Anchor-free TDOA Self-Localization

Johannes Wendeberg[1], Fabian Höflinger[2], Christian Schindelhauer[1], and Leonard Reindl[2]

[1] *Department of Computer Science, University of Freiburg, Germany*, {wendeber, schindel}@informatik.uni-freiburg.de
[2] *IMTEK, University of Freiburg, Germany*, {fabian.hoeflinger, reindl}@imtek.uni-freiburg.de

***Abstract*—We present an approach for the localization of passive receiver nodes in a synchronized communication network. The positions of the nodes are arbitrary and unknown. The only source of information is the time differences of arrival (TDOA) when environmental sound or ultrasound signals are received. The discrete signals occur at unknown positions and times, but they can be distinguished. The goal is to determine the relative positions of all receiver nodes and implicitly the positions and times of the environmental signals.**

Our novel approach solves iteratively a non-linear optimization problem of time differences of arrival by a physical spring-mass simulation. Here, our algorithm shows a smaller tendency to get stuck in local minima than a non-linear least-squares approach.

The approach is tested in numerous simulations and in a real-world setting where we demonstrate and evaluate a tracking system for a moving ultrasound beacon without the need to initially calibrate the positions of the receivers. Using our approach we estimate the trajectory of a moving model train with a precision in the range of centimeters.

Keywords: TDOA, anchor-free, localization, ultrasound tracking

I. Introduction

Localization using infrastructures like global navigation satellite systems or GSM multilateration depends on the availability of external systems. These can fail due to environmental conditions (indoor locations, in the forest, on mountains), or they could be deactivated for political reasons. Besides, most infrastructural location services are too imprecise for indoor localization.

Our anchor-free approach does not rely on external infrastructures. We address the problem of self-localization of four or more receivers using the time differences of arrival (TDOA) of acoustic signals from the environment – of which we do not know the positions of origin. A sound source could be a finger snapping, coughing, or the tick sound of a metronome. All we assume is that we can distinguish the sounds.

TDOA data of audible sound can be obtained by discrete timestamping [1], [2] or by cross correlation of signals [3]. Ultrasound is used in [4], [5]. Usually, the receivers' positions are known. Then, estimating a sender's position using time differences of arrival can be addressed in closed form equations or by iterative approaches. Moses et al. use TDOA with additional angle information to locate unknown sender and receiver positions [6]. This would require expensive receiver arrays or directed receivers.

Localization without anchors and relying only on TDOA can be solved if assumptions on the signal positions are made, e.g. the signals originate from far away [2], [7]. A very elegant approach was proposed in [8] where the special case of ten microphones in space is solved in a linear approach.

Close to our problem setting is the approach of Biswas and Thrun [1]. No assumptions about the signal positions are required and only TDOA information is used to iteratively refine a Bayesian network. However, the correct solution cannot be found in every case.

We present an iterative approach based on a spring-mass simulation to solve the problem of anchorless localization with only TDOA information with a probability of more than 99 %.

II. Iterative Cone Alignment

We consider the problem of self-localization of receivers using only TDOA information from unknown signal sources. In a communication network n receivers are located at unknown positions $\mathbf{M}_i \, (i = 1, \ldots, n)$ in $p = \{2, 3\}$-dimensional Euclidean space $\mathbb{R}^p$. Now m signals occur at arbitrary positions $\mathbf{S}_j \, (j = 1, \ldots, m)$ at unknown time points t_j. The wavefront of every signal propagates with constant signal velocity c, starting at the signals' origins $\mathbf{S}_j$ at time t_j. The signals arrive at the receivers at time points T_{ij} which can be measured. The propagation is described by the constraint equation

$$c\,(T_{ij} - t_j) = ||\mathbf{M}_i - \mathbf{S}_j|| \quad . \tag{1}$$

An equation system is formed by the equations for n receivers and m senders. For a minimum number of signals and receivers the system is unique or overdetermined. When we distribute the equations we get squared and mixed terms. According to [2] and [8] it does not seem likely to find efficient solutions to the problem in general.

Non-linear approaches can solve the problem. However, in some cases the iterative methods run into local minima from which they cannot recover, even with repeated attempts [1].

We now present our novel Cone Alignment algorithm which is based on a spring-mass simulation. We quantify the chance of running into local minima and we use our algorithm to increase the probability of solving the problem. In the following we omit the indices i, j for clarity.

From Eq. (1) we know that $T = t + \frac{1}{c}\|\mathbf{M} - \mathbf{S}\|$. The equation describes a cone in p+1-dimensional space (Fig. 1) where the signal time t is added as a dimension. The vector $(\mathbf{M}, T)$ is the apex of the cone, $(\mathbf{S}, t)$ describes a signal at position $\mathbf{S}$ at time point t. If for all receivers $\mathbf{M}_1, \ldots, \mathbf{M}_n$ and signal sources $\mathbf{S}_1, \ldots, \mathbf{S}_m$ these equations are satisfied we receive a possible solution of the given problem. Of course, this does not necessarily imply that we found the correct solution as the problem might be underconstrained.

We use an error function to describe the potential energy of springs. We define $\Phi((\mathbf{D}, t_\mathbf{D})) := c\,t - \|\mathbf{D}\|$.

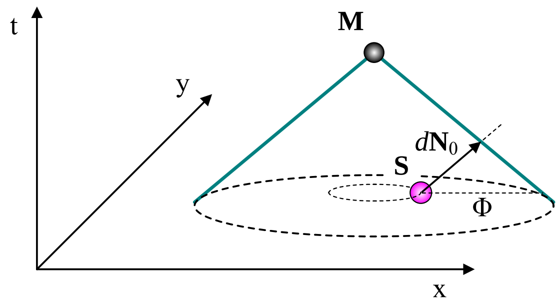

Fig. 1. Cone representation of Eq. (1) in 2D space. The signal source $\mathbf{S}$ resides offside the cone surface of receiver $\mathbf{M}$ and therefore it is not *valid* and $\Phi \neq 0$. The direction vector $\mathbf{N}_0$ intersects the cone to restore validity.

If the error function Φ gives a non-zero value, which we call an *invalid* location, one can change both the position and time $(\mathbf{S}, t)$ of the signal source and the position vector $\mathbf{M}$ of the receiver by moving it in p+1-dimensional space in order to recover a valid position. Receiver time T is fixed by definition. We define:

$$\mathbf{N} := \left(\frac{\mathbf{S} - \mathbf{M}}{||\mathbf{S} - \mathbf{M}||}, \frac{1}{c} \right) \tag{2}$$

The normalized direction vector $\mathbf{N}_0 := \frac{\mathbf{N}}{||\mathbf{N}||}$ describes the shortest path from $\mathbf{S}$ to the cone surface of $\mathbf{M}$ in respect of signal velocity c.

For the case that $t > T + \frac{1}{c}\|\mathbf{M} - \mathbf{S}\|$ and thus $\mathbf{N}_0$ does not intersect the cone, we choose $\mathbf{N}_0 := (\vec{0}, -1)$ pointing along the time axis ensuring an intersection.

By construction there is a scalar $d \in \mathbb{R}$ such that $\Phi((\mathbf{M}, T) - (\mathbf{S}, t) + d\mathbf{N}_0) = 0$. d equals the distance along $\mathbf{N}_0$ between $(\mathbf{S}, t)$ and the cone surface (Fig. 1). It can be computed by

$$d := \left(1 - \frac{\Phi((\mathbf{M}, T) - (\mathbf{S}, t) + \mathbf{N}_0)}{\Phi((\mathbf{M}, T) - (\mathbf{S}, t))} \right)^{-1} . \tag{3}$$

We calculate a force to minimize d using the spring equation $\mathbf{F} = -k\, d\, \mathbf{N}_0$ where k is a constant describing the spring stiffness. Applying $\mathbf{F}$ to every receiver particle and $-\mathbf{F}$ to the corresponding signal particle changes the locations and time points to minimize the spring extension and hereby the potential energy of the spring-mass sytem. In the case of success all relations become valid.

We manipulate the signal and receiver positions with a simulation of this spring-mass system. It is based on *particles* which are tuples $(\mathbf{x}_t, \mathbf{v}_t, m)$ representing the receivers and signals in p+1-dimensional space at discrete simulation times t. They have physical properties position $\mathbf{x}$, velocity $\mathbf{v}$ and mass m and obey Newton's law of inertia.

Since we have no anchor points we cannot directly compare our calculation results to the real positions ("ground truth"), i.e. the final translation and rotation of the signal sources and the receiver network are not determined. For an evaluation of the quality of the algorithm we use singular value decomposition (SVD) to generate a transformation to align our found positions with the real-world positions.

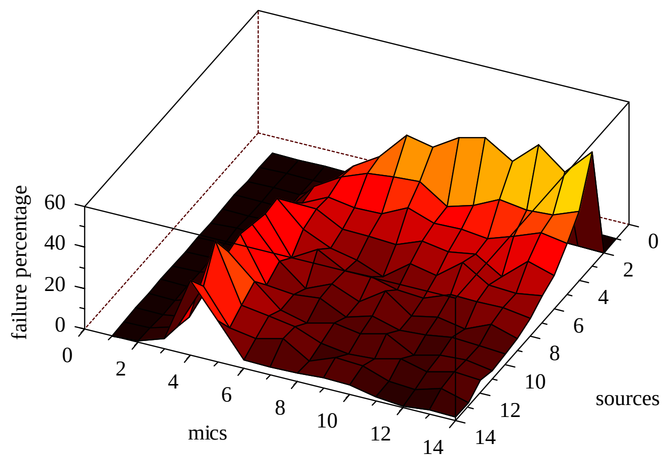

Fig. 2. Distribution of local minima in percent for two dimensions. For four receivers and for three signal sources the risk of ending in a local minimum is exceedingly high.

III. Simulation

We have implemented the algorithm in C++. Simulations were run in both the two-dimensional and the three-dimensional case. For the signal velocity we choose the speed of sound at 20 °C, which is c = 343 m/s.

For any numbers of microphones and sound sources, where $n, m \leq 14$, we created 100 random scenarios in the plane and in space with an edge length of 1000 m. Timestamps were calculated and passed to our algorithm. As an abort condition of the algorithm we chose an error threshold ϵ. If the threshold could not be reached after a maximum number of iteration steps the run was marked as not successful.

In some cases the localization algorithm failed and got stuck in a local minimum of the error function, see Fig. 2. This opposes errors in finding positions due to under-determined scenarios. Local minima occur in uniquely determined or over-determined scenarios. The failure rate converges to zero with increasing number of microphones and signals.

In a visual representation we saw that items were blocked on the wrong side of a line or a plane. We implemented an algorithm that mirrored them on the other side by way of trial. This resolved local minima in some but not in all cases.

Furthermore, we ran experiments with simulated TDOA error. Here, the jitter in timestamping the signals at the receivers is assumed to be Gaussian distributed. Errors of a standard deviation up to 200 ms were tested, which is a spatial equivalent of 70 m. With increasing TDOA error the average distance from the real positions grew linearly and the risk of local minima increased.

We have compared our algorithm to a non-linear least-squares fit using gradient descent, which is a common approach to non-linear problems. It is briefly mentioned with regard to this problem in [8]. Both algorithms, the least-squares fit and the Cone Alignment algorithm, are executed with Newton's method following, which speeds up convergence. For the least-squares fit we find the distribution of local

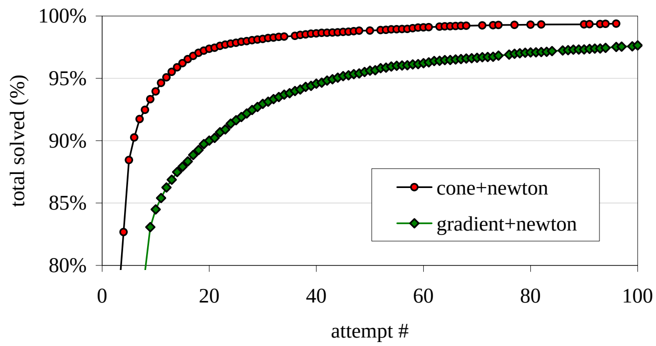

Fig. 3. In the 4M / 6S minimum case in planar space the Cone Alignment algorithm solves 99.4% of a total of 8000 random scenarios. Using gradient descent we achieve only 97.6 % after max. 100 attempts per scenario.

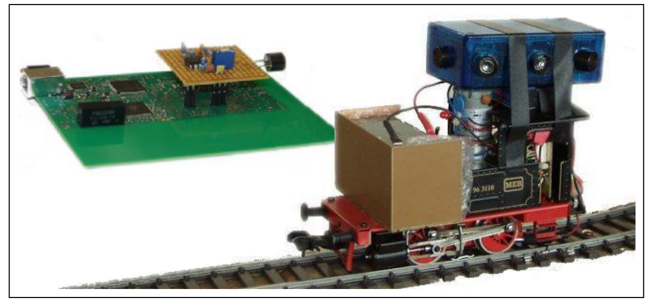

Fig. 4. *Left*: Receiver platine with ultrasound capsule. *Right*: Beacon with eight ultrasound capsules facing in all directions on top of the model train.

minima similar to the results of the Cone method. We observe regions with high risk of local minima, especially in the case of four microphones and for three signal sources.

We focus on the prominent case of four microphones in a plane, a minimum case where solutions are unique. Using Cone Alignment we achieve a failure rate which undermatches the failure rate of the gradient method by approximately 15 % for varying numbers of at least six signals.

We suppose that the gradient descent method fails to escape local minima, as it can only decrease in its error function. In contrast, the particles of the spring-mass simulation gather momentum while relaxing the spring constraints. In this way, barriers can be overcome towards a smaller minimum. As we implemented particle velocity as an imitation of physical springs we did not have to optimize a momentum parameter.

We have run repeated executions of our spring-mass simulation with randomized initial values, which increases the probability to find a solution. In the minimum case of 4 microphones and 6 signal sources in the plane we achieve a success rate of 99.4 % after 100 repeats with randomized initialization. Only 0.6 % of all cases remain stuck and unsolvable, see Fig. 3. As we can split larger scenarios into subsets of this size and merge them after solving a subset, we can solve larger scenarios in the same way. This form of repeating should also work for the other minimum cases, for 5 / 4 and for 7 / 3 microphones and sound signals, and for the three-dimensional case.

In the case of the gradient descent method and Newton's method combined we could not achieve such a high success rate. After 100 repeats still 2.4 % of all scenarios fail to be solved, which is more by a factor of four.

IV. Real-world experiments

We have tested this theoretical approach in several real-world experiments. We use a network with laptops as network nodes. Our software establishes TCP/IP-communication via wireless network (WLAN) between the laptops and it provides precise time synchronization up to an order of 0.1 ms. On every computer we record audio signals, either audible sound or ultrasound. In the case of audible sound we use the built-in microphones. Ultrasound signals are received with external receiver devices which we have built and which are connected to the laptops. From the discrete audio signals we calculate the time points of arrival using the synchronized time. They are exchanged to every participating computer and the positions are computed locally.

In a first experiment we placed eight receiver devices on a green field on our campus in an area of 30 m. With their built-in microphones they recorded audio signals produced by an assistant who walked beneath the computers while clapping two wooden bars at arbitrary locations.

The spring-mass simulation got the times of the clapping and computed the relative locations of microphones and sound signals. We achieved an average location error (Euclidean distance) of the microphones of 28 cm (σ = 14 cm). The average error of the signal positions was 39 cm (σ = 28 cm).

Now, we present a tracking system for moving targets using our algorithm. It can quickly be set up, without the need to measure the positions of the devices. This is in contrast to many commercially available tracking systems which are expensive and need to be calibrated. Of course, when the positions of at least three of the devices are specified, the relative coordinates that we obtain can be converted to absolute coordinates.

Our ultrasound tracking system consists of a sender beacon and receivers that record and process the signals from the beacon (Fig. 4). It has been assembled from off-the-shelf components and underprices most commercially available tracking systems.

The beacon creates short ultrasound pulses. With eight ultrasound capsules facing in all directions it creates an approximately isotropic signal. The beacon can be carried by a person or we attach it to a moving unit, for example a model car or a model aircraft. It is battery powered so it can be used independently from line voltage.

The receiver devices record the signal with their ultrasound microphones. The analog signal is then amplified and digitized. Over a serial connection the data is forwarded to a processing computer. Here, the data stream is searched for signal peaks, as in the case of audio signals.

In an experiment we track a moving model train. On a very simple trajectory, an oval of the dimensions 3.9 m $\times$ 1.8 m, the train travels with a velocity of about 0.5 m/s. The ultrasound beacon has been attached to the roof of the model

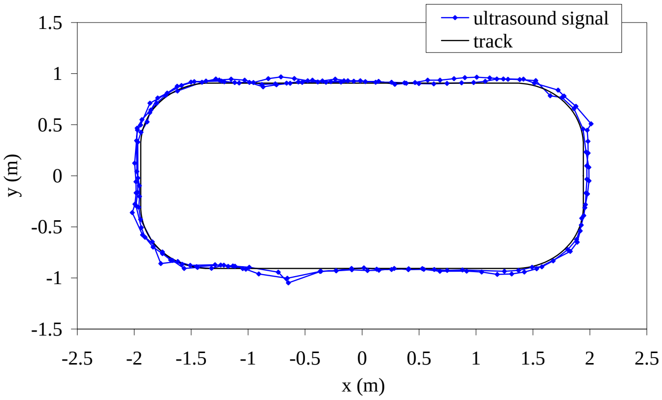

Fig. 5. Trajectory of the model train with the ultrasound beacon attached. The five receivers are located outside the track. The root mean square (RMS) error of the trajectory is 2.5 cm compared to the track.

train, see Fig. 4. Five receivers are placed roughly in an oval around the track, at a distance of 4 – 6 m. As we conduct a 2D experiment we place the receivers at the same height as the beacon. Of course, our algorithm can be used in three-dimensional settings. Then, we distribute the receiver devices in space.

Next, the ultrasound capsules are roughly oriented towards the oval track and connected to adjacent laptops. With our software running they can find each other in a Wi-Fi network and synchronize their clocks. Using a measuring tape we measure the positions of the ultrasound capsules up to a precision of 3 cm. For the dimensions of the train track we describe the geometrical shape of the track.

After approximately three rounds the spring-mass algorithm got the TDOA data as the only input. We calculate both the unknown ultrasound receiver positions and the trajectory of the train on the track. Comparing the data we find it well matching the ground truth data. However, we observe some overestimation. The receivers show an average deviation from the real positions of 44.5 cm (σ = 7.7 cm).

The overestimation is weakly pronounced for the trajectory of the model train (Fig. 5). We observe only a small overestimation which results in a root mean square (RMS) track error of 2.5 cm.

V. Conclusions

We have addressed the problem of self-localization using nothing but TDOA information with our novel Cone Alignment algorithm. The iterative spring-mass simulation solves the problem of relative localization in terms of energy minimization. Particles obeying Newton's law of inertia gather momentum while spring constraints are relaxed.

Like all iterative approaches for this problem the algorithm suffers from the risk of local minima. We have quantified the success rate of our algorithm and we have increased the probability of solving the scenario of 4 microphones and 6 signals to 99.4 %. Here, the algorithm outmatches the non-linear least squares approach, especially in the minimum case of four receivers in a plane.

In our real-world experiments we have proven the viability of our approach. We have located the receiver positions with unknown audio signals from the surroundings. Furthermore, using our algorithm we have created a quick setup reference system for vehicle tracking where precise indoor locations in the order of centimeters are provided. There is no need to measure the positions of the reference receivers. As the sole tasks we attach the ultrasound beacon to the moving vehicle and place the receivers at generally distributed, but arbitrarily chosen, positions in the room.

A. Future work

In graphical representations of the problem we have seen that we could solve the problem of local minima in some cases by flipping the particles. In this way, we might further increase the success rate of the algorithm.

We also plan to improve the practical aspects of our localization scheme. For many scenarios the assumption of discrete, distinguishable sound events is impractical. We envisage speaker tracking and locating ourselves by passing cars. This requires to calculate the TDOA by comparing audio signals using cross correlation. We expect this will extend the number of application scenarios for our technique.

Furthermore, we aim to improve the prediction of moving signals. Especially under the assumption of spatial coherence of signals we plan to apply filtering techniques like the Kalman filter or a Monte Carlo simulation. Then, we can estimate the beacon's position and interpolate in case of missing signals.

Of great interest is also the question of unsynchronized localization. This would simplify the approach immensely and would be helpful especially for unreliable network connections and for mobile networks like GSM or UMTS.

References

[1] R. Biswas and S. Thrun. A Passive Approach to Sensor Network Localization. In *Proceedings of the International Conference on Intelligent Robots and Systems, 2004. (IROS 2004). 2004 IEEE/RSJ*, volume 2, pages 1544–1549, 2004.

[2] 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.

[3] Jean-Marc Valin, François Michaud, Jean Rouat, and Dominic Létourneau. Robust Sound Source Localization Using a Microphone Array on a Mobile Robot. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, pages 1228–1233, 2003.

[4] 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.

[5] Nissanka B. Priyantha, Anit Chakraborty, and Hari 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.

[6] 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.

[7] S. Thrun. Affine Structure From Sound. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, Cambridge, MA, 2005. MIT Press.

[8] 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. ICASSP 2008.*, pages 2445–2448. IEEE, 2008.