Range sensor data fusion and position estimation for the iLoc indoor localisation system

S. Knauth
Stuttgart University of Applied Sciences
HFT Stuttgart, Faculty for
Geomatics, Computer Science and Mathematics
Schellingstr. 24, D-70174 Stuttgart, Germany
stefan.knauth@hft-stuttgart.de

C. Jost and A. Klapproth
Lucerne University of Applied Sciences
iHomeLab - the Swiss think tank and lab
for intelligent living and building automation
Technikumstrasse 21
CH-6048 Horw, Switzerland
christian.jost@hslu.ch info@iHomeLab.ch

Abstract

An ultrasound indoor localisation system typically comprises an infrastructure formed by a plurality of fixed nodes, so called beacons. By ultrasound time of flight measurements, the range between a mobile node and some of the fixed beacons at known positions is detected. Given a minimum of 3 such range measurements, the postion of the mobile node can be determined by trilateration. To achieve good accuracy and coverage, even under difficult or disturbed conditions, it is common practice to include more than 3 measurements to calculate a position estimation. The calculation result will somehow average over the reported values. Inclusion of inacurrate measurements will degrade the result. We present an algorithm for selection of the three "most believed" range measurements out of the obtained ranges. The algorithm is employed in the iLoc indoor localisation sytsem.

1. Introduction

Ultrasound localisation systems are known for a while, for example the "CRICKET" [3], "CALAMARI" [5] or "BAT" [4] systems. They provide high and reliable accuracy, achieved with moderate effort, when compared to newer approaches like ultrawideband systems. The capabilities of embedded systems have evolved considerably since the time of development opf these systems. The iLoc system [1] is a newly developed ultrasound positioning system optimized for size, cost, deployment effort and accuracy.

The iLoc indoor localisation system is an RF- and ultrasound based indoor localisation system which is used for visitor tracking in the iHomeLab laboratory: Visitors get an electronic name badge. This badge can be localized with an average accuracy of less than 10 cm deviation of its spatial position, by means of reference nodes distributed in the lab rooms. In this paper we report on the employed strategies for position estimation based on the iLoc ultrasound range sensor readings.

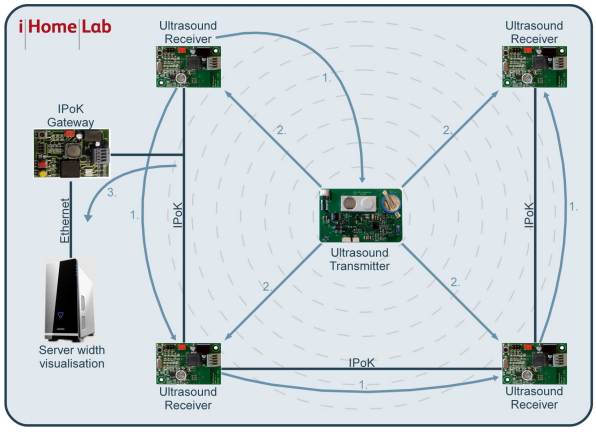

Figure 1. Setup overview. 4 reference nodes are shown. The upper left receiver sends out a synchronisation signal (arrows labelled "1") by wire ("IPoK") to the other reference nodes and by radio to the mobile node (center). The mobile node emits an ultrasound pulse (arrows "2") and the reference nodes record the reception time.

In Fig. 1 the principle of operation is shown. A dedicated reference node sends out a synchronisation signal to the other reference nodes at a rate of about 20 Hz, such that their local clocks are synchronised to about 1..2 μs. The signal is transferred via a serial multipoint connection, the so-called "IPoK" (IP-over-Klingeldraht) two-wire bus [2], which also supplies power to the reference nodes. By radio, the synchronisation signal is sent to the mobile nodes (tags) which shall be localised. The signal contains a sequence number indicating which tag, upon reception of the RF packet, shall transmit an ultrasound pulse.

This ultrasound pulse may be received by some reference nodes, which are therefore equipped with an ultrasound receiver. The reception time is recorded by the reference nodes and is transferred via the IPoK bus to an ethernet gateway. A PC collects the reception times and calculates time-of-flight data and finally evaluates the 3D spatial position of the tags.

The given ultrasound transmitters, receivers and electronics allow a detection of pulses over a distance of about 16 meters, under optimal conditions. This corresponds to a maximum time-of-flight of about 50 μs, enabling a repetition rate of about 20 Hz. Using 20 mobile nodes, each node may send one pulse per second and thus be localized with this rate. Principally, 3 range measurements at different positions allow the determination of the tag position. Practically, the density of reference nodes has been chosen be much higher such that the distance to the most far node does not exceed about 5 meters, and that each point is covered by more than 3 nodes. This is due to the fact that the ultrasound signal needs a line-of-sight for propagation and gets shaded for example by the body of the wearer of the tag or of other visitors in the same room. About 70 Nodes are deployed in the two rooms, as shown in figure 2. A more detailled description of the system is given in [1].

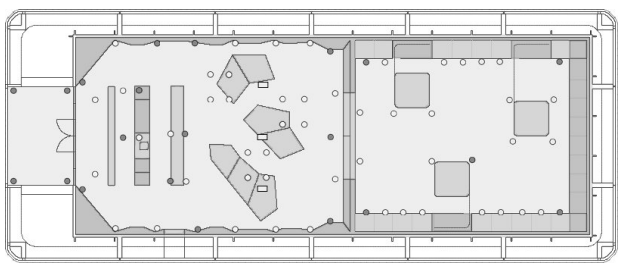

Figure 2. Node deployment in the environment (Filled- and clear circles).

2. Position calculation

For a mobile node, time-of-flight data is converted to distances by using the appropriate sound velocity and a system-induced constant time offset, as described in Eq. (1):.

$$s_i = c_{sound} \cdot (t_i - t_{ofs}) \tag{1}$$

Here c_{sound} is the sound velocity, t_i is the measured ultrasound time-of-flight, t_{ofs} is the mentioned constant time offset to correct for delays in the generation- and detection chain, and s_i is the calculated distance between the mobile node and reference node i at position R_i,

Typically for each node a set of 10..20 values is obtained, meaning that about 10..20 reference nodes were able to pick up an ultrasound signal. The position is determined by trilateration, i.e. by solving the equation system

$$|R_i - P| = s_i \tag{2}$$

where P is the unknown position of the mobile node, R_i is the position of a reference node i, and s_i is the ultrasound-detected range between the mobile node P and the reference node i.

Since only 3 values are needed to solve the trilateration equation system (2), an often-used possibility to calculate the position for the over-estimated equation system is to perform a least-square calculation or fit of the values. This method did not lead to the desired accuracy since it averages over all measurements and includes "bad" measurements with the same weight as "correct" range measurements.

We investigated the nature of errors which occur in the measurements. A linear averaging method or fit performs well on gaussian error distributions in the measurement values. But the obtained range values showed a different behaviour: Repeated measurements for a given tag position showed that about 95 percent of the "correct" values fell into an interval of 2 cm. Errors did occur most likely from non-line-of-sight measurements, and from reflections, and led to results being some 10 cm or even more away from the correct result Also acoustic noise lead to wrong measurements, which manifested in reporting of arbitrary range values. Therefore we chose a different approach to combine the range measurements to a result (for an overview see fig 3): We obtain n range measurements with distance values s_i, i = 1..n, each from a different node and its position R_i. For a trilateration calculation, 3 ranges are sufficient. So we calculate positions $P_{ijk}(s_i, s_j, s_k, R_i, R_j, R_k)$, for all possible permutations (i, j, k) of the n range values. The number of such tuples is (n over 3).

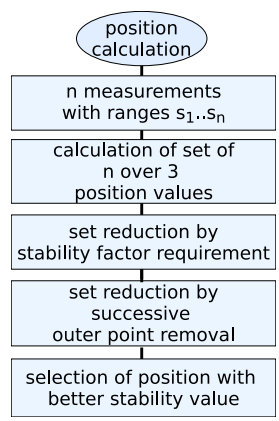

Figure 3. Range sensor data processing steps

The first step to reduce the number of results is to perform a boundary check, We select only those results which make sense i.e. only those which mark a position which can be accessed by a visitor. This also removes the mirrored results which are given by the trilateration calculation.

3. Conclusions

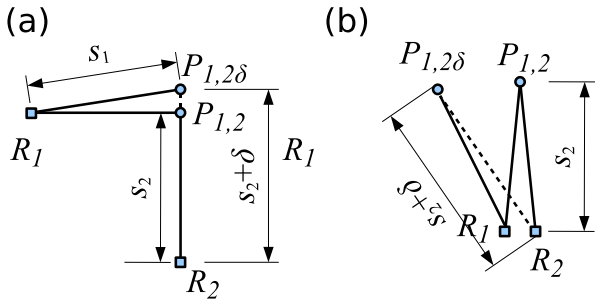

Figure 4. Geometric considerations: The variation of s_2 by δ leads to displacements of the obtained position. In (a) the "good" case is displayed, where the displacement is about the value of δ. Under worse geometric conditions (b), a range modification may lead to multiplied position displacement. Such geometries lead to a degradation of the position result accuracy

After this step, for each obtained position result P_{ijk}, a so-called stability factor f is calculated. Easy spoken, the stability factor expresses how strong a measurement error will influence the calculated position. The stability factor f is the divergence of the position P_{ijk} when varying $s(i), s(j)$, and $s(k)$. In order to evaluate f, besides calculating P_{ijk}, also three positions each with a short range variation δ (δ was typically set to 1 cm) are calculated:

$$P_{(i\delta)jk} = P(s_i + \delta, s_j, s_k) \tag{3}$$

$$P_{i(j\delta)k} = P(s_i, s_j + \delta, s_k) \tag{4}$$

$$P_{ij(k\delta)} = P(s_i, s_j, s_k + \delta) \tag{5}$$

These positions are positions of the mobile node, if the corresponding measurement s is varied by δ. The total displacement vector F when subsequently varying all 3 measurements is given by

$$F = \begin{pmatrix} |P_{(i\delta)jk} - P_{ijk}| \\ |P_{i(j\delta)k} - P_{ijk}| \\ |P_{ij(k\delta)} - P_{ijk}| \end{pmatrix} \tag{6}$$

The normalized length f of this vector F

$$f = \frac{1}{\sqrt{3}\delta}(|F|) \tag{7}$$

is our "stability factor". f is a measure for the error multiplication given by the geometry of the 3 selected reference nodes R_i, R_j, R_k with respect to the mobile node position P. Under optimal geometric conditions, i.e. when the reference nodes somewhat enclose the mobile node, a modification of an observed range s_i by δ will result in a position derivation of the order of δ. This will lead to a value of f close to 1. The value for f increases for less ideal geometric constellations (see Fig. 4). Only positions, for which the stability factor does not exceed a defined threshold (typically 4), are taken into account for the further calculation. By this selection process, the set of obtained permutations is already considerably reduced.

It turned out that it is still not desirable to take an averaged value of the remaining positions as the position estimation result. As already mentioned, the range measurements, if not disturbed, are highly accurate, and 3 accurate measurements would fulfil our requirement to have a position accuracy below 10 cm. So, instead of averaging or mean-squares procedures, we use a selection algorithm to select the three range values which are most likely to give the best position estimation.

The algorithm works as follows: The remaining position vectors are averaged (center of gravity), and the position which is most far away from this average position is removed from the set of calculated positions. This step, including the averaging, is repeated as long as there are more than two positions in the set. Finally, the position with the lower stability factor is chosen as the real position of the tag.

This is a quite complicated and compute-intensive algorithm for selecting the 3 "best" range values out of a set of ranges, but it operates quite stable also under problematic conditions, i. e. a multiplicity of range values where only a few of them carries useful information. A main difference to other approaches is that here no averaging or otherwise merging of the measurements is performed but the most likely value triplet s_i, s_j, s_k is selected out of the measurements. The approach has been chosen because we observed that a reported ranges is either very accurate or is more or less random. The obtained positioning error obtained by the method is well below 10 cm (standard derivation and absolute position).

4. Results

In Fig. 5 results for the described algorithm are shown. The figure is taken from a set of about 1500 subsequent measurement cycles, with at most 8 of a set of 9 reference nodes reporting TOF values. For comparison, also position errors obtained from averaging over all observed TOF value triplets, only applying the boundary condition, are shown. The rightmost values include all measurements lying outside of the graphs X-Axis. In the experiment, the mobile node stayed at a fixed position. The recording time was about 5 minutes. During the recording of the measurement values, the sound propagation was disturbed by noise, people walking around, and placing of ultrasound shields and reflectors. The high overall accuracy of the reported measurement values has been achieved by careful adjustment of the sound velocity and position data of the reference nodes. The advantage of the selection algorithm is clearly visible, especially when regarding the outliers which would produce unwanted result hopping.

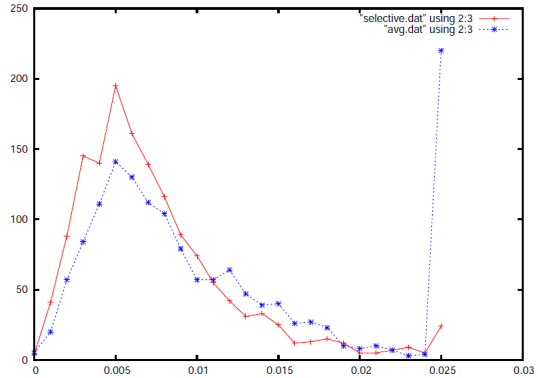

Figure 5. Error Distribution of the localisation algorithm: Solid line (red, +) indicates positions calculated with the presented selection algorithm, dashed line (blue, *) indicates positions which were obtained by an averaging algorithm.

5. Conclusions

We presented and discussed an algorithm for position estimation in indoor localisation systems. The iLoc positioning system, for which the algorithm has been developed, exhibits some particular properties with respect to the obeyed range sensor data:

- Due to the nature of ultrasound ranging, the error distribution of the range measurement values is not gaussian distributed but more or less is either quite accurate, delayed by reflections, or random by noise.
- The system collects a large number of data, each mobile node may be detected by 25+ sensors at a time.
- The update rate is about 1 Hz, and the system is intended to considerably update its obtained position estimation within this time frame.
- The position calculation process is performed with a PC connected to the system, rather than on the resource-constrained nodes itself.

It turned out that the above peculiarities and requirements were met best with the presented approach:

- For each possible combination of 3 reference nodes, calculate a position estimate by trilateration. Of course do only include nodes which did report a range value for the particular node.
- From the set of all position estimates, remove those who lie outside a defined boundary, given for example by the room geometry.
- From the remaining set, remove those position estimates which do not meet the stability requirement (f shall be below a certain threshold).
- From the remaining set, remove the "most out" point. Repeat this process until only 2 position estimates are left (There can be no "outer" point in a set of two points).
- From the remaining two position estimates, select the one with the better (lower) stability factor as the systems result for the position estimation.

The described algorithm performs well also under quite disturbed circumstances, i. e. people walking around and making noise. It requires a considerable number of reference nodes to ensure that the ultrasound emitter at the mobile node has line-of-sight to at least 3 nodes.

For dynamic localisation, an unwanted effect to occur is the effect of "hopping" positions in visualisations. Typically, such "hopping" induced by, for example, pick up of random noise or dynamic propagation conditions, introduced for example by moving people, can only be removed by some sort of averaging, which would lead to a delayed response of the system. The presented approach delivers fast updates without "hopping", at an observed absolute ranging accuracy of about 10 cm, in the building.

Further work to be carried out on the algorithm itself includes some measures to initially reduce the number of ranges to a smaller number such that the calculation hardware performance requirements can be reduced and the position calculation may take place in the gateway processor. Also there will be included some post-calculation position estimators, to more accurately support also mobile nodes which are in a deep sleep mode and only rarely transmit position update pulses.

Acknowledgements

This work is funded by Hasler Stiftung, Berne, Switzerland, in the frame of the "ManCom" programme.

References

[1] S. Knauth, C. Jost, and A. Klapproth. iloc: a localisation system for visitor tracking and guidance. In *Proceedings of the Embedded World Conference 2009*, Nuremberg, Germany, Mar. 2009.

[2] S. Knauth, R. Kistler, C. Jost, and A. Klapproth. Sarbau - an ip-fieldbus based building automation network. In *Proc. 11th IEEE Intl. Conference on Emerging Technologies and Factory Automation (ETFA '08)*, Hamburg, Germany, Oct. 2008.

[3] A. Smith, H. Balakrishnan, M. Goraczko, and N. B. Priyantha. Tracking Moving Devices with the Cricket Location System. In *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.

[4] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, Oct. 1997.

[5] K. Whitehouse, F. Jiang, C. Karlof, A. Woo, and D. Culler. Sensor field localization: A deployment and empirical analysis. *UC Berkeley Technical Report UCB//CSD-04-1349*, Apr. 2004.