Network Synchronization and Localization based on Stolen Signals

Christian Schindelhauer[1], Zvi Lotker[2], and Johannes Wendeberg[1]

[1] Department of Computer Science, University of Freiburg, Germany, {schindel, wendeber}@informatik.uni-freiburg.de
[2] Department of Communication Systems Engineering, Ben-Gurion University of the Negev, Israel, zvilo@bgu.ac.il

Abstract. We consider an anchor-free, relative localization and synchronization problem where a set of n receiver nodes and m wireless signal sources are independently, uniformly, and randomly distributed in a disk in the plane. The signals can be distinguished and their capture times can be measured. At the beginning neither the positions of the signal sources and receivers are known nor the sending moments of the signals. Now each receiver captures each signal after its constant speed journey over the unknown distance between signal source and receiver position. Given these nm capture times the task is to compute the relative distances between all synchronized receivers. In a more generalized setting the receiver nodes have no synchronized clocks and need to be synchronized from the capture times of the stolen signals.
For unsynchronized receivers we can compute in time $\mathcal{O}(nm)$ an approximation of the positions and the clock offset within an absolute error of $\mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ with probability $1 - m^{-c} - e^{-c'n}$ (for any $c \in \mathcal{O}(1)$ and some $c' > 0$).
For synchronized receivers we can compute in time $O(nm)$ an approximation of the correct relative positions within an absolute error margin of $\mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ with probability $1 - m^{-c} - e^{-c'n}$. This error bound holds also for unsynchronized receivers if we consider a normal distribution of the sound signals, or if the sound signals are randomly distributed in a surrounding larger disk.
If the receiver nodes are connected via an ad hoc network we present a distributed algorithm which needs at most $O(nm \log n)$ messages in total to compute the approximate positions and clock offsets for the network within an absolute error of $\mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ with probability $1-n^{-c}$ if $m > n$.

1 Introduction

Localization and synchronization are fundamental and well researched problems. In this paper we take a fresh look at this problem. Basic principles in localization are synchronized clocks and anchor points – points with known positions. Some localization methods like DECCA, LORAN and cellular localization have a fixed and known set of anchor points. In other methods, like GPS, the anchor points are moving, but communicate their position to the receivers. In some methods, like WLAN based communication, the position of the anchor points, i.e. the WLAN base stations, needs to be learned. In this paper we do not use anchors at all.

For localization one can use the direction, the runtime or the strength (received signal strength indicator – RSSI) of signals. Runtime based schemes may know the time of arrival (TOA), where the time while the signal travels is known, or the time difference of arrival (TDOA) where the time difference of the signal arriving at two receivers is used.

In this paper we use the time difference of arrival of abundantly available, distinguishable signal sources of unknown location and timing, called "stolen" signals, which can be received at a set of receivers. Assuming that the senders and receivers are on the plane the task is to find the locations of all receivers. Furthermore, we consider the case where the receivers are unsynchronized and try to synchronize their clocks from these stolen signals.

As an application we envisage wireless sensor networks in a noisy area utilizing otherwise interfering signals, e.g. a sensor network with microphones within a swamp with quaking frogs or laptop computers which receive encrypted signals from other WLAN clients and base stations of unknown locations. Then, position information can be used for geometric routing.

In this work we will steal the received sound or radio signals to synchronize the clocks of our network and to compute the locations of our network nodes. For this, we assume that a subset of the senders of the stolen signals are randomly distributed around the receivers. For simplicity we assume a uniform distribution in a disk of same or larger size. The localization and synchronization approach is briefly introduced in [1].

After collecting all the time information from all receivers we want to compute the time offsets and positions of all nodes without knowing where or when the stolen signals are produced. We only assume that we can distinguish stolen signals and they reach all nodes of our network. We are also interested in a distributed algorithm for an ad hoc network minimizing the number of messages.

Problem Setting

Given n synchronized receiver nodes $r_1, \ldots, r_n \in \mathbb{R}^2$ and m signals $s_1, \ldots, s_m \in \mathbb{R}^2$ that are produced at unknown time points $t_{s_1}, \ldots, t_{s_m}$. The signals travel with fixed speed, which we normalize to 1, and are received at time t_{r_i,s_j} for signal s_j and receiver r_i.

Given t_{r_i,s_j} as the only nm inputs we have the following nm equations:

$$t_{r_i,s_j} - t_{s_j} = \|r_i - s_j\|$$

where $(t_{s_j})_{j\in[m]}$, $(s_j)_{j\in[m]}$ are unknown and $(r_i)_{i\in[n]}$ need to be computed.

Since no locations are based at the beginning translation, rotation and mirroring symmetries occur. This can be easily resolved by choosing one receiver as the origin $(0, 0)$, assuming a second receiver lying on the x-axis and a third receiver having a positive y coordinate. Further complications are possible measurement inaccuracies for the time.

The given problem is a non-linear non-convex optimization problem for which no efficient solution for the general case has been known so far. Non-linear non-convex optimization is known to be NP-hard. However, for this specific problem no computational complexity results are known.

2 Related Work

Localization of wireless sensor networks is a broad and intense research topic, where one can distinguish *range-based* and *range-free* approaches.

Range-based approaches include techniques based on RSSI [2][3] or *time of arrival* (TOA, "time of flight") [4][5] to acquire distance information between nodes. The DILOC algorithm uses barycentric coordinates [6]. In many cases a first rough estimation is refined in iterative steps [7][8][9]. Usually, range-based systems require expensive measurement equipment in terms of power consumption and money.

We use the term of *time of arrival* to denote a range measurement by sending a signal to a transponder and measuring the time of the signal flight. In contrast, the term *time differences of arrival* (TDOA) describes the reception of an unknown signal without any given range information. In some contributions a set of receivers is used to locate one or more beacons by evaluation of the TDOA [10][11]. Maybe closest to our problem setting is the iterative solution of Biswas and Thrun [12]. They also implement a distributed approach [13]. A very elegant solution for a fixed number of 10 microphones in three-dimensional space is shown by Pollefeys and Nister [14].

Range-free systems do not require the expensive augmentations that range-based systems do. In the centralized approach [15] the connectivity matrix between nodes is evaluated and a set of distance constraints is generated, which leads to a convex optimization problem. A general disadvantage of centralized algorithms is the lack of scalability and communication overhead. Distributed algorithms avoid this issue [16][17]. A common representation for the communication ranges of nodes in range-free approaches are unit disk graphs (UDG) [18].

In [19] we present a technique for robust distance estimation between microphones by evaluating the timing information of sharp sound signals. We assume synchronized receivers and that signals originate from a far distance, but we have no further information about their location. We consider this a range-based approach because we estimate distances between receivers using the time differences of signals between nodes.

A question that occurs in many wireless sensor network schemes is synchronization. Many synchronization algorithms rely on the exchange of synchronization messages between nodes in the network, assuming that the message delay is symmetric. Another method uses an external radio signal from a base station (e.g DCF77) or a satellite system (e.g. GPS) carrying the current time information. In some approaches a network is assumed to be synchronized in round-based algorithms [20]. An overview of techniques and synchronization issues is given in [21]. Our TDOA-based distance estimation approach [19] implements a synchronization protocol based on the Network Time Protocol algorithm.

Most of the referred algorithms perform effectively in a very specific environment and on safe ground conditions. There are attempts to survey the numerous approaches and to compare them quantitatively [22] and qualitatively [23]. Or the Cramér-Rao bound is calculated to determine the lower variance bounds of a position estimator [24][25][26]. Few is actually known about the general solvability of localization problems in wireless sensor networks. Stéwenius examines the required minimum of microphones and signal sources for convergence towards unique solutions [27]. Eren et al. inspect the uniqueness of ranged networks by analyzing the graph rigidity [28].

3 Estimating Distances

The localization problem that we face is vastly overconstrained for large n and m. While we have $n+m$ unknown receiver and sender locations, m unknown signal time points and n unknown clock offsets between the receivers, we face nm equations on the other side. So, the clue for an efficient solution of the problem is to concentrate on the most helpful information.

For this we consider only two receivers. As we have pointed out in [19] it is possible to estimate the distance between two receiver nodes if the signals are uniformly distributed on a circle around the receivers at a large distance. Here, we show that this method also results in a reasonable estimation if the signals are distributed in the same disk where the receivers lie.

Max-Min-Technique Given two vertices i, j $(1 \leq i < j \leq n)$ and the relative time differences of the stolen signals: $t_{r_i,s_k} - t_{r_j,s_k}$ for all stolen signals $s_1, \ldots, s_m$, we compute the *estimated distance $d_{i,j}$ between i and j* as

- $d_{i,j} := \max_k\{|t_{r_i,s_k} - t_{r_j,s_k}|\}$ if the receivers are synchronized and as
- $d_{i,j} := \frac{1}{2}\left(\max_k\{t_{r_i,s_k} - t_{r_j,s_k}\} - \min_k\{t_{r_i,s_k} - t_{r_j,s_k}\}\right)$ if the receivers are not synchronized. The estimated relative time offset will be computed using the time signal $k^* := \arg\max_k\{t_{r_i,s_k} - t_{r_j,s_k}\}$. Then, $t_{r_i,s_{k*}} - t_{r_j,s_{k*}} - d_{i,j}$ yields the approximation of the correction for the clocks at i and j.

Clearly, this estimation is only an approximation. But a surprisingly good one. First, note that in both cases the estimation is always upper-bounded by the real distance: $\|r_i - r_j\| \geq d_{i,j}$. We now describe a sufficient condition for the accuracy of the estimator. For this we define the ϵ-critical area.

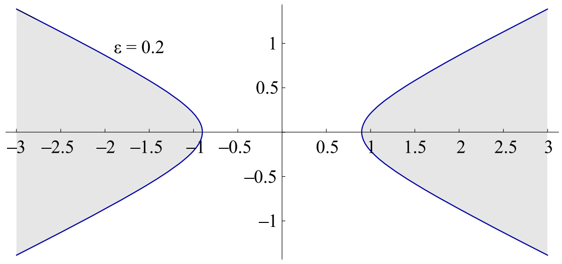

Fig. 1. The 0.2-critical areas of two nodes at $(-1, 0)$ and $(1, 0)$ are on the left and right side of the hyperbolas.

Definition 1. *The ϵ-critical area of two nodes (u, v) is the set of points p in the plane where*

$$\|u - v\| - (\|p - v\| - \|p - u\|) \ \leq \epsilon \ .$$

This convex area is bounded by a hyperbola containing the point u, see Fig. 1. If in this critical area signals are produced, then the distance estimation is accurate up to an absolute error of ϵ.

Lemma 1. *If in both of the ϵ-critical areas of (u, v) and (v, u) signals are produced, then the Max-Min distance estimation $d_{u,v}$ is in the interval $d_{u,v} \in [\|u - v\| - 2\epsilon, \|u - v\|]$. The time offset between the clocks of u and v can be computed up to an absolute error margin of 2ϵ.*

If at least in one of the ϵ-critical areas of (u, v) and (v, u) a signal is produced, then for synchronized receivers the Max-Min distance estimation $d_{u,v}$ is in the interval $d_{u,v} \in [\|u - v\| - \epsilon, \|u - v\|]$.

These signals can be found in time $O(m)$.

Proof. The proof of the accuracy of the distance estimators follows from the definition of the critical areas. For the accuracy of the time offset consider that one clock u is assumed to be correct, then the other node's clock offset is chosen such that the signal arrives later at time $d_{u,v}$ if the signal was detected at the ϵ-critical area of u.

The best signals can be found by computing the minimum or maximum of the differences of the time points at the receivers u and v. $\square$

Lemma 2. *For two receivers u, v with $\ell := \|u - v\|$ the intersection of the ϵ-critical area (v, u) of a disk with center $\frac{1}{2}(u + v)$ and radius r has*

- *at least an area of* $\min\{\pi\ell^2, \frac{1}{2}\epsilon^2\}$ *if* $r = \ell$ *and*
- *at least an area of* $\min\{\pi r^2, (r - \ell)^2\sqrt{\epsilon/\ell}\}$ *if* $r > \ell$.

Since the critical areas are rather large there is a good chance that a signal could be found in one of these areas.

Theorem 1. *For m stolen signals the Max-Min distance estimator for two receiver nodes u, v with distance $\ell := \|u - v\|$ within the disk with center $(0,0)$ and radius 1 outputs a result $d_{u,v}$ with $d_{u,v} \in [\|u - v\| - \epsilon, \|u - v\|]$ with probability $1 - p$, where for ϵ and p we have:*

1. *If u and v are unsynchronized and the m signal sources are uniformly distributed in the unit disk we have $\epsilon = \mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ and $p = \frac{1}{m^c}$ for any $c > 1$.*
2. *If u and v are unsynchronized and the m signal sources are independently normal distributed with mean $(0,0)$ and variance 1 we have $\epsilon = \mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ and $p = \frac{1}{m^c}$ for any $c > 1$.*
3. *If u and v are unsynchronized, u and v are not close to the unit disk boundary, i.e. $|u| < 1-k$ and $|v| < 1-k$ for some constant $k > 0$, and the m signal sources are uniformly distributed in the unit disk we have $\epsilon = \mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ and $p = \frac{1}{m^{ck^2}}$ for any $c > 1$.*
4. *If u and v are synchronized, u or v are not close to the unit disk boundary, i.e. $|u| < 1 - k$ or $|v| < 1 - k$ for some constant $k > 0$, and the m signal sources are uniformly distributed in the unit disk we have $\epsilon = \mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ and $p = \frac{1}{m^{ck^2}}$ for any $c > 1$.*

Using this information we do not need to consider the signals at all, again. From now on, we will only use the distance estimation information and compute the locations of the receiver nodes. Now the goal is to avoid any further loss of precision when we compute coordinates out of the distance estimates.

4 Centralized Localization and Synchronization

Now we discuss how the distance estimation can be converted into cartesian coordinates without increasing the inaccuracy by more than a constant factor. The usual approaches to reconstruction of node positions from distances are iterative force-directed algorithms [29] or non-linear optimization schemes to minimize a function

$$\min_{r_i, r_j} \left(\sum_{i=1}^{n} \sum_{j=i+1}^{n} \|r_i - r_j\|^2 - d_{i,j}^2 \right)$$

where $d_{i,j}$ denotes the distances yielded by the Max-Min approximation technique. Examples are the gradient descent method, Newton's method [30] or the Levenberg-Marquardt algorithm. The common problem of all these methods is their lack of reliability. They cannot guarantee successful convergence to the correct network topology and they are prone to local minima of the error function. In such cases the induced error is disproportionately higher than one would expect from changes in parameters.

We require an algorithm with constant propagation of error where the induced uncertainty can be bounded below a function of the input error ϵ. For this we have to consider the rigidity and precision.

- Rigidity: If the number of receivers is small or the accuracy is high, then different topologies are valid solutions to the problem. This problem is known as the rigidity problem [28]. We will prove that our distance estimations are so precise that this problem can occur only with a very small probability.
- Precision: In some situations small measurement errors of the distance result in much larger changes of the coordinates. Sometimes, there seems to be no valid solution. We will prove that for any triangle (with non-collinear points), the problems can be solved if the distance estimation error is small enough. The coordinates will suffer from a higher estimation error. This increase can be bounded by a constant factor if the triangles are not too extreme. Furthermore, the probability that such receivers exist grows exponentially with the number of receivers.

Assuming that all distance estimations are precise up to an additive error of at most ϵ_0 we will present algorithms which produce an output with an additive error of at most $\epsilon = \mathcal{O}(\epsilon_0)$ with probability $1 - e^{-cn}$ for n receiver nodes and a constant $c > 0$. In this section we assume that a central node has complete knowledge of the nm capture times of all nodes. In fact our algorithms use only the $\binom{n}{2}$ distance estimations $d_{u,v}$ for all receiver nodes u, v.

Our basic method for localization is *bilateration with a symmetry breaker*. Given two anchor points u, v where $u = (0, 0)$ and $v = (d_{u,v}, 0)$ and the estimated distances $d_{u,p}, d_{v,p} \geq 0$ we want to compute the location of a point p such that $\|u - p\| = r_1$ and $\|v - p\| = r_2$. We know that the given distances are only an approximation of the real distances which could be longer by an additive term of ϵ_0. Of course, there are two symmetric solutions for p. So, we also assume a third anchor point w with given coordinates, called symmetry breaker, which is used for deciding which solution is valid.

At the beginning we assume that d, r_1, r_2 are the correct values of the triangle distances and compute the coordinates of p by

$$p_{1,2} = (d_{u,p} \cos \alpha_u, \pm d_{u,p} \sin \alpha_u) , \quad \text{where} \quad \cos \alpha_u = \frac{d_{u,p}^2 - d_{v,p}^2 + d_{u,v}^2}{2 d_{u,p} d_{u,v}} .$$

However, this method fails if $d_{u,p}, d_{v,p}, d_{u,v}$ do not fulfill the triangle inequality. Then, we are not able to find the locations. For deciding between p_1 and p_2 we use the symmetry breaker w. If $| \, \|w - p_1\| - d_{w,p}| \leq | \, \|w - p_2\| - d_{w,p}|$ we choose p_1, and p_2 otherwise.

Using only bilateration it is not possible to locate all points in the plane. But if we have three anchor points we have three possibilities to apply bilateration. The third point is used as a symmetry breaker. Actually every triangle can be used for the localization as long as the points are not collinear.

Theorem 2. *For every set of non-collinear points u, v, w and for every disk D of radius r containing u, v, w there exists an $\epsilon_0 > 0$ such that each point in D*

can be located with an absolute error of $\epsilon' \leq c_{u,v,w,r} \cdot \epsilon$*, if* $\epsilon \leq \epsilon_0$*. Here,* ϵ *is the precision of the distance measurements and* $c_{u,v,w,r}$ *is a constant which depends solely on* u, v, w *and the disk radius* r*.*

So, the distance information provides enough rigidity if the number of signals is large enough, since the larger the number of signals the smaller the starting error ϵ_0.

The factor $c_{u,v,w,r}$ describes the loss of quality of the localization depending on u, v, w and r.

1. $c_{u,v,w,r}$ increases with growing disk radius r.
2. $c_{u,v,w,r}$ decreases with growing minimum edge length.
3. If a triangle angle of u, v, w approaches 0 or π, then $c_{u,v,w,r}$ also increases.

So, best results can be achieved if the edge lengths are large and if they are the same. Since we are only interested in asymptotic results we use the following corollary.

Corollary 1. *Fix some* $0 < \delta_1 < \pi/6$*,* $\delta_1 < \delta_2 < \pi$ *and* $0 < r_0 \leq r$*. Then there is some* $\epsilon_0 > 0$ *and a constant* c *such that all triangles, where all inner angles are in the interval* $[\delta_1, \delta_2]$ *and all edge lengths are at least* r_0*, can be used for localization of all points in the disk of radius* r *within an accuracy of* $c \cdot \epsilon$*. This localization is based on distances which are only known with some absolute precision of* $\epsilon < \epsilon_0$*.*

In the case of *unsynchronized receivers* we experience an accuracy of $\epsilon = \mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ with high probability. It remains to find the best base triangle based on the distance estimations. This can be done by computing all $\binom{n}{2}$ distances within time $\mathcal{O}(n^2 m)$ and testing all $\binom{n}{3} = \mathcal{O}(n^3)$ triangles.

Since we look for any triangle obeying the properties of inner angles and some reasonable minimum edge length r_0, one can use a faster approach.

Algorithm 2 Finding a base triangle

```
Start with an arbitrary node s
Find the node u maximizing d_{s,u}
Find the node v maximizing d_{u,v}
Find the node w maximizing min{d_{u,w}, d_{v,w}}
Use u, v, w as a base triangle for trilateration of all other points
```

Note that each step of this algorithm can be solved by estimating $\mathcal{O}(n)$ distances. Each distance estimation needs time $\mathcal{O}(m)$. So the overall running time is $\mathcal{O}(nm)$.

Using this algorithm a centralized algorithm can solve the localization in nearly all cases for sufficiently large m and n.

Theorem 3. *For n receivers and a subset of m signals produced uniformly distributed in a disk the nodes can be synchronized and localized in time $\mathcal{O}(m)$ with an accuracy of $\epsilon = \mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ in running time $\mathcal{O}(m+n)$ with probability $1 - \frac{1}{m^c} - e^{-nc'}$ for any c and some $c' > 0$.*

1. *If the m signals are produced independently with a Gaussian normal distribution with mean $(0, 0)$ and variance 1 or*
2. *if the m signals are independently and uniformly produced in a disk with radius $r > 1$ and center $(0, 0)$*

then the receiver nodes can be be synchronized and localized in time $\mathcal{O}(m+n)$ with a maximum error of $\epsilon = \mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ with probability $1 - \frac{1}{m^c} - e^{-nc'}$ for any c and some $c' > 0$.

The proof follows by combining the distance estimation with the triangulation results. The exponential bound for the receivers follows from the observation that there is a constant probability for three receiver nodes to satisfy the triangle property. Adding three more receiver nodes independently results in the multiplicative decrease of this failing probability, thus leading to an exponential probability function with respect to n.

Of course these bounds also hold for synchronized receivers. We now concentrate on the case where signals and receivers are produced in the same disk, since the accuracy can be increased considerably for this case. The key point is to find a second base triangle where all receiver nodes have a constant distance to the boundary of the disk. Then, the distance estimations to all other points are accurate to an error of $\mathcal{O}(\frac{\log^2 m}{m^2})$.

Algorithm 3 Finding an inner base triangle

1: Find a base triangle u_0, v_0, w_0
2: Compute the coordinates of all receiver nodes with low precision
3: Based on this information find a base triangle u, v, w with a minimum distance of $\frac{1}{8}$ to the border of the disk

The distance $\frac{1}{8}$ is an arbitrary non-zero choice. Decreasing this distance will increase the distance estimation error but will increase the probability of finding such triangles.

Theorem 4. *For n synchronized receivers and a subset of m signals produced uniformly distributed in a disk the nodes can be localized in time $\mathcal{O}(m)$ with an accuracy of $\epsilon = \mathcal{O}\left(\frac{\log^2 m}{m^2}\right)$ in running time $\mathcal{O}(m+n)$ with probability $1 - \frac{1}{m^c} - e^{-nc'}$ for any c and some $c' > 0$.*

The proof is analogous to the one above.

Our centralized Algorithm 2 can be extended for distributed localization, as for example in a wireless sensor network. In such an ad hoc network each node broadcasts its capture times, which costs $\mathcal{O}(nm)$ messages for each receiver node with a total communication complexity of $\mathcal{O}(n^2m)$. The upper bound for the message broadcast in the distributed network is described by the following theorem.

Theorem 5. *For n receiver nodes and $m > n$ signals a distributed algorithm requires $\mathcal{O}(nm \log n)$ messages and computes the coordinates and the clock offset of all receiver nodes within an error of $\epsilon = \mathcal{O}\left(\sqrt{\frac{\log m}{m}}\right)$ with probability $1 - n^{-c}$ for any constant c, if the receiver and signal nodes are independently, randomly, and uniformly distributed in the same disk.*

5 Outlook

While our focus is the localization and synchronization of the receiver nodes, it is straight-forward to determine the time and position of the signals using the receivers as anchor points.

Synchronization based on stolen signals can provide a helpful feature for wireless sensor networks or ad hoc networks. While the speed of light could be too fast for an accurate localization, it is always a good source for synchronizing clocks. Our approach solves the problem that distant nodes might suffer from the delay of the synchronization signal since we compensate with other stolen signals.

A remarkable property of this localization problem is the decrease of complexity with increasing problem size. If the number of receivers and stolen signals increases, the precision of the approximation improves, while the algorithm's running time remains linear. In case the number of signals is too large, a set of node can agree to consider only a random subset of signals. While the accuracy decreases, the number of messages and the computational effort can be reduced to fit the wireless network's capabilites.

On the other side, the problem is very complex when few signals and receivers are known. We have observed this for approximation methods based on iterative improvement of local solutions like force-directed algorithms and gradient-based search in previous work. For four receivers, which is the minimum number of receivers in the plane to solve this problem, all considered methods can run into local minima. In the successful cases they converge slower to the solution than in large scenarios. If one also reduces the number of signals to the absolute minimum of six signals in the plane (where we conjecture that the solution is unique for non-degenerated input) it appears to be the hardest problem setting. It is an open problem how to solve this localization problem for four receivers and six signals in the plane.

Acknowledgment

This work has partly been supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) within the Research Training Group 1103 (Embedded Microsystems).

References

1. C. Schindelhauer, Z. Lotker, and J. Wendeberg. Brief Announcement: Network Synchronization and Localization Based on Stolen Signals. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011* (to appear), 2011.
2. M. L. Sichitiu and V. Ramadurai. Localization of Wireless Sensor Networks with a Mobile Beacon. In *Proceedings of the First IEEE Conference on Mobile Ad-hoc and Sensor Systems*, pages 174–183, 2004.
3. B. Ferris, D. Hähnel, and D. Fox. Gaussian Processes for Signal Strength-Based Location Estimation. In *Proceedings of Robotics: Science and Systems Conference (RSS)*, 2006.
4. N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket Location-Support System. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43, 2000.
5. A. Savvides, C. C. Han, and M. B. Strivastava. Dynamic Fine-Grained Localization in Ad-Hoc Networks of Sensors. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 166–179. ACM, 2001.
6. U. A. Khan, S. Kar, and J. M. F. Moura. Distributed Sensor Localization in Random Environments using Minimal Number of Anchor Nodes. *IEEE Transactions on Signal Processing*, 57(5):2000–2016, 2009.
7. P. Biswas and Y. Ye. Semidefinite Programming for Ad Hoc Wireless Sensor Network Localization. In *IPSN '04: Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 46–54, New York, NY, USA, 2004. ACM.
8. C. Savarese, J. Rabaey, and K. Langendoen. Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks. In *USENIX technical annual conference*, volume 2. Monterey, CA, 2002.
9. A. Savvides, H. Park, and M. B. Srivastava. The Bits and Flops of the N-hop Multilateration Primitive For Node Localization Problems. In *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, page 121. ACM, 2002.
10. D. J. Torrieri. Statistical Theory of Passive Location Systems. *IEEE Transactions on Aerospace and Electronic Systems*, AES-20(2):183–198, March 1984.
11. M. D. Gillette and H. F. Silverman. A Linear Closed-Form Algorithm for Source Localization From Time-Differences of Arrival. *Signal Processing Letters, IEEE*, 15:1–4, 2008.
12. R. Biswas and S. Thrun. A Passive Approach to Sensor Network Localization. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004. (IROS 2004).*, volume 2, pages 1544–1549, 2004.
13. R. Biswas and S. Thrun. A Distributed Approach to Passive Localization for Sensor Networks. In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, page 1248. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.

14. M. Pollefeys and D. Nister. Direct computation of sound and microphone locations from time-difference-of-arrival data. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2008.*, pages 2445–2448. IEEE, 2008.
15. L. Doherty and L. El Ghaoui. Convex Position Estimation in Wireless Sensor Networks. In *INFOCOM 2001. Proceedings of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, volume 3, pages 1655–1663. IEEE, 2002.
16. S. N. Simic and S. Sastry. Distributed Localization in Wireless Ad Hoc Networks. *UC Berkeley ERL report*, 2001.
17. Z. Lotker, M. M. de Albeniz, and S. Pérénnes. Range-Free Ranking in Sensors Networks and Its Applications to Localization. In *Ad-Hoc, Mobile, and Wireless Networks*, volume 3158 of *Lecture Notes in Computer Science*, pages 158–171. Springer, 2004.
18. A. Kröller, S. Fekete, D. Pfisterer, and S. Fischer. Deterministic boundary recognition and topology extraction for large sensor networks. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1000–1009. ACM, 2006.
19. T. Janson, C. Schindelhauer, and J. Wendeberg. Self-Localization Application for iPhone using only Ambient Sound Signals. In *Proceedings of the 2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 259–268, Nov. 2010.
20. A. Anta, M. Mosteiro, and C. Thraves. Deterministic Recurrent Communication and Synchronization in Restricted Sensor Networks. *INRIA RR*, 486277, 2010.
21. B. Sundararaman, U. Buy, and A.D. Kshemkalyani. Clock Synchronization for Wireless Sensor Networks: A Survey. *Ad Hoc Networks*, 3(3):281–323, 2005.
22. K. Langendoen and N. Reijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Computer Networks*, 43(4):499–518, 2003.
23. G. Stupp and M. Sidi. The expected uncertainty of range-free localization protocols in sensor networks. *Theoretical Computer Science*, 344(1):86–99, 2005.
24. A. Savvides, W. L. Garber, R. L. Moses, and M. B. Srivastava. An Analysis of Error Inducing Parameters in Multihop Sensor Node Localization. *IEEE Transactions on Mobile Computing*, pages 567–577, 2005.
25. S. Dulman, P. Havinga, A. Baggio, and K. Langendoen. Revisiting the Cramer-Rao Bound for Localization Algorithms. *4th IEEE/ACM DCOSS Work-in-progress paper*, 2008.
26. R. L. Moses, D. Krishnamurthy, and R. M. Patterson. A Self-Localization Method for Wireless Sensor Networks. *EURASIP Journal on Advances in Signal Processing*, pages 348–358, 2003.
27. H. Stewénius. *Gröbner Basis Methods for Minimal Problems in Computer Vision.* PhD thesis, Lund University, April 2005.
28. T. Eren, D. K. Goldenberg, W. Whiteley, Y. R. Yang, A. S. Morse, B. D. O. Anderson, and P. N. Belhumeur. Rigidity, Computation, and Randomization in Network Localization. In *Twenty-third annual Joint Conference of the IEEE Computer and Communications Societies*, volume 4, pages 2673–2684. IEEE, 2004.
29. A. Efrat, D. Forrester, A. Iyer, S. G. Kobourov, C. Erten, and O. Kilic. Force-Directed Approaches to Sensor Localization. *ACM Transactions on Sensor Networks (TOSN)*, 7(3):1–25, 2010.
30. G. C. Calafiore, L. Carlone, and M. Wei. A Distributed Gauss-Newton Approach for Range-based Localization of Multi Agent Formations. In *2010 IEEE International Symposium on Computer-Aided Control System Design (CACSD)*, pages 1152–1157. IEEE, 2010.