Removing Systematic Error in Node Localisation Using Scalable Data Fusion

Albert Krohn[1], Mike Hazas[2], and Michael Beigl[3]

[1] Telecooperation Office, University of Karlsruhe, Germany, krohn@teco.edu
[2] Lancaster University, United Kingdom
[3] Distributed and Ubiquitous Computing Group, University of Braunschweig, Germany, beigl@ibr.cs.tu-bs.de

Abstract. Methods for node localisation in sensor networks usually rely upon the measurement of received strength, time-of-arrival, and/or angle-of-arrival of an incoming signal. In this paper, we propose a method for achieving higher accuracy by combining redundant measurements taken by different nodes. This method is aimed at compensating for the systematic errors which are dependent on the specific nodes used, as well as their spatial configuration. Utilising a technique for data fusion on the physical layer, the time complexity of the method *is constant and independent of the number of participating nodes*. Thus, adding more nodes generally increases accuracy but does not require additional time to report measurement results. Our data analysis and simulation models are based on extensive experiments with real ultrasound positioning hardware. The simulations show that the ninety-fifth percentile positioning error can be improved by a factor of three for a network of fifty nodes.

1 Introduction

In sensor networks, knowledge of physical node topology is important for packet routing, power control, location annotation of gathered sensor data, and to satisfy demands for specific application areas, such as mobile computing. A variety of solutions for node localisation have been proposed, concentrating on both sensing technologies and algorithms. These techniques have traditionally traded off performance attributes (such as accuracy and location update rate) with resource requirements (such as sensor node cost, power consumption, computational complexity, and communication overhead).

There has been a significant amount of effort invested in improving the accuracy of the location result while keeping the algorithm resource requirements to a reasonable level. More recently, some researchers have become concerned with characterising how the location result is affected by the error of the raw sensor measurement (such as radio signal strength, range, or angle-of-arrival) [1, 2]. In this paper, we introduce a new method wherein the systematic measurement error typically found in localisation systems can be dramatically reduced without incurring additional computational complexity or communication overhead. In fact, as the number of nodes within sensing range increases, the communication overhead of the method stays constant, and the overall accuracy gain improves.

1.1 Related Work

A number of systems have been developed for tracking people and objects [3]. These have employed a variety of sensing media, including infrared, ultrasound, RF signal strength, ultra-wideband radio, computer vision, and physical contact. However, many of these systems have been designed for static deployment using dedicated infrastructure.

Other systems have been specifically designed for the processing, communication, power constraints, and dynamic deployments imposed upon sensor networks. Primarily, these have relied upon creating ranging estimates using the received signal strength indication (RSSI) provided by a sensor node's radio transceiver [4, 5], or creating finer-grained (centimetre-scale) ranging estimates by measuring the time-of-flight of an acoustic signal [6–8]. Recently, it has been shown that competitive ranging accuracies are possible by measuring the phase offset between two slightly different radio carrier frequencies [9]. Some authors have also proposed measuring the angle-of-arrival of a signal, to be used to compute node orientation [7], or as a quantity to be used in triangulation to produce location results [10] (as opposed to *trilateration* which is applied using range measurements).

There has been a large focus on developing and comparing algorithms which take sensed ranges and compute location results with varying degrees of distributed operation [11]. Some of the algorithms do not require anchor nodes whose location is previously surveyed or known—these algorithms produce a set of *relative* coordinates which describe the spatial layout of the nodes with respect to one another, but not to any global reference point [12, 13]. Relative location solutions are often sufficient for the purposes of network routing or analysing how sensor readings vary across the network (e.g. temperature and humidity in environmental monitoring). Researchers have also analysed the effects of ranging noise on the location results [1, 14]. One paper in particular showed for multi-hop networks how ranging and angle-of-arrival measurement error creates accuracy bounds for the location result [2].

1.2 Motivation

Much in the spirit of the latter papers, our approach is informed by analysis of measurement error. However, rather than try to produce increasingly robust (and possibly resource-intensive) algorithms, we chose to address the measurement error, specifically the *systematic error* that is common in many sensor node localisation systems. This paper's contribution is twofold. First, it presents a method to reduce systematic errors usable for many sensor measurement settings. Second, a lightweight framework is proposed which requires minimal hardware, software and network resources and is able to support such error reduction even in scenarios where there are a large number of sensor devices.

Systematic errors appear as constant and repeatable error within measured quantities. This systematic error component can arise due to sensor decalibration (e.g. misalignment of an acoustic sensor) or environmental interference which is

constant for that spatial distribution of nodes (e.g. a metal object is present near one of the nodes, affecting its radio RSSI measurements).

Because a large component of the systematic error is hardware-dependent or spatially-dependent and thus different for each node in the system, we propose fusing the measurements taken by different nodes together in order to "average out" this measurement offset. This can be done when the nodes are capable of independently computing estimates of the same spatial quantities. For example, if all nodes in a system are capable of independently estimating the range between two given nodes, then all of these range measurements could be combined to produce an inter-node range whose systematic error component is much reduced.

As more nodes join the system and are capable of measuring the same physical quantities, then the communication overhead required for each node to report its measurement can become prohibitively expensive, especially if measurement updates need to be provided at a high rate. To avoid this, our method accomplishes the averaging by using a physical layer data fusion technique called *synchronous distributed jam signalling* (SDJS). This allows all nodes to report their measurements for a given quantity simultaneously. Thus, the communication overhead for our method is constant and independent of the number of nodes in the system. The error models, proof-of-concept, and simulations are based on extensive data gathered using an ultrasonic relative positioning system [15].

1.3 Target application scenarios

Because our method relies upon averaging of measurements from different nodes to remove large systematic errors, it naturally grows more effective when there are greater numbers of nodes within measurement range of one another. High node densities are often envisioned for application scenarios such as environmental habitat monitoring and battlefield surveillance, and it is typical to consider high node densities in analyses of node localisation algorithms and error [2, 11, 13].

Another application area which envisions high node densities is that of mobile and ubiquitous computing. In such scenarios, it is assumed that everyday objects such as mobile phones, furniture, or even coffee cups are embedded with microcontrollers, sensors, and wireless communication. These *augmented objects* are thus able to sense, compute, communicate, and work together to aid people in their everyday tasks. Contextual data is often needed for this, and one of the most important types of context data is location information. Thus, accurate localisation is an important goal for the dense settings of sensors envisioned for mobile and ubiquitous computing.

2 System Operation

This section discusses the attributes and operation of the type of positioning sensor network that we consider in this paper. We begin with a notation for

relative positioning systems and then discuss how nodes can locally estimate the same physical quantity. Throughout this paper, we focus on implementation and analyses of 2D localisation systems, such as the hardware platform in figure 1, but similar methods could be applied for 3D systems.

2.1 Relative Positions

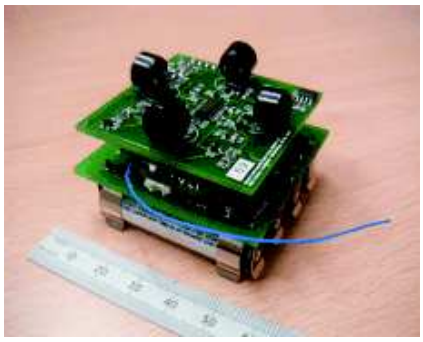

Fig. 1. A node used for ultrasonic relative positioning

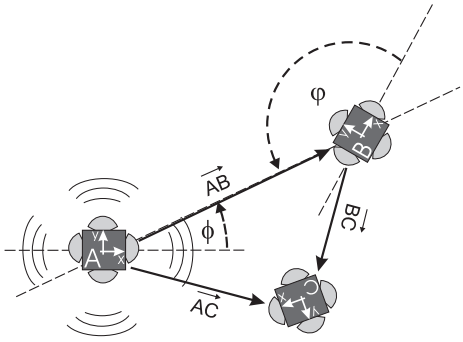

Fig. 2. Relative positioning of nodes

Figure 2 depicts three sensor nodes which take measurements that can be used to estimate the nodes' relative positions. Each node in figure 2, denoted with a capital letter, defines its own local coordinate system, referenced to its sensor hardware. In figure 2 this is marked with the local (x,y) - coordinate systems that belong to each of the nodes. Relative positioning systems normally do not use anchor nodes or global coordinate systems, so we have to clarify a relative position (which is a vector) by always including the according reference coordinate system in the notation.

For example, $\overrightarrow{AB}_A$ expresses the position of B (meaning the vector A to B) in the view of a local coordinate system of A. If we determine relative positions via multiple hops, we can just add the vectors together as long as we assure to use the same reference coordinate system for all vectors. Different views of positions (and therefore vectors) can be translated by simply transforming the according coordinate systems. The position of C in figure 2 in the view of A ($=\overrightarrow{AC}_A$) can be combined through step by step summation of a vector chain from A to C:

$$\begin{aligned} \overrightarrow{AC}_A &= \overrightarrow{AB}_A + \overrightarrow{BC}_A \\ \overrightarrow{BC}_A &= \begin{pmatrix} \cos\rho & -\sin\rho \\ \sin\rho & +\cos\rho \end{pmatrix} \overrightarrow{BC}_B \\ \rho &= \phi + (180^o - \varphi) \end{aligned} \tag{1}$$

2.2 Creating Local Estimates of the Same Physical Quantity

As described above, relative positions and orientations can be computed by applying trigonometry to measured ranges, angles-of-arrival, and angles-of-emission. The method proposed in this paper relies upon multiple nodes being able to independently create local estimates of a given physical quantity so that they can then be fused to improve accuracy. We now give two examples of how multiple distributed estimates of the same physical quantities would be created:

Ranging For example, many observing nodes might independently estimate the range between two beaconing nodes A and B. To do this, an observing node (for example C) would need to minimally measure its range to A and B (d_{AC} and d_{BC}), as well as the angles-of-arrival of the ranging signals from both A and B (ϕ_{AC} and ϕ_{BC}). Applying the Law of Cosines, the observing node C could then estimate d_{AB} with only these locally-measured quantities. Using experimental data, Section 4.2 demonstrates that ranges locally estimated in this way by three observing nodes can be fused together to yield an improved range result.

Positioning Alternatively, many nodes might each estimate the relative position of B with respect to A. As explained below in Section 5, an observing node is able to independently estimate this physical quantity by locally measuring the ranges to A and B, the angles-of-arrival of the ranging signals, as well as the angles-of-emission of the ranging signals. By way of simulations, Section 6 shows how an improved estimate of the relative locations of A and B can be created by fusing estimates from many observing nodes.

3 A Model for Measurement Errors

In general for measurement systems, the process of measuring can be understood as an operation based on the ground truth. For example, suppose we take a sensor observation $\tilde{r}$. It is based on the actual physical quantity being observed, the "ground truth" r. Incorporated in the observation is also an error due to the system status (e.g. spatial configuration of the nodes, or environmental conditions) which is described by a vector we want to call $\mathbf{t}$. Additionally, there is an independent part n influencing the measurement, which appears as random noise on the measurements. Thus, the overall sensor observation can be described as the sum

$$\tilde{r} = r + f(\mathbf{t}, r) + n. \tag{2}$$

This model divides the influences on our measurement into the two types of *systematic* ($f(\mathbf{t}, r)$) and *statistical* (n) errors. *Statistical errors* come from random sources of noise, such as the thermal noise present in sensors and their supporting circuitry, or (particularly relevant to ultrasonic sensing) acoustic noise created by uncorrelated physical events in the environment. Statistical errors

cannot be avoided but can be reduced by techniques such as cooling of electronic parts or aggressive filtering of the sensor signal. *Systematic errors* in a measurement system occur for a variety of reasons, including component ageing, wear-and-tear, or changing environmental conditions that affect the measurement process in a constant and repeatable way. To handle the systematic errors, the traditional approach is to either *model* the errors or *calibrate* the affected parts of the system. Modelling systematic error is e.g. done in ultrasonic ranging systems by taking temperature and humidity measurements to gain an accurate estimate of the speed of sound in air.[4] Methods for calibration on the other hand vary in difficulty and effectiveness, but a thorough calibration can be a tedious task requiring high precision. Moreover, re-calibration is often required during the life-time of a product, and this can also be very difficult and expensive to manage.

By analysing measurements from our ultrasonic localisation system, we discovered that a large component of the systematic error is dependent upon the *spatial configuration* of the nodes. Spatially-dependent systematic error has been shown to exist for other sensing modalities as well; examples include infrared light intensity [16] and radio RSSI [17]. We would expect that our methods can be applied to sensor network localisation systems based on infrared light, ultrasound, audible sound, or radio signals.

4 Removing Error in Measurements

In this paper, we focus on the low-labour technique of averaging; device-specific calibration and modelling can always be carried out and applied in addition to averaging. According to our error model (2), multiple measurements taken in the same system conditions will help to reduce the statistical error n (assuming the statistical error is non-biased). By contrast, the systematic part $f(\mathbf{t}, r)$ can not be averaged out in this way, since it represents the constant and repeatable error component.

4.1 Removing systematic error

One of the core ideas presented in this paper is a natural extension to averaging in location system. It is obvious that averaging of multiple measurements helps against statistical errors such as thermal noise in analogue components. In the same way, we want to use averaging on the systematic error. We can achieve this by *varying* the system state during multiple measurements. In our application of positioning systems, this would e.g. imply to vary the system states like temperature, orientation etc. In other words, if we are able to measure the same physical value under *varying* system states, we will not get a repeatable, but also varying systematic error, and thus the systematic error loses its *systematic* nature.

[4] If the estimate of the speed of sound is too fast or too slow, then all measurements will suffer from a systematic over– or under-ranging.

4.2 Analysis of experimental measurements

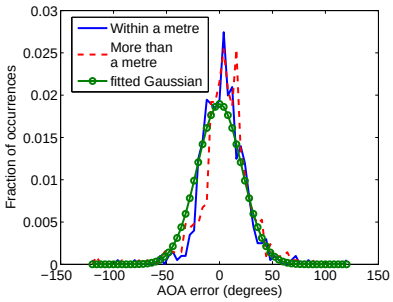

Fig. 3. Mean error for angles-of-arrival detected by five nodes in many different spatial configurations. The error characteristic in this plot is for measurements where the incoming ultrasound pulses were sensed on two or more of the node's transducers; this condition was true for over seventy percent of the successful measurements taken.

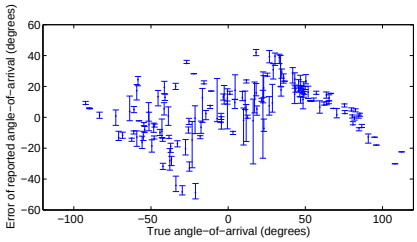

Fig. 4. Angle-of-arrival measurements taken by a single node in different spatial configurations. The error bars indicate the tenth and ninetieth percentile levels of the observed errors at a given angle. The angle-of-arrival error offsets for a given system state tend to be tightly constrained, and thus systematic.

In figure 3, we see a plot of the errors of the angle-of-arrival of ultrasound pulses that were measured experimentally; our experimental analysis is based on data gathered using five nodes placed in fifty spatial configurations, with favourable line-of-sight conditions [15]. The distribution is very similar to the Gaussian indicated in the same figure. But for a given node spatial arrangement, we are more interested in the specific error distribution rather than the total distribution for all nodes/configurations. Figure 4 shows the angle-of-arrival measurement errors reported by one node over many experimental settings. Each error bar shows the distribution of angle-of-arrival errors measured from a particular node in a particular location. While the error distribution for different nodes in different spatial configurations varies, the spread of errors for a given node in a given spatial configuration tends to be small. Thus, if measurements are taken by a node in the same system state ($\mathbf{t}$ is constant), the error tends not to vary much; it is repeatable, or systematic. This holds true for readings taken by all five nodes used in the experiments; ninety percent of the ninetieth percentile angle errors are within 16° of the median error.

For a particular node's repeated measurements, Figure 4 shows that the angle-of-arrival error tends to have a systematic offset. However, in order to motivate the use of averaging to compensate for systematic errors, it is necessary to show that the systematic errors of the nodes in a given spatial configuration are independent. Using the experimental data collected from five nodes in fifty spatial configurations, the range between two of the nodes was estimated by

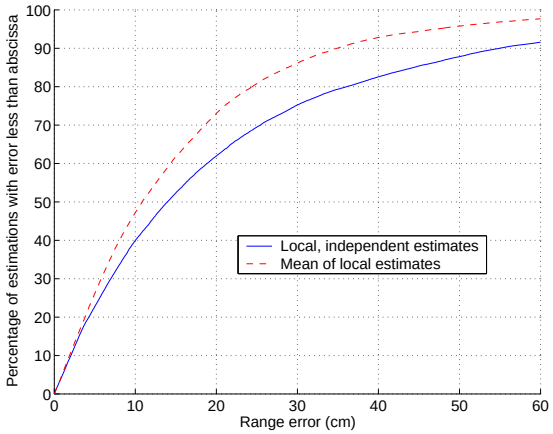

Fig. 5. Removal of systematic error using data collected from three nodes to estimate the range between two others

using data collected locally by the other three. The three estimates were then averaged together to create a fused estimate of the range between the two nodes. Figure 5 illustrates the result of fusing three locally-obtained range estimates. The ninetieth percentile accuracy improves by about 15 cm. The graph demonstrates that the device– and location-dependent systematic errors in the system are not significantly correlated with one another; averaging the locally-estimated quantities, despite their individual systematic offsets, yields an improved result.

5 Multiple, Simultaneous Measurements

In our analysis of the measurements, we found that a large component of the systematic error depends on the particular signal angles-of-arrival at each node, which is in turn dictated by the relative locations of the nodes. Instead of carrying out a device-specific calibration, we propose choosing a physical quantity which is to be measured by many nodes to overcome the systematic nature. Because each node has its own systematic but specific error offset (dependent on the node's location, orientation, and particular hardware), the random spatial configuration of the nodes can be exploited to yield several estimates of the same physical quantity which have different systematic error offsets. These independent estimates can then be averaged together for an improved result.

Localisation nodes which are capable of measuring (1) the time-of-flight (i.e. a range), (2) the angle-of-emission, and (3) the angle of arrival of the signal can, using simple trigonometry, locally estimate the relative positions of any other two beaconing nodes (figure 2). Note additionally that in many sensor systems, any signals used for localisation (such as infrared, ultrasound, or ultra-wideband radio) are of a *broadcast* nature.

Thus, multiple measurements can be accomplished simultaneously by different nodes: whenever a node emits a localisation signal, a number N of other nodes can take a measurement. Provided these N nodes can all estimate the same quantity from their measurements (such as the range between two nodes A and B, or their relative location coordinates), then the scaling of taking measurements can be improved from $O(N)$ to $O(1)$.

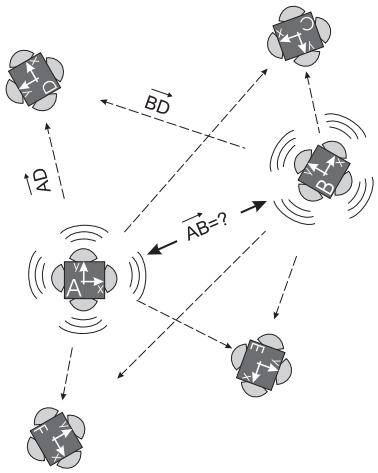

Fig. 6. Performing multiple measurements in a relative positioning system based on ultrasound

As shown in figure 6, if A and B both successively emit ranging signals, the other nodes present can individually and simultaneously compute results under different system states which can then be combined for improved positioning of A and B. For example, when A transmits, another node D can measure A's relative position using the angle-of-arrival and range to A. When B transmits, D similarly estimates its relative location. Since D has additionally logged the angle-of-emission of the signal from A, D can perform a rotation (1) of its locally-referenced coordinate system to express the location of B relative to A:

$$\overrightarrow{AB}_A = (\overrightarrow{AD}_D - \overrightarrow{BD}_D)_A. \tag{3}$$

Other nodes in the environment (C, E, and F) can do the same, simultaneously arriving at the relative position of B with respect to the local coordinate system of A ($\overrightarrow{AB}_A$). Thus after an ultrasonic emission from both A and B, there are $N = 4$ measurements of $\overrightarrow{AB}_A$. These estimates individually held by the nodes C, D, E, and F now need to be combined, ideally in a way which scales favourably as N increases.

5.1 Data Fusion with Multi-SDJS

We now hold N measurements $\tilde{r}_i$ of the same vector from A to B, which we want to collect for a data process such as averaging. With multi-SDJS, we are capable of collecting these vectors in a very efficient way. The important property is that the time required to perform SDJS is independent of the number of participants. Increasing the number of participants in multi-SDJS allows more observations to be incorporated into the fused data estimate, without impacting the amount of time required for the fusion process.

For the ease of understanding, we explain the protocol directly with our application of collecting the measured vectors from A to B; we called the measurement $\tilde{r}_i$. The SDJS process we utilise operates with *scalar* values in an a-priori *known interval*. Therefore, the $\overrightarrow{AB}_A$ measurements are expressed as polar coordinates $(\tilde{r}_i,\ \tilde{\varphi}_i)^T$, and multi-SDJS is performed for each component.[5] For the SDJS process, we assume that the nodes are within radio range of one another, since they have been able to perform the localisation measurements.

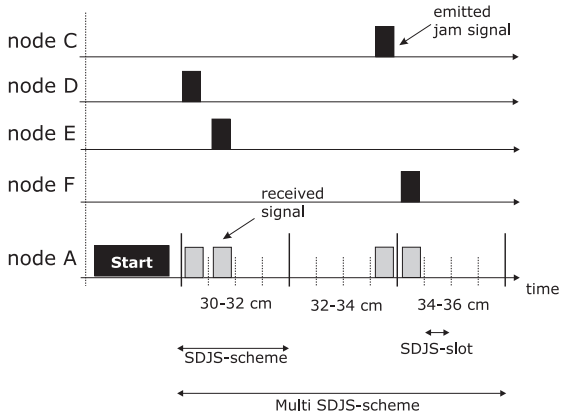

Fig. 7. The multi-SDJS process for collecting data using signalling on the physical layer of the nodes' radio

Figure 7 shows how the collecting of data works. Node A initiates the multi-SDJS scheme by broadcasting a start packet, which contains information about the upcoming SDJS communication. In this case, it specifies that three SDJS schemes will follow, each scheme containing four SDJS slots. The schemes correspond to the ranges 30–32 cm, 32–34 cm, and 34–36 cm. To lower the necessary

[5] Note that averaging can only work using a polar coordinate representation $(r, \varphi)^T$, as the distribution of the measurement (as that shown in figure 10) is not bias-free in $(x, y)^T$. This is because the raw ultrasonic measurements themselves produce relative distances and angles, and not (x, y) coordinate results.

slots and schemes, the ranges are based on the distance estimation that the initiating node A has.

After the start packet, each node selects the SDJS scheme which corresponds to their measurement. For example, if node D has a measurement of 31.2 cm, it would select the first SDJS scheme. Then, during the time of the first SDJS scheme, node D (and also node E) *randomly* select one SDJS-slot and emit a jam signal. Using these emissions, from D and E signal to A that they have a measurement that lies in the interval from 30 to 32 cm. Node A will receives the signals and can estimate the number of nodes based on counted signals and collisions. Node A does not simply *count* the number of jam signals but has to include possible collisions into its estimation. For a discussion on the estimation theory, the reader is referred elsewhere [18].

6 Simulation

Our simulation is based on an indoor scenario of a 5m by 5m room with up to 200 nodes.[6] The simulation error models for the range and angle-of-arrival accuracy were derived from our experiments, in which over half a million measurements were taken using five ultrasonic positioning nodes placed in fifty different spatial configurations [15]. The following models were used for the simulation:

Angles-of-arrival: a Gaussian systematic error with $\mu = 0°$ and $\sigma = 14°$; a Gaussian statistical error with $\mu = 0°$ and $\sigma = 1.4°$

Angles-of-emission: a uniform distribution in the interval $[-45° \ldots + 45°]$

Ranges: a right-sided half-Gaussian systematic error with $\mu = 0$ cm and $\sigma = 3.5$ cm; a Gaussian statistical error with $\mu = 0$ mm and $\sigma = 4.9$ mm

Simulation model for the SDJS process: The SDJS process influences the averaging process of the distributed measurements. It quantises the numbers (we used a resolution of 2 cm and 2˚) and then the number estimation of the SDJS process can also introduce errors due to an inadequate range of possible values or due to using too few slots per scheme. For our simulation, we assumed that the SDJS process would be suitably parameterised such that the errors introduced with SDJS are negligible.

In the simulations, we compare "direct measurements" with the measuring process results of our framework. The direct measurements assume a measurement that is based on signals and data exchange only between the endpoint nodes. We always take the best possible solution here and let B transmit and A receive as the model for the angle-of-arrival (Gaussian, $\sigma = 14^o$) is much more narrow than the transmit model (uniform $\pm 45^o$).

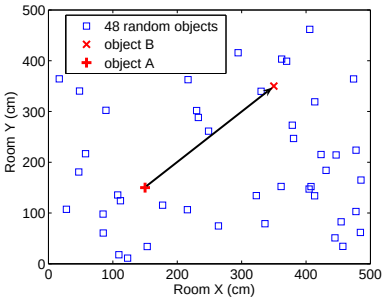

Fig. 8. A simulation run with 50 nodes

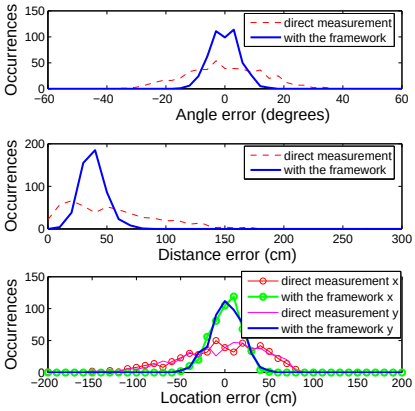

Fig. 9. The improvements using the framework for the measurements

6.1 Acquisition of the relative position between two nodes

Figure 8 shows a simulated room with 50 nodes. We measured the relative location from node A $(150, 150)^T$ to B $(350, 350)^T$. We simulated 1000 runs of this scenario, varying the position of the 48 "helper" nodes. We looked at the error for the relative position of node B, meaning the error of $\overrightarrow{AB}_A$. In figure 9, we compare the results of direct measurements and our framework for the case of 50 nodes. The three plots show the relative occurrences of error values for the angle between A and B, the absolute position error of B and the error for the Cartesian coordinates (x,y) of the room. All three plots compare the distribution of the results of the direct measurements with those of our framework.

For the direct measurements, the maximum error in positive x direction is given when node B is (wrongly) positioned on the same y-position like A in the room, then the x-position $\overrightarrow{AB}_A$ of B is 432 cm plus some noise, resulting in a x-position error of 82 cm plus noise. We can see this limit around 100 cm in the x-position errors of figure 9. The distribution of angles are symmetric to zero, which we expected. For all values, we see how the framework improves the overall results. The distribution are much more narrow. The error on the absolute location error gives the distribution of the distance of the results from the target position of B. We also see here, how the mean was shifted towards zero.

Figure 10 gives another illustration of one of our simulation runs. It shows how the distribution of the direct measurements are placed on a segment of a circle, as the error in the angle φ is much higher than on the radius r. After the averaging process of the 48 helper nodes, the distribution is much more

[6] The simulation was done using OMNet++. See `http://www.omnetpp.org/`.

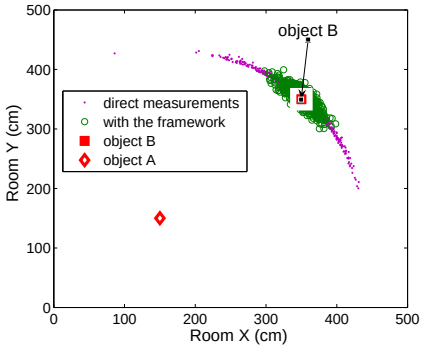

Fig. 10. 1000 simulations run with 50 nodes

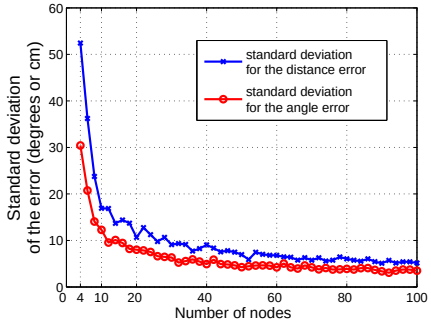

Fig. 11. Gradual improvements of accuracy with increasing number of nodes N

concentrated around the target node B. We simulated 1000 runs, each run with a new topology, only the (x,y)-position of node A and B were held constant. Their orientation varied as well.

In table 1, we find the statistical characteristics of the distributions for the normal direct measurement process and the averaged values through our framework. The improvements are exceedingly visible in the decreasing of the standard deviation of the distributions. With the averaging, we achieve much more narrow distributions. The (x,y) coordinates describe the relative position error of node B in the view of node A. Φ denotes the angle under which B was estimated from the view of A. The row labelled "Euclidian" shows the absolute Euclidian position error in the location estimation of B from its real position at (350,350).

(μ/σ)	direct meas.	with framework (N=48)
Φ (degrees)	(-0.3/ 13.4)	(-0.1/ 4.9)
Euclidian (cm)	(53.6/ 38.8)	(38.5/ 10.8)
x (cm)	(-4.7/ 47.2)	(1.4/ 17.2)
y (cm)	(-2.6/ 46.0)	(1.6/ 17.7)

Table 1. Statistical characteristics for N=0, 48

Table 2, demonstrates how adding large numbers of nodes improves the results of the system. Numbers are based on a simulation of the same situation, but with 198 "helper" nodes. One can see, how the standard deviations σ of the resulting statistics decrease and the distribution becomes even more narrow. The tabulated ninety-fifth percentiles give the error values at which the corresponding cumulative distribution reached 0.95. In other words, 95% of all measurements will have an error equal or smaller than that value.

(σ/95%)	direct	N=48	N=198
x (cm)	(47.2/90.5)	(17.2/ 32.4)	(9.9/18.6)
y (cm)	(46.0/86.5)	(17.7/33.5)	(10.0/20.6)
Φ (degrees)	(13.4/35.6)	(4.9/9.8)	(2.8/5.5)
Euclidian (cm)	(38.8/125)	(10.8/56.3)	(5.2/44.1)

Table 2. statistical characteristics for N=0, 48, 198

As a next step, we now want to look at the dependency of our achieved accuracy and the number of nodes. Table 2 already gives a hint on the clear convergence with increasing number of nodes N. We again simulated the same situation but now with a variable number of nodes from $N = [1..200]$. For each number of nodes, we simulated 100 random topologies in the 5 by 5 m room but kept nodes A and B at (x,y)-position (150,150) and (350,350) with random orientation. We numerically extracted mean and standard deviations from the results of the framework and plotted them against the number of nodes in Figure 11. We see how the standard deviation of the measurements decreases the more nodes that are involved. The gradient is especially high for the low number of nodes ($N < 10$). That means, that even with a comparably small number of nodes, we can already achieve great improvements in accuracy. But as the framework process time is independent of the number of participants, one would always include all available measurements from all nodes in the area.

7 Conclusion

This paper has described and evaluated a method for reducing device-dependent and spatially-dependent systematic error in localisation systems. Using this method, many nodes simultaneously gather measurements to create estimates of the same physical quantity, such as the relative location of one node with respect to another. In so doing, the systematic errors can be treated as statistical or random, since they depend on the nodes' independent locations and orientations. Averaging can then be applied to these independent estimates to improve the overall result.

Using the SDJS protocol, the time required for sensor measurements and data fusion is constant and *independent* of the number of nodes involved; the accuracy increases with a higher number of nodes involved, as long as this number stays in reasonable bounds. Some quantitative measures have been previously published [18]. As an example calculation, we want to parameterise a typical SDJS process for our application. With an initial estimate of the location of B relative to A, we may need for the next estimate a range of ± 50 cm for the SDJS process with a resolution of 1 cm, resulting in one hundred SDJS schemes. Assuming a slot time with 18 µs (based on physical parameter from IEEE 802.11a) and fifty slots per value, we get a total time of $100 \cdot 50 \cdot 18\mu s = 90.0$ ms for the averaging process. Using fifty slots per value would allow the network to scale to upwards of 300 nodes, and the SDJS process would still return accurate fused estimates.

With these properties, our method is ideally suited for low-cost hardware that is densely distributed in the environment. With the method, we address both statistical and even severe systematic errors of the hardware through the fusion of measurements taken from different nodes in different locations and orientations.

Using a large number of measurements captured from five ultrasonic localisation nodes, we demonstrated that the systematic errors of locally-estimated ranges between any two nodes are independent and thus the overall error can be vastly reduced by fusing the local estimates; even with just three nodes performing the process, the ninetieth percentile range error can be improved by about 20 cm. With the error models derived from these real measurements, we also showed in simulations that the accuracy can be greatly improved using even more nodes. For this, we simulated a 5×5 m room with both 50 and 200 nodes and quantified the accuracy improvement. For example, the standard deviation of the determined relative angle was reduced from $\sigma_{\Phi,2} = 13.4$ for the direct measurement over $\sigma_{\Phi,50} = 4.9$ with a setting of 50 nodes to only $\sigma_{\Phi,200} = 2.8$ when using 200 nodes in the room.

Acknowledgements

The work presented in this paper was undertaken as part of a collaborative project, "Relate: Relative Positioning of Mobile Objects in Ad Hoc Networks." It is funded by the European Commission's FP6 IST Programme, project number 013790.

References

1. Kamin Whitehouse, Chris Karlof, Alec Woo, Fred Jiang, and David Culler. The effects of ranging noise on multihop localization: An empirical study. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 73–80, Los Angeles, USA, April 2005.
2. Andreas Savvides, Wendy L. Garber, Randolph L. Moses, and Mani B. Srivastava. An analysis of error inducing parameters in multihop sensor node localization. *IEEE Transactions on Mobile Computing*, 4(6):567–577, November 2005.
3. Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, August 2001.
4. Jeffrey Hightower, Roy Want, and Gaetano Borriello. SpotON: An indoor 3d location sensing technology based on RF signal strength. UW CSE 00-02-02, University of Washington, Department of Computer Science and Engineering, Seattle, WA, February 2000.
5. Kiran Yedavalli, Bhaskar Krishnamachari, Sharmila Ravula, and Bhaskar Srinivasan. Ecolocation: A sequence based technique for RF localization in wireless sensor networks. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 285–292, Los Angeles, USA, April 2005.

6. Andreas Savvides, Chih-Chieh Han, and Mani B. Srivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the Seventh International Conference on Mobile Computing and Networking (MobiCom)*, pages 166–179, Rome, Italy, July 2001.
7. Nissanka B. Priyantha, Allen K. L. Miu, Hari Balakrishnan, and Seth Teller. The Cricket Compass for context-aware mobile applications. In *Proceedings of the Seventh International Conference on Mobile Computing and Networking (MobiCom)*, Rome, Italy, July 2001.
8. Masateru Minami, Yasuhiro Fukuju, Kazuki Hirasawa, Shigeaki Yokoyama, Moriyuki Mizumachi, Hiroyuki Morikawa, and Tomonori Aoyama. DOLPHIN: a practical approach for implementing a fully distributed indoor ultrasonic positioning system. In *Proceedings of the Sixth International Conference on Ubiquitous Computing (UbiComp)*, pages 347–365, Nottingham, UK, September 2004. Springer.
9. Miklós Maróti, Branislav Kusý, György Balogh, Péter Völgyesi, András Nádas, Károly Molnár, Sebestyén Dóra, and Ákos Lédeczi. Radio interferometric geolocation. In *Proceedings of the Third International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 1–12, San Diego, USA, November 2005.
10. Dragoş Niculescu and Badri Nath. Ad hoc positioning system (APS) using AOA. In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 1734–1743, San Francisco, USA, March 2003.
11. Koen Langendoen and Niels Reijers. Distributed localization in wireless sensor networks: A quantitative comparison. *Computer Networks*, 43:499–518, August 2003.
12. Srdjan Čapkun, Maher Hamdi, and Jean-Pierre Hubaux. GPS-free positioning in mobile ad-hoc networks. *Cluster Computing Journal*, 5(2):157–167, April 2002.
13. Yi Shang and Wheeler Ruml. Improved MDS-based localization. In *Proceedings of the Twenty-Third Conference of the IEEE Communications Society (INFOCOM)*, Hong Kong, March 2004.
14. David Moore, John Leonard, Daniela Rus, and Seth Teller. Robust distributed network localization with noisy range measurements. In *Proceedings of the Second International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 50–61, Baltimore, USA, November 2004.
15. Mike Hazas, Christian Kray, Hans Gellersen, Henoc Agbota, Gerd Kortuem, and Albert Krohn. A relative positioning system for co-located mobile devices. In *Proceedings of the Third International Conference on Mobile Systems, Applications, and Services (MobiSys)*, Seattle, USA, June 6-8 2005.
16. Albert Krohn, Michael Beigl, Mike Hazas, Hans Gellersen, and Albrecht Schmidt. Using fine-grained infrared positioning to support the surface-based activities of mobile users. In *Proceedings of the Fifth International Workshop on Smart Appliances and Wearable Computing (IWSAWC)*, Columbus, USA, 2005.
17. John Krumm and Eric Horvitz. LOCADIO: inferring motion and location from Wi-Fi signal strengths. In *Proceedings of the First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (Mobiquitous)*, pages 4–13, Boston, USA, August 2004.
18. Albert Krohn, Tobias Zimmer, Michael Beigl, and Christian Decker. Collaborative sensing in a retail store using synchronous distributed jam signalling. In *Proceedings of the Third International Conference on Pervasive Computing*, Munich, Germany, 2005.