



HAL
open science

Collaborative Sensor Calibration: Convergence Study

Léo Le Taro, Hervé Rivano

► **To cite this version:**

Léo Le Taro, Hervé Rivano. Collaborative Sensor Calibration: Convergence Study. CoRes 2017-2ème Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication, May 2017, Quiberon, France. pp.4. hal-01517793v1

HAL Id: hal-01517793

<https://inria.hal.science/hal-01517793v1>

Submitted on 3 May 2017 (v1), last revised 15 May 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Collaborative Sensor Calibration: Convergence Study

Léo Le Taro[†] and Hervé Rivano

Univ Lyon, INSA Lyon, Inria, CITI, F-69621 Villeurbanne, France

Air pollution is an increasingly concerning issue in urban areas because of its impact on citizens' health. To tackle pollution effectively, accurate monitoring is a must. Precise stations managed by governmental or specialised institutions do exist, but they are both costly and bulky, which limits the potential to deploy them densely. However, recent progress in micro, connected sensors brings new alternative deployment schemes for dense monitoring by low-cost, imprecise sensors. For such a deployment to be relevant relative to urban air quality monitoring aspects, we are concerned with maintaining the system's properties over time. Indeed, one of the major drawbacks of cheap sensors is their drift: chemical properties degrade over time and alter the measurement accuracy. We challenge this issue by designing distributed, online recalibration procedures. We present a simulation framework modelling a mobile wireless sensor network and we assess the system's measurement confidence using trust propagation paradigms. As WSN calibrations translate as information exchange between sensors, we also study means of limiting the number of such transmissions by skipping the calibrations deemed least profitable to the system.

Keywords: Wireless Sensor Networks, Air Pollution Monitoring, Distributed Algorithms, Sensor Calibration

1 Introduction

Urban pollution monitoring is traditionally carried out using governmental air quality stations, precise but costly and sparsely deployed. In large cities, this limits monitoring resolution to neighbourhood level, actual estimation being computed by numerical models. Recent advances in nanotechnologies is giving birth to small, affordable electrochemical sensors. One of the major research challenges is to achieve a higher spatial resolution with dense and mobile deployments of such low-cost sensors. To keep an exploitable measurement accuracy, sensor calibration must be considered [TPS⁺05]. However, low-cost electrochemical sensors such as ones we use to measure NO₂ concentrations degrade over time [KMM⁺15], hence recalibration is required if the system is to remain usable over a long deployment. recalibrating nodes in-place is called non-blind calibration [RBB⁺06]. Ground truths, i.e. reference measurements, are needed to adjust the gain and offset of the low-cost nodes. However, such an approach does not scale because it is infeasible to move a high-quality reference to periodically visit hundreds of sensors.

Another paradigm, denoted blind calibration, assumes unknown ground truths and develops techniques to predict and compensate errors in measurements. Blind calibration methods rely on the underlying signal of interest being either band-limited (i.e. varying smoothly in time and space) or sparse. Yet NO₂ fields prove to be neither, exhibiting large spatiotemporal variations [MLE⁺15], thus negating the possibility to exploit such properties.

Mobile sensing is gaining more attention as recent studies found that a few mobile nodes on well-selected routes can reflect data as accurately as many static ones [CD15]. On top of it, mobile sensor networks offer opportunities for multihop calibration, i.e. freshly calibrated sensors may in turn calibrate others. Work has been conducted to minimise error propagation in multihop calibration [SHT15].

In this paper, we investigate the performance of such multi-hop calibration techniques, in particular their capability to maintain the system's overall accuracy over extended periods of time.

[†]Ph.D. funded by LabEx IMU.

2 Problem Statement

Given a heterogeneous set of sensing units, comprised of precise base stations and low quality sensors, assuming all low quality sensors are initially uncalibrated, we are wondering whether the system converges to an exploitable state, if so, we are concerned with the time required to reach a permanent regime.

Secondly, we are interested in limiting the number of energy-hungry transmissions between sensors, and wonder how saving data exchanges would affect the system's accuracy.

3 Model

We model our system as a discrete time and space process in which sensors follow a stochastic mobility pattern. There are $S = \{1..s\}$ initially uncalibrated mobile sensors randomly moving within a space of $P = \{1..p\}$ positions. Mobile sensors move following a uniform distribution: at any given time, a sensor moves to $x \in P$ with probability $1/p$. Several sensors may share the same position.

To simulate the presence of precise institutional or governmental air quality stations in a real-life urban scenario, let $R \subset P$ be the subset of r positions featuring a static reference station, assumed perfectly calibrated and reliable.

Each mobile sensor a is assigned a trust value $C_a \in [0..1]$, 0 meaning completely inaccurate and 1 meaning completely accurate. Unless recalibrated, mobile sensors degrade and lower their trust following an exponential decay of rate $\gamma \in [0..1]$: $C_a(t) = C_a(0) \cdot e^{-\gamma t}$.

Reference stations retain a constant trust of 1. Initially, all mobile sensors start with a trust of 0.

Mobile sensors are said to have a *rendez-vous* with a reference station or another mobile sensor when, at the same time instant, they stand at the same position as, respectively, a reference station or another mobile sensor.

4 Implementation

Each sensor executes the following pseudo-code algorithm at every time step:

```

Data:  $\gamma, q$ 
 $C \leftarrow C \cdot e^{-\gamma}$ ;
foreach  $n$  in find_neighbours( $p, t$ ) do
  if  $n$  is a reference station then
     $C \leftarrow 1$ ;
  else if  $n$  is a mobile sensor and  $C + q < n.C$  then
     $C \leftarrow n.C$ ;
  end
end

```

Algorithm 1: Sensor trust updating process

We can observe that mobile sensors unconditionally recalibrate to reference stations, however calibration to a peer is conditioned by the variable q . In the real world, each calibration means wirelessly exchanging information. We wish to carry out only meaningful calibrations, with a high benefit. This threshold q allows skipping the calibrations deemed not worthwhile, as they would not lead to a noticeable increase of the trust.

As a corollary, a q of 1 means mobile sensors do not cooperate between each other and only recalibrate to reference stations. Conversely, a q of 0 means that a sensor unconditionally update their trust to the maximum of its neighbors'.

5 Binary Calibration: Analytical Evaluation

Binary calibration is the concept of considering either fully uncalibrated or fully calibrated sensors, with trust values being respectively 0 or 1. In our framework, binary calibration is equivalent to letting γ be zero.

5.1 Sensors do not cooperate

Each sensor has probability r/p at each time step to encounter a reference station and become calibrated. The probability for such a sensor to remain uncalibrated at t follows a geometric progression of rate $1 - r/p$: $P(T > t) = (1 - r/p)^t$.

The system's global average trust can be derived: $\bar{C}(t) = 1 - P(T > t) = 1 - (1 - r/p)^t$

5.2 Sensors cooperate

Being calibrated at t means:

- either having calibrated between 1 and $t - 1$ (noted $T \leq t - 1$);
- or *not* having calibrated between 1 and $t - 1$ *but* having a rendez-vous at exactly t (noted $T = t$).

Both cases are mutually exclusive, let us sum their probabilities, keeping in mind sensors start the simulation uncalibrated (time step 0):

$$P(T = 0) = 0 \quad ; \quad P(T = t) = 1 - \left(1 - \frac{1}{p}\right)^{r+sP(T \leq t-1)}$$

$$P(T \leq t) = P(T \leq t-1) + (1 - P(T \leq t-1))P(T = t)$$

The rest of this paper presents results of our simulation framework whose purposes are to validate our binary calibration analysis, as well as study cases of non-binary calibration, with $\gamma > 0$.

6 Non-binary Calibration: Performance Evaluation

In our system, individual sensors' trust occasionally jump when they recalibrate, then continuously variate to lower values because of decay. How does the average trust of all sensors evolve under these conditions?

Fig. 1 shows that after a short hysteresis-shaped transient phase, the system's average trust converges to a permanent regime where the trust remains stable. Almost all sensors' trusts evolve between 0.5 and 0.95 despite a pessimistic decay rate $\gamma = 10^{-3}$. Cooperative recalibration is therefore able to maintain the accuracy of the system.

The existence of a permanent regime raises the question of how much time is necessary to reach it. We define the time to converge as the time $t_{0.9}$ by which the average trust of the system $\bar{C}(t_{0.9})$ exceeds 0.9. $t_{0.9}$ is plotted against r in Fig. 2, which validates our framework against the analytic evaluations conducted in in Sec. 5 for binary calibration schemes ($\gamma = 0$). It shows that cooperation between sensors have a much stronger impact on the convergence speed than the number of reference stations.




Fig. 1: System state over time

($p = 10000, s = 150, r = 1, \gamma = 10^{-3}, q = 0$)

Curves represent a 95% confidence interval over 100 simulations.

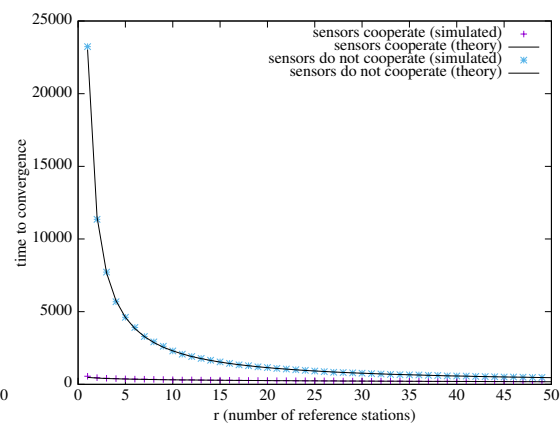


Fig. 2: Time to converge

($p = 10000, s = 150, \gamma = 0$)

Assuming sensors exchange information once for each recalibration, neglecting the cost of polling one's neighbour's trust, Fig. 3 depicts the trade-off between the number of wireless transmissions and the accuracy of the system, i.e. the average trust. The recalibration threshold q is the parameter that restricts the transmission to the case of "useful" recalibrations. Numbers were gathered during the permanent regime. An increase of q degrades the system's trust as expected. However, using even a low q dramatically limits the number of wireless transmissions while sacrificing very little trust.

7 Conclusion and Perspectives

In this paper, we have presented a simulation framework capable of providing insight on what we can expect from auto-calibrating low-cost electrochemical sensor networks. Cooperation between sensors make the system




Fig. 3: Tradeoff between trust and transmission cost
($p = 10000, s = 150, r = 1$)

converge quickly to and maintain a steady average accuracy despite the individual decay of the sensors, without requiring many reference stations. It is also shown that many calibrations opportunities are worth a very marginal improvement of the accuracy.

Current and future work shall integrate more realistic mobility patterns into the model. One could expect that an urban mobility pattern induce correlation between subsets of sensors, hence a less stable “permanent” regime. The impact of the mobility model and the deployment of reference stations on the spatial mapping of trust values should then be investigated. Besides, realistic mobility models should take into account the distance between calibrating and calibrated nodes into the trust propagation function. Finally, calibration theory requires a non-correlated set of simultaneous calibrating and calibrated measurements. Such a set could be collected over a sequence of rendezvous, at the cost of a more complex interpretation of the trust.

References

- [CD15] Carlos T Calafate and Bertrand Ducourthial. On the use of mobile sensors for estimating city-wide pollution levels. In *2015 International Wireless Communications and Mobile Computing Conference (IWCMC)*, pages 262–267. IEEE, 2015.
- [KMM⁺15] Prashant Kumar, Lidia Morawska, Claudio Martani, George Biskos, Marina Neophytou, Silvana Di Sabatino, Margaret Bell, Leslie Norford, and Rex Britter. The rise of low-cost sensing for managing air pollution in cities. *Environment international*, 75:199–205, 2015.
- [MLE⁺15] Sharon Moltchanov, Ilan Levy, Yael Etzion, Uri Lerner, David M Broday, and Barak Fishbain. On the feasibility of measuring urban air pollution by wireless distributed sensor networks. *Science of The Total Environment*, 502:537–547, 2015.
- [RBB⁺06] Nithya Ramanathan, Laura Balzano, Marci Burt, Deborah Estrin, Tom Harmon, Charlie Harvey, Jenny Jay, Eddie Kohler, Sarah Rothenberg, and Mani Srivastava. Rapid deployment with confidence: Calibration and fault detection in environmental sensor networks. *Center for Embedded Network Sensing*, 2006.
- [SHT15] Olga Saukh, David Hasenfratz, and Lothar Thiele. Reducing multi-hop calibration errors in large-scale mobile sensor networks. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, pages 274–285. ACM, 2015.
- [TPS⁺05] Gilman Tolle, Joseph Polastre, Robert Szewczyk, David Culler, Neil Turner, Kevin Tu, Stephen Burgess, Todd Dawson, Phil Buonadonna, David Gay, et al. A macroscope in the redwoods. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 51–63. ACM, 2005.