginning of each day, thus t denotes the past full days since deployment. For the systematic error model to be more realistic, the parameters are interpolated between two subsequent days in order to obtain a continuous function of time.

The simulation setup regarding the scheme parameters is shown in Table 1.

Number of mobile nodes	50
Number of reference stations	0
Spatial grid form	basic (quadratic)
Spatial cell length	100 m
Temporal cell length	30 min
Calibration window size	10
Minimum data range	35
Daily validity loss	0.13
Validity boost	0.00003
Validity discretization step	0.0003

Table 1. Scheme parameters of simulation setup.

Calibration Gain The calibration gain, a measure for the effectiveness of a calibration, is computed as the ratio between the difference of the normalized mean squared error (NMSE) between uncalibrated and calibrated measurements, normalized by the uncalibrated NMSE:

$$\text{calibrationGain} = \frac{NMSE_{\text{uncalib}} - NMSE_{\text{calib}}}{NMSE_{\text{uncalib}}} \quad (7)$$

The mean squared error is a standard metric to quantify measurement errors [23]. The NMSE, the mean squared error normalized by the ground truth, is calculated as follows:

$$NMSE = \frac{1}{n} \sum_i^n \frac{(m_i - gt_i)^2}{gt_i^2} \quad (8)$$

summing over all nodes at all time steps.

For the sake of representation NMSE and validity over time in Figure 2(a) and Figure 2(b) were created by temporally binning the data with 150 bins, hence the angular course.

In Figure 2(a) we see a calibration course of a single exemplary node. The NMSE of calibrated measurements (solid line) in comparison to uncalibrated measurements (dotted line) is improved at nearly every point in time.

Figure 2(b) shows the calibration course averaged over all nodes of the same simulation. The NMSE of calibrated measurements increases much slower and remains nearly constant despite sensor aging. Generally, the quality of the calibrated measurements is significantly better than the uncalibrated measurement. While the calibrated NMSE ranges around 0.8, the uncalibrated NMSE fluctuates around 1.4, yielding a calibration gain of 65%.

The periodicity of the uncalibrated error can be explained by the systematic error model, which interpolates between daily chosen parameters. Remarkably, this periodicity vanishes in the calibrated error course, indicating that the remaining error is for the most part of statistical origin.

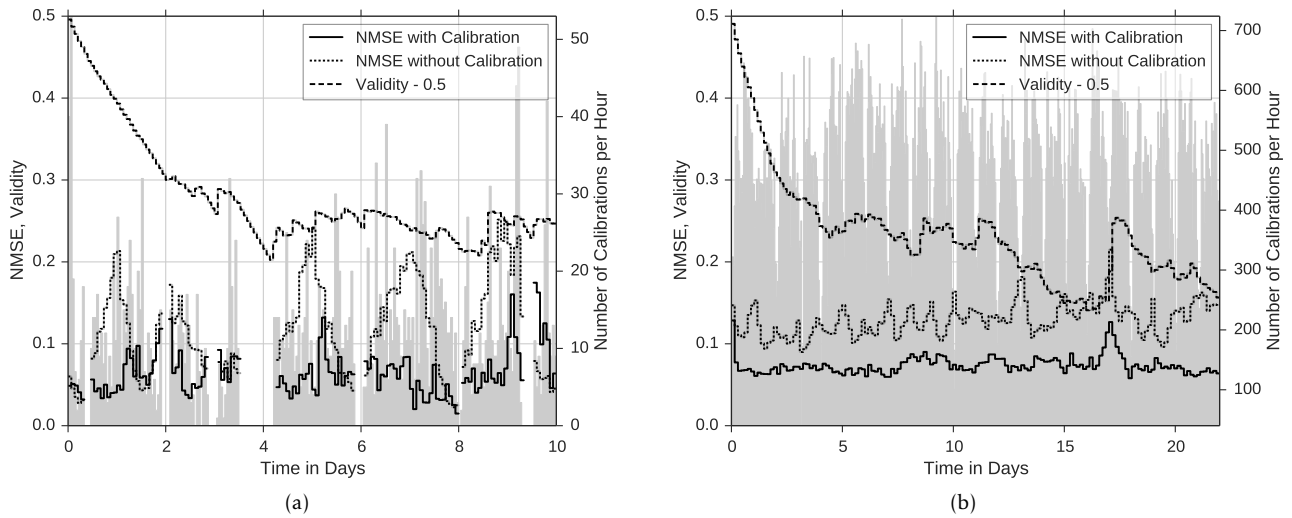


Figure 2. NMSE with and without calibration and validity (note numerical shift). Gray bars in the background represent number of calibrations per hour. (a) Exemplary course of a single node for 10 days. (b) Course averaged over all 50 nodes for 22 days.

System Life Time In Figure 2(a) the exponential validity loss is best recognizable at times where no calibrations are present, especially at day 4. This loss is slowed when calibration processes are happening. At areas of high calibration density, depicted by the gray shaded bars in Figure 2(a), the validity stabilizes, increases or even jumps due to the implemented validity boosts.

The global validity in Figure 2(b) drops with advancing time, as the validity boosts are not able to handle the global loss. Still, in times with a high number of calibrations, the validity rises again as the boosts dominate. Figure 2(b) shows the trend that the validity ranges between 0.65 and 0.75 from day five on, with highs and lows. If the global validity drops under a specific threshold, it is assumed that the system is not able to recover itself and it stops yielding reasonable data. This marks the end of the system's life.

The expectable system life time without calibration is determined by the validity threshold and half-life. If a critical node density and subsequently a sufficient number of calibrations is reached, the life time is significantly prolonged. With a sufficient amount of reference stations, this could enable a hypothetically infinite system life time. The requirements for such a state are highly dependent on the data set and the validity configuration boost.

Reference Stations The impact of reference stations is shown in Figure 3. In order to achieve reliable results, multiple simulations are fused in the diagram. The reference stations were placed strategically at the most frequented locations. It is obvious that the

deployment of more reference stations results in better calibration gains. However, the difference compared to a setting without reference stations diminishes with an increasing number of nodes, resulting from the utilization of rendezvous among imperfect nodes. This shows that our rendezvous-based approach performs best when deployed in a greater scale, and that the accuracy can compete with reference stations.

Identification via Rendezvous The risk of trace reconstruction via rendezvous increases with low

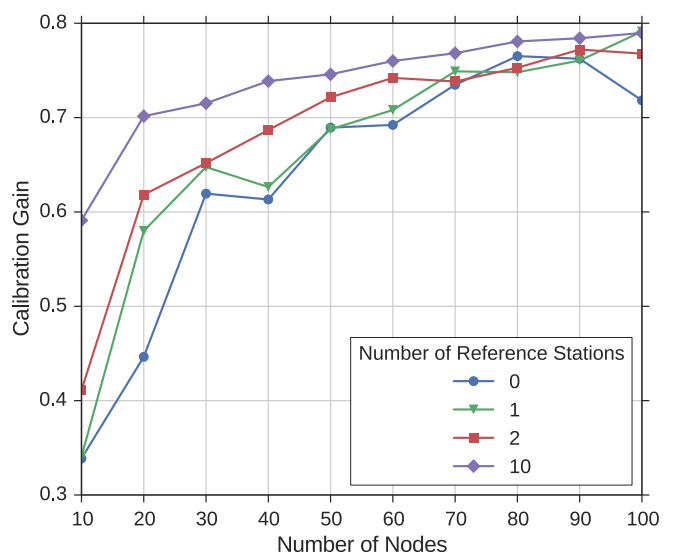


Figure 3. Competitiveness of pure rendezvous-based calibration is shown by impact of reference stations.

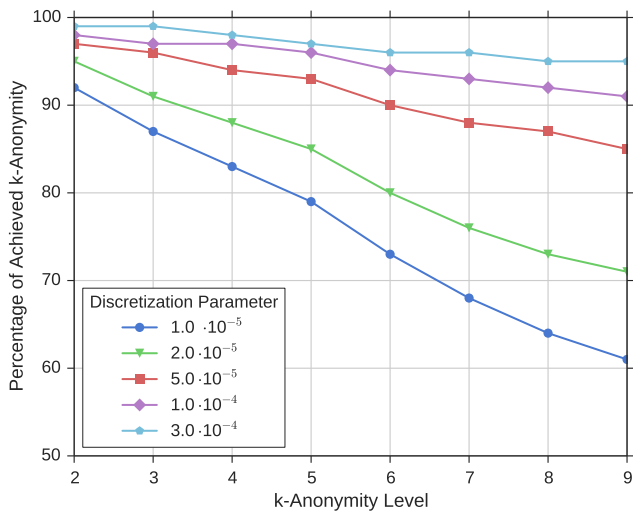


Figure 4. K-anonymity in dependency of validity discretization, depicting the decrease of anonymity with smaller discretization steps.

measurement density and high validity diversity. The validity diversity can be measured with k -anonymity [24]. Here, k represents the number of validities that are discretized to the same value: the lower the density and the higher the diversity, the lower k .

Figure 4 shows the percentage of achieved k -anonymity for different discretization steps. It can be seen that the smaller the discretization steps, the lower the percentage of achieved k -anonymity and consequently the higher the potential privacy risk. As the validity is only utilized as a weight, even a discretization with the largest of the tested steps is not expected to impair the calibration performance significantly. Thus, given a reasonable validity discretization, a successful privacy protection with respect to trace reconstruction via rendezvous is feasible.

6. Discussion

In the course of our evaluation we decided on certain parameters (proximity thresholds, grid size, etc.) for our simulation. While these assumptions were certainly not made arbitrarily and are in line with previous research, we would like to discuss in this section whether our scheme generalizes to other settings or what needs to be adjusted when applying it to different scenarios.

Our spatio-temporal parameters were chosen as previous research suggested for ozone [19] (see above). This choice of course is dependent on the phenomenon (i.e. environmental parameter) that is sensed, respectively its homogeneity and dispersion behavior. Different pollutants or phenomena dictate a different choice of parameters. The same is true for the

general environment: A city with street canyons may call for other proximity thresholds than an open area in nature.

An important prerequisite in this context is that the sensing system somehow should ensure that the same phenomenon is actually being measured in the first place and that measurement takes place under the same circumstances. If, for example, one sensor is used to measure the temperature in open sunlight and another in the shadow, that means they actually are not measuring the same parameter and of course this makes the readings incomparable. Another example would be the usage of air quality sensors in greatly different sampling contexts (e.g. standing vs. riding a motorcycle), in which the difference in speed could lead to an invalid air sample in the latter case. However, such problems need to be addressed at a different level, such as training of participants or outlier detection. Co-location and proper handling of measurement equipment could also be incentivized through the use of game elements [25].

Finally, the actual mobility patterns that are likely to be exhibited in the sensing scenario may differ from the ones used in our simulation. We used taxicab traces as basis because they reflect the movement of real people through a real urban environment. On the other hand, the authors are aware that if sensors were actually deployed on the taxicabs, privacy problems would probably be secondary. Still, the general properties of the mobility data should be realistic for the underlying scenario: Everyday people traversing the public spaces they live in.

All of these are aspects that need to be taken into account, both when designing a Participatory Environmental Sensing application and when determining the parameters for using our scheme to calibrate sensors within them. Nevertheless, we do not see that any of this would invalidate the general applicability of our scheme to different environmental sensing scenarios.

7. Conclusion

In this work, we presented a novel privacy-protecting calibration scheme for participatory environmental sensing that combines collaborative blind macro-calibration with Private Proximity Testing, personalized exclusion zones, spatial generalization, pseudonymization and MIX networks. This enables the calibration of low-cost sensors based on rendezvous and the exchange of measurements between them. We evaluated our scheme on 22 days of simulated data, which combines real-world mobility traces with modeled calibration errors. The results show that our method is capable of achieving significant calibration gain even without reliable reference stations present and protecting the users' location privacy at the same time.

8. Future Work

As simulation-based evaluations always have their weaknesses, an evaluation in a real-world setting would certainly be preferable in the future. On the down side, such an evaluation entails significant resource costs, which makes it a very difficult option.

Other future work could focus on the introduction of additional measures. This could include a modified validity definition, to e.g. incorporate sensor types or the user's measurement reputation. However, such approaches could also entail additional privacy implications that need to be considered. Further enhancement could be achieved by the incorporation of location tags [15] for the prevention of attacks based on spoofed locations.

Acknowledgement. This research has been partially funded by the German Federal Ministry of Education and Research (BMBF) as part of projects *Software Campus* (grant no. 01IS12051) and *SDI-X* (grant no. 01IS15035A). This work is an extended version of the paper “*Private Rendezvous-based Calibration of Low-Cost Sensors for Participatory Environmental Sensing*” [26], published in the proceedings of the *2nd EAI International Conference on IoT in Urban Space (Urb-IoT 2016)*.

References

- [1] WHO (2014), 7 million premature deaths annually linked to air pollution, <http://www.who.int/mediacentre/news/releases/2014/air-pollution>. Accessed 02/25/16.
- [2] KULARATNA, N. and SUDANTHA, B.H. (2008) An Environmental Air Pollution Monitoring System Based on the IEEE 1451 Standard for Low Cost Requirements. *IEEE Sensors* 8(4). doi:10.1109/JSEN.2008.917477.
- [3] BUDDE, M., EL MASRI, R., RIEDEL, T. and BEIGL, M. (2013) Enabling low-cost particulate matter measurement for participatory sensing scenarios. In *MUM'13*.
- [4] BUDDE, M., ZHANG, L. and BEIGL, M. (2014) Distributed, low-cost particulate matter sensing: scenarios, challenges, approaches. *ProScience* 1.
- [5] HASENFRATZ, D., SAUKH, O. and THIELE, L. (2012) On-the-Fly calibration of Low-Cost gas sensors. In *Wireless Sensor Networks, LNCS 7158*. doi:10.1007/978-3-642-28169-3_15, URL http://dx.doi.org/10.1007/978-3-642-28169-3_15.
- [6] WIESNER, K., DORFMEISTER, F. and LINNHOFF-POPIEN, C. (2014) Privacy-Preserving calibration for participatory sensing. In *Mobiquitous'13, Rev. Sel. Papers*.
- [7] WIESNER, K. and DORFMEISTER, F. (2014) PRICAPS: A system for Privacy-Preserving calibration in participatory sensing networks .
- [8] KAZEMI, L. and SHAHABI, C. (2013) Tapas: Trustworthy privacy-aware participatory sensing. *Knowledge and information systems* 37(1): 105–128.
- [9] BALZANO, L. and NOWAK, R. (2007) Blind calibration of sensor networks. In *IPSN'07 (ACM)*: 79–88.
- [10] ŠIKŠNYS, L., THOMSEN, J.R., ŠALTENIS, S., YIU, M.L. and ANDERSEN, O. (2009) A location privacy aware friend locator. In *Advances in Spatial and Temporal Databases*.
- [11] USHIDA, M., YAMAOKA, Y., ITOH, K. and TSUDA, H. (2014) New Privacy-Preserving method for matching location data. In *IMIS'14*: 594–599.
- [12] CHAUM, D.L. (1981) Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* 24(2): 84–90.
- [13] CHRISTIN, D., REINHARDT, A., KANHERE, S.S. and HOLLICK, M. (2011) A survey on privacy in mobile participatory sensing applications. *Journal of Systems and Software* 84(11): 1928 – 1946. doi:<http://dx.doi.org/10.1016/j.jss.2011.06.073>, URL <http://www.sciencedirect.com/science/article/pii/S0164121211001701>.
- [14] BUDDE, M., KÖPKE, M. and BEIGL, M. (2015) Robust in-situ data reconstruction from poisson noise for low-cost, mobile, non-expert environmental sensing. In *ISWC'15*.
- [15] NARAYANAN, A., THIAGARAJAN, N., LAKHANI, M., HAMBURG, M. and BONEH, D. (2011) Location privacy via private proximity testing. In *NDSS*.
- [16] DINGLEDINE, R., MATHEWSON, N. and SYVERSON, P. (2004) *Tor: The second-generation onion router*. Tech. rep.
- [17] PIORKOWSKI, M., SARAFIJANOVIC-DJUKIC, N. and GROSSGLAUSER, M. (2009), The epfl/mobility dataset (v. 2009-02-24), <http://crawdad.org/epfl/mobility/>.
- [18] SPENCER, K. (2014), `OpenSimplexNoise.java`, <http://gist.github.com/KdotJPG/b1270127455a94ac5d19>. Accessed 02/25/16.
- [19] HASENFRATZ, D. (2015) Enabling Large-Scale urban air quality monitoring with mobile sensor nodes .
- [20] THE EUROPEAN PARLIAMENT AND THE COUNCIL OF THE EUROPEAN UNION (2008), Directive 2008/50/ec.
- [21] BYCHKOVSKIY, V., MEGERIAN, S., ESTRIN, D. and POTKONJAK, M. (2003) A collaborative approach to In-Place sensor calibration. In *IPSN'03*.
- [22] XIANG, Y., BAI, L., PIEDRAHITA, R., DICK, R.P., LV, Q., HANNIGAN, M. and SHANG, L. (2012) Collaborative calibration and sensor placement for mobile sensor networks. In *IPSN'12 (ACM)*: 73–84.
- [23] SAUKH, O., HASENFRATZ, D. and THIELE, L. (2015) Reducing Multi-Hop calibration errors in Large-Scale mobile sensor networks. In *IPSN'15 (ACM)*.
- [24] SWEENEY, L. (2002) k-Anonymity: A Model for Protecting Privacy. *Uncertainty, Fuzziness and Knowledge-Based Systems* 10(05): 557–570.
- [25] BUDDE, M., ÖXLER, R., BEIGL, M. and HOLOPAINEN, J. (2016) Sensified gaming – design patterns and game design elements for gameful environmental sensing. In *ACE2016*.
- [26] MARKERT, J.F., BUDDE, M., SCHINDLER, G., KLUG, M. and BEIGL, M. (2016) Private rendezvous-based calibration of low-cost sensors for participatory environmental sensing. In *UrbIoT'16*.