Robust Device-Free Localisation with Few Anchors

Francesco Potortì
CNR-ISTI
Via G. Moruzzi, 1
Pisa, 56124 ITALY
Potorti@isti.cnr.it

Pietro Cassarà
CNR-ISTI
Via G. Moruzzi, 1
Pisa, 56124 ITALY
pietro.cassara@isti.cnr.it

Filippo Palumbo
CNR-ISTI
Via G. Moruzzi, 1
Pisa, 56124 ITALY
filippo.palumbo@isti.cnr.it

UbiComp '18 Adjunct, Sept 8-12, 2018, Singapore, Singapore.
ACM XXX-X-XXXX-XXXX-X/XX/XX...$15.00.

Abstract
Radio-Frequency based device-free localisation systems are able to pinpoint people's location in a given area without their cooperation. They work by analysing the perturbations that the presence of a person causes on the communications exchanged by a high number of radio devices installed around the area. Typical numbers are some tens of devices for an area the size of a single-family house, with an accuracy around one meter and a high sensitivity to even the slightest movement. The literature about device-free localisation systems typically concentrates on improving sensitivity, accuracy and discriminating power, without worrying too much about the number of involved radio devices. In this paper, for the first time, we demonstrate that device-free localisation can work with reduced performance with as low as four anchors in an environment composed by two wide rooms. Being able to use so few anchors opens the possibility of several use cases where installing tens of devices is not desirable.

Author Keywords
Device-free Localisation, Performance Evaluation

ACM Classification Keywords
H.3.4 [Systems and Software]: Performance evaluation.

Today, and more and more tomorrow, we expect that a significant number of small IoT devices communicating through RF will be present in most indoor environments. Device-free localisation can be obtained "parasitically", without installing any dedicated devices, by appropriate software configuration of generic deployed devices, a scenario that opens an enormous range of possible applications.

Introduction

One pillar of human behavioural studies is identifying the user position. In past decades, a wide range of sensing technologies, mainly wearable ones, has been used to gather this information. In order to obtain reliable models of a user's behaviour, it is critical for the sensing technology to be onobtrusive. Natural candidates are Radio-Frequency (RF) Device-Free Localisation (DFL) systems, which have been investigated for around ten years [6]. They exhibit the interesting capability of identifying the position of a target person in a given area without any cooperation from the target, i.e., without the person wearing any passive or active device to ease localisation. An RF-DFL system is able to identify the position of a person with an error of less than 1 m in optimal conditions [2]. Onobtrusive indoor localisation is at the base of many systems that offer insightful information on our everyday lives, daily patterns and behaviour. Specifically, it is a key feature in scenarios where wearing a continuously monitoring device is not acceptable [4].

Capabilities of RF-DFL are similar to that of camera-based localisation systems, with some important differences. A camera-based system is highly accurate in positioning and can detect many more characteristics of the target than an RF-DFL system, like identity and activities. While these are powerful capabilities, they can be a disadvantage in contexts where privacy is a concern. An RF-DFL systems is limited to detecting the presence of a small number of persons and tracking them individually if they do not come too close each other; on the other hand, it has the powerful ability to work in situations of low or no visibility, and even through walls, as long as its radios can keep communicating with each other. RF-DFL systems are also less accurate at detecting moving targets with respect to static ones, and have no ability to recognise objects. Installing and maintaining a camera-based system can be complex and expensive because cameras are not very cheap and require a relatively high-bandwidth communications medium to send their video stream to an external computer. RF-DFL systems, on the other hand, require very cheap devices with relatively low battery consumption, but each needs to be precisely located at installation time. Most notably, many devices, in the order of some tens, are needed to obtain the high accuracy typical of these systems. The high number of required fixed devices (anchors) needed for an RF-DFL system to work is the main motivation of this paper. Generally speaking, the literature about RF-DFL systems has been concerned with improving system accuracy and robustness [5]. On the other hand, this paper is concerned with lowering the number of anchors as much as possible while still obtaining useful performance.

In the following we show that, in our experimental setup, we manage to trade accuracy for number of anchors down to as low as four anchors. This work is an improvement of the one presented in [7], where we managed to go down to eight anchors using three different RF-DFL methods. As far as we know, this is the first time that a solid experimental demonstration is given that using RF-DFL is possible with only four anchors in a 75 m^2 environment. This result opens the possibility of using RF-DFL to many applications that are precluded from using the usual RF-DFL systems that need tens of installed RF devices.

Working principles

RF-DFL systems we are concerned with are based on exchanging Received Signal Strength (RSS) measurements among a set of fixed devices installed in the environment (anchors). The way the RSS information is used to infer the position of the target user is a matter of

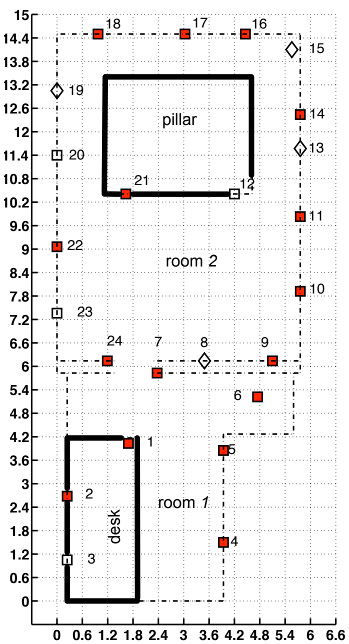

Figure 1: Localisation area. The pillar has wooden walls, the dotted lines are the walls of the area, composed by two communicating rooms.

research, and current solutions are still in their youth. The literature shows two main approaches, respectively based on classification and on Radio Tomographic Imaging (RTI). In both cases, anchors installed in the environment exchange data packets; each anchor measures the RSS from other anchors and broadcast these measurements in the packets they send themselves. By collecting data from the packets exchanged by anchors, it is possible to produce a real-time picture of the RSS measured on all links between any two anchors. This RSS picture is perturbed when the environment changes, and a person is more than enough to create a distinctly measurable perturbation. Classification-based approaches identify some features of the RSS picture and use machine-learning algorithms [1, 8]. They require a training phase, and consequently are expensive to deploy and update, but can be easily tuned to the specific needs of the task at hand and give high-accuracy results. Radio tomographic imaging was proposed by Wilson and Patwari in [2] and later refined in various ways. The idea is to use the tomography concept: let's define a very simple function which predicts the perturbation that a target presence induces in the RSS picture, given the target position; by inverting this function one can find the target position given the RSS picture perturbation. This method works surprisingly well and is quite flexible.

Our experimentations are based on the classification method described in [8], denoted as CLAS, plus two different RTI-based methods. The former is based on absolute changes in the RSS picture (called shadow-based in [2]), denoted as SRTI, the latter is based on the variance of the RSS picture [9] and denoted in the following as VRTI. We demonstrate for the first time that, for all three RF-DFL systems, accuracy performance gracefully degrades with diminishing number of anchors down to only four anchors. We comment on the results produced by the three methods, using the same set of experimental data and the same procedure adopted in [7].

Experimental procedure

The measurement campaign took place in two rooms for a total area of 75 m^2 (see figure 1). The path followed by the target user was fixed and easily reproducible, because it was indicated by marks stuck on the floor, one per step. The target moved along the path shown in figure 2 at a regular speed of one step per second, with the help of a metronome, and stayed still for 5 s on the points indicated with a circle. Anchors are IRIS Motes from Crossbow, based on the 2.4 GHz RF transceiver AT86RF230, compliant with IEEE 802.15.4 standard, most of which were hung on the room walls. They run a TinyOS application inspired by the SpinQueue algorithm [1]. Three measurements were done, using 16, 20 and 24 anchors. In figure 1 the 16 anchors are shown with red squares. The four black squares were added for the 20 anchors measurement, and finally the four diamonds were added for the 24 anchors measurement. Packets were sent by all anchors in a round-robin fashion on a different IEEE 802.15.4 communication channel each round, for a total of four different channels. After four rounds, the whole procedure started again. The protocol included error checking and auto-correction features based on timeouts which managed packet losses and lost synchronisation in the token-passing procedure, obtaining a rate of ∼60 packets per second. We made a total of three *measurements*, with 24, 20 and 16 anchors. As depicted in figure 2, each measurement consisted of a target person moving through 85 *positions* for 3 *repetitions*, for a total of about 250 positions on each of which the *positioning error* was computed in each of the three measurements.

[1] http://span.ece.utah.edu/spin

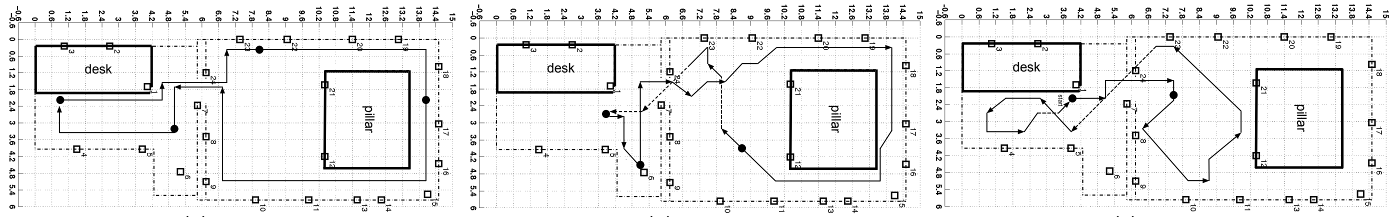

(a) Path 1 (b) Path 2 (c) Path 3

Figure 2: Paths walked by the target user

Using few anchors in different configurations
The purpose of this work is to understand whether it is possible to obtain useful positioning results even with few anchors. More specifically, it is important to answer two questions:

- is there a minimum number of anchors below which performance drops dramatically (threshold effect) or else performance degrades gracefully with diminishing number of anchors?
- is performance strongly dependent on anchor placement in the environment, so that it is paramount to have a solid method to identify good spots for installing anchors, or having some simple rules of thumb is enough to get robust performance, for example just avoid placing anchors too near each other?

We start from the experimental data gathered in [7] and compute the performance of the three methods ignoring data from and to a subset of anchors, which is equivalent to using a smaller number of anchors than available in the experimental traces. Note that this procedure is conservative from a performance evaluation point of view. In reality, reducing the number of anchors increases the token speed, and consequently the number of samples and the method accuracy, while with our procedure the round-robin speed of the token remains the same as in the cases with more anchors. Ignoring some anchors can be done in many ways, for example if we start from a 24-anchors network and we want to evaluate the performance of a 10-anchors subnetwork, we are faced with about 2 million possible 10-anchors subnet configurations. In the following, we exploit this diversity to reach two goals: (i) obtaining a reliable estimate of the *expected performance* of a 10-anchor subnet, something which is not easy to do by simply installing 10 anchors in different ways; (ii) obtaining a reliable estimate of the *performance variability* when installing 10 anchors following some given installation rule. The installation rule is key to this procedure. We chose something that does not require any location-specific computations or specialised training: just distribute the anchors so that they are more or less at a uniform distance each other. In practical terms, the procedure is the following. Let be N the total number of installed anchors. In our case we made three measurements, with $N = 24$, $N = 20$, $N = 16$. If we want to measure the performance with, say, 19 anchors, we can discard 5 anchors from the initial

We remove anchors so that remaining ones are more or less uniformly distributed in the environment, which means no specific installation training or measurement is required. We show that anchor placement is not critical under this assumption.

Configuration size	Measurements		
	24	20	16
24	1		
23	24		
22	276		
21	2000		
20	10146	1	
19	37944	20	
18	107440	190	
17	233088	1120	
16	387855	4525	1
15	490776	13104	16
14	463320	27690	120
13	315888	42680	544
12	147578	47190	1628
11	43344	36100	3312
10	7812	17906	4560
9	4336	5140	4096
8	3321	1155	2214
7	696	1020	624
6	564	320	256
5	216	204	144
4	114	95	76

Table 1: Number of possible anchor configurations.

24. There are 42504 ways to discard 5 anchors from the initial 24, but many of these are of little significance, for example those that consider 5 consecutive anchors. We choose stretches of consecutive removed anchors so that the stretch lengths lie in a small range. We choose a range of 2, meaning that the lengths of the longest and shortest stretch of consecutive removed anchors have a difference not greater than 2. Table 1 lists the number of possible configurations for each number of anchors starting from the three measurements done with 24, 20 and 16 anchors. When more than 1000 anchor configurations are possible, we randomly choose 1000 of them with uniform probability to reduce computation time.

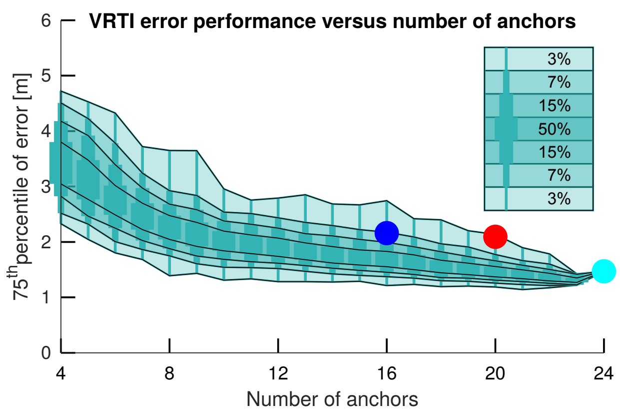

Figure 3: Error performance (75th percentile of error) for 15890 different anchor configurations. A variable number of anchors is ignored starting from the 24-anchors measurement. The red and blue dots represent the 20- and 16-anchors measurements.

As an initial example, let's use the VRTI algorithm in the 24-anchors measurement. We use the 75th percentile of error as a performance measure. In figure 3, we depict the performance of a total of 15890 different anchor configurations chosen from all the possible configurations of 24 anchors with some of them disabled, as detailed in table 1. The figure shows different shaded regions separated from percentiles lines of 0, 3, 10, 25, 75, 90, 97, 100, thus summarising the performance of all the considered configurations, where the performance measure is the 75th percentile of error, which we consider the most significant and robust [3]. Figure 3 is interesting in a number of ways. First of all, consider the inner area, the darkest one which starts with the vertex at 24 anchors and grows progressively wider with diminishing number of anchors. This is the area where 50% of cases fall. This means that, when removing anchors from the whole 24-anchors configuration, in half of cases the 75th percentile of error is inside this area. Notice how the area gently slopes towards greater errors with diminishing number of anchors: this is a very strong indication that the VRTI method behaves smoothly when reducing the number of anchors, meaning that with this method it is possible to consider trade offs between localisation accuracy and number of installed devices, down to 4 anchors. This is a significant result, considering that it was demonstrated only down to 8 anchors in [7]. Secondly, the results are consistent within the experiment: looking at the lightest areas on top, one sees that significant deviations from the most common performance measure are found in only 3% or 10% of cases. This means that in the vast majority of cases, the way the anchors are placed in the environment is not that important, assuming that one does not install devices in bunches. Again, this result is a significant extension of what is reported in [7].

Figure 4 summarises the performance of all the configurations stemming from the three algorithms. The areas are delimited by the 10th and 90th percentiles of the performance measure; in other words, in 80% of the

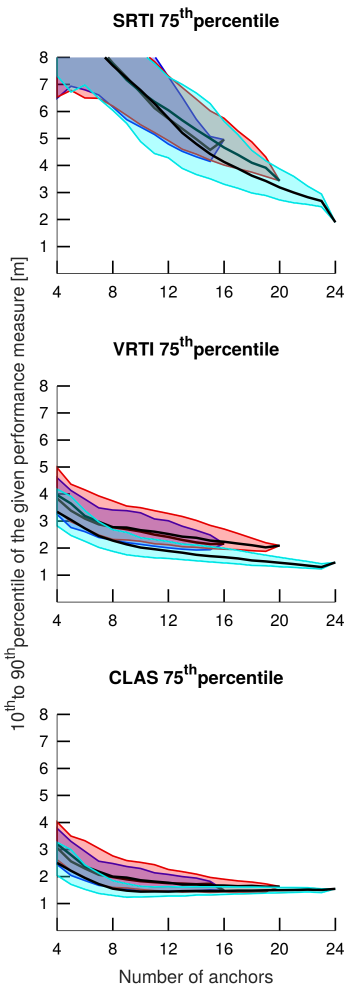

Figure 4: 75th percentile of error for the three algorithms. A variable number of anchors is ignored starting from the 24-, 20- and 16-anchors measurements.

anchor configurations, the performance measure (75th percentile) falls inside the coloured area. Even more interesting is the fact that the graceful degradation of performance with diminishing number of anchors that we had observed in figure 3 can be observed in all situations.

This brings us to two important results: all three algorithms can 1) be used with smaller numbers of anchors than normally used in the literature, at the price of a gradual loss of accuracy, and 2) all three algorithms are relatively insensitive to the positioning of anchors in the environment.

Conclusion and research challenges

Using a reliable measurement procedure, we prove that methods for RF-DFL can be used with a lower number of anchors than what is commonly described in the literature. In a two-room area, performance of RF-DFL methods degrades gracefully down to only 4 anchors. We also show that anchor placement is not critical.

In future IoT scenarios, an RF-DFL system could be "parasitically" built on top of existing devices, by exploiting the occasional communications between them. Realising this scenario requires several steps: DFL working with few anchors; DFL working with intermittent or even occasional communications; a common protocol that the devices can use to exchange Received Signal Strength (RSS) information in a secure way. This paper shows that the first step is possible; the next research step is to analyse the case when anchors do not transmit continuously.

References

[1] C. Morelli, M. Nicolini, V. Rampa, and U. Spagnolini. Hidden Markov Models for Radio Localization in Mixed LOS/NLOS Conditions. *IEEE Transaction on Signal Processing 5*, 4 (April 2007), 1525–1542.

[2] J. Wilson and N. Patwari. Radio Tomographic Imaging with Wireless Networks. *IEEE Transaction on Mobile Computing 9*, 5 (May 2010), 621–632.

[3] Barsocchi, P., Chessa, S., Furfari, F., and Potortì, F. Evaluating AAL solutions through competitive benchmarking: the localization competition. *IEEE Pervasive Computing Magazine 12*, 4 (Oct.–Dec. 2013), 72–79.

[4] Barsocchi, P., Cimino, M. G., Ferro, E., Lazzeri, A., Palumbo, F., and Vaglini, G. Monitoring elderly behavior via indoor position-based stigmergy. *Pervasive and Mobile Computing 23* (2015), 26–42.

[5] Kaltiokallio, O., Bocca, M., and Patwari, N. A fade level-based spatial model for radio tomographic imaging. *IEEE Transactions on Mobile Computing 13*, 6 (2014), 1159–1172.

[6] Pirzada, N., Nayan, M. Y., Hassan, F. S. M. F., and Khan, M. A. Device-free localization technique for indoor detection and tracking of human body: A survey. *Procedia - Social and Behavioral Sciences 129* (2014), 422 – 429. 2nd International Conference on Innovation, Management and Technology Research.

[7] Potortì, F., Cassarà, P., and Barsocchi, P. Device-free indoor localisation with small numbers of anchors. *IET Wireless Sensor Systems* (2018). In print.

[8] Viani, F., Martinelli, M., Ioriatti, L., Benedetti, M., and Massa, A. Passive real-time localization through wireless sensor networks. In *Proc. IEEE Intl. Conf. IGARSS* (Cape Town, RSA, July 2009), 718–721.

[9] Wilson, J., and Patwari, N. See-through walls: Motion tracking using variance-based radio tomography networks. *Mobile Computing, IEEE Transactions on 10*, 5 (May 2011), 612–621.