# Lost in Space
# Or Positioning in Sensor Networks


Michael O'Dell
University of Zurich
Switzerland
odellm@access.unizh.ch

Regina O'Dell
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
bregina@tik.ee.ethz.ch

Mirjam Wattenhofer
Department of Computer Science
ETH Zurich, Switzerland
mirjam.wattenhofer@inf.ethz.ch

Roger Wattenhofer
Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland
wattenhofer@tik.ee.ethz.ch



## ABSTRACT

We discuss the success chances of a real-world implementation of a positioning system for a wireless sensor network. While much research has been done in the area of node localization, the method of choice to verify the results has been in theory or by simulation. To realize the visions of future sensor networks, we take very basic sensor nodes (prototyping the smart dust idea) and implement a series of experiments to determine their suitability for use in positioning algorithms. Specifically, we look at the accuracy and effectiveness of direct distance measurements based on different signal strengths. We found that even though links are stable and symmetric over time, positioning in the real world is yet in its infant stage and that current theoretical models of sensor networks do not apply well in unspecialized hardware.


## 1. INTRODUCTION

The future of wireless sensor networks is often envisioned as a large collection of small, if not to say tiny, low-power devices. The challenges that come with this paradigm are twofold. On the one hand, engineers are trying to build hardware that conforms to this "smart dust" ideal. On the other hand, researchers need to construct models and algorithms that are suited to the very stringent requirements that these devices impose: low energy consumption, low storage and computational capabilities, and basic radio transmission hardware. With current state-of-the-art technology, such wireless sensor networks exhibit unstable links whose changes are often unpredictable.

One of the primary needs for a sensor node is to know its position so that the sensor data can be coupled to a location. Additionally, network coordinates allow for efficient routing known as geometric routing algorithms [9]. Then the problem known as *localization* (or positioning) is, given that some specially equipped nodes know their positions, for all the nodes in the network to determine their location based on the positions of the *seeds* (a.k.a. anchors) and possibly the network connectivity structure.

In this paper we focus on determining the potential of using *minimal hardware requirements* for the task of positioning in a wireless sensor network. The question is how well can a node localize itself if we only have small, low-power, low-storage devices at our disposal. The answer will allow engineers of such a network to determine whether to invest in more specialized and thus bulkier and more expensive hardware, or if the position error is tolerable in their applications.

There has been some work on investigating the physical characteristics of real-world sensor networks [5, 16]. Our findings corroborate their conclusions (as in [5]) in that the link characteristics are far from the theoretical models in use, such as the unit disk graph [4, 11] or the quasi unit disk graph [10]. However, these works are either too general in nature in that they survey the detailed link stability but not its effect on positioning ([5] looks at flooding, [16] at routing). Or they only consider two nodes at a time, whereas our findings in Section 3 indicate that there are big differences between the experimental results if only two nodes are involved versus an entire network.

There has been a significant amount of research done on the theoretical side of positioning and virtual coordinates algorithms, see [2, 3, 7, 11, 12, 14], to name a few. The underlying assumption in all of these is that the network can be represented as a static graph, usually even a unit disk graph. Physical wireless links, in contrast, are prone to instabilities. So even if the nodes are not moving, the neighborhood of a node might change rather dramatically and, more importantly, usually unpredictably, over a short period of time. This can be due to the sometimes drastic effects of intereference, scattering, or dampening of signals. If sensor networks are to be deployed in uncontrolled environments, then these effects simply cannot be ignored. We will further discuss these phenomena in Section 3 on the basis of the measurements we have taken.

Numerous of the above papers also have particular hardware requirements which assume fairly accurate measurements on the part of the nodes. Examples include measuring the time of flight (or time difference of arrival (TDoA) as in [8, 13]), the angle of arrival (AoA) [12], the received signal strength (RSS) [1]. Measuring the time of flight in reasonably-sized ad hoc networks presupposes very accurate hardware which can detect differences in the nano second range: If we assume approximately the speed of light, the time it takes for a signal to cross the distance of a few meters amounts to some tens of nano seconds. In this paper, we have also opted to use the signal strength measured in terms of packet loss at different powers as a first indication of the distance between two nodes. However, instead of requiring a particular hardware component dedicated to the precise power measurement of incoming signals, we have even less stringent hardware requirements. This is described in more detail in Section 2.

Another line of research has been the development of hardware suited to the specific purpose of sensor localization. Papers in this area include Cricket [13], RADAR [1], or also [4, 15]. Some of these work only indoors (Cricket, RADAR), some only outdoors (GPS). Cricket has a separate ultrasonic component and RADAR was tested on a larger scale with more available computing power (laptops). With current or foreseeable technology, a node cannot support fairly sophisticated positioning hardware *in addition to* the sensor and actuators that carry out the intended purpose, all at the smallest scale. Thus, we wanted to investigate the potential of very limited hardware for positioning in sensor networks. We also do not impose any indoor/outdoor restrictions.

Our contributions are on one hand results on the stability, symmetry and distance relationship of the wireless radio links, and on the other hand we have observed that the gap in the measurements between two nodes in a lab and the nodes in a network is significant enough to render any localization attempts useless at this point. Put in another way, the distance-to-power correlation is strongly and unpredictably environment dependant.

## 2. HARDWARE PLATFORM

The hardware we have used for our experiments is the ESB/2 platform from the scatterweb project, now its own company [6]. The nodes are built from standard components, consisting of a chip with a 32kHz CPU, 2kB of RAM, and a low power consumption radio transceiver, along with numerous sensors and actuators such as infrared, temperature, vibration, microphone, beep, and LED.

In the current version available to us, the nodes can adjust their transmission power $x$ from the application (for $1 \leq x \leq 100$ percent), but the transceiver is not able to directly measure the received power, only whether the signal is above a threshold. The way that the power is adjusted at the sender is via a potentiometer which controls the current to the transceiver. It has been brought to our attention that it is now possible to read out the received signal strength on the ESBs directly and this modification is part of future work in this area.

What we do instead is a "software version" of RSS by measuring the packet loss while varying the transmission power at the sender. In order to determine an approximation to the received signal strength, an anchor node writes its sending signal level into a packet, and the receiving node could then read out this value, take the minimum over all received packets, and thus determine the lowest signal level at which it can still "hear" the anchor. While this way of measuring the transmission power is certainly not the most precise way, it fulfills the natural assumption that greater perceived received signal strength means that the sender needed to use more power to reach the receiving node, thus the receiver is farther away. The exact correlation needs to be determined, but the important point we want to verify is that the same input level on the sender should reach the same distance given similar conditions.

A critical issue with these nodes is the susceptability of the radio signal strength to various outside influences. Two nodes might be as close as a couple of centimeters, but placed near a wall or close to some underground cable, or as far as a few meters, and both times the best received signal strength will be the same. Here, we build on the preliminary measurements of the transmission range as a function of the signal strength in various settings from project website [6].

## 3. RESULTS

We will now discuss our measurements, their results and the motivations that led from one experiment to another. In the following we use the terms sensor node and node interchangeably.

### 3.1 In the Lab

As discussed in Section 1 we want to implement popular positioning heuristics on real sensor nodes. Towards this goal we first need to obtain some data on the correlation between the power level received and the distance of the nodes without obstacles.

Our experimental setup is the following: In the corridor of our lab, an anchor node transmits 100 packets at each power level and a receiving node placed at a specified distance measures the number of packets it receives. This experiment is repeated for inter-node distances ranging from 1cm to 120cm, as in a first step we want to explore the accuracy of the distance-to-power relation on a small scale. The minimum power level at which a packet is received at a given distance is plotted in Figure 1. While the data points do not lie on the theoretically assumed parabola, they are almost monotone with slight deviations of at most three levels and the curve exhibits a certain regularity.

Most applications for wireless networks will, however, not be satisfied with a single packet arriving with some low probability. Therefore we furthermore test the link quality if we require that at least $x$ percent of the packets arrives at the listening node. The results for $x = 90$ are shown in Figure 2. The graph looks similar for all other values of $x$ (down to 50).

The deviation of the data set to the best-fit curve is in both experiments not negligible (about 10 units in the latter case), as can be seen in Figure 1 and 2. However, the data set is still well-behaved in the sense that a curve is discernable and this curve is almost monotone.

The conclusion that can be drawn from these experiments is that in a controlled environment with a clear line of sight, the distance-to-power function with a specified stability can be approximated to a certain extend by a monotonely in-

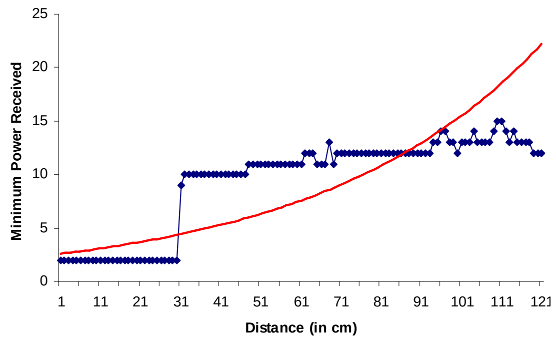


**Figure 1: The minimum power level which was received at the given distance.**

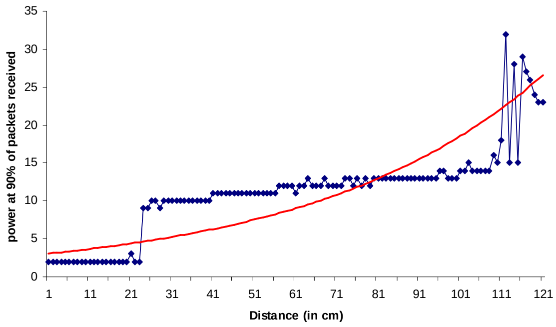


**Figure 2: The lowest power level at which at least** 90% **of the packets were received.**

creasing function thus supporting theoretical assumptions. In such a setting, a possible localization scheme could measure the packet loss for different power levels from anchor nodes (which know their position) and then use the inverse of the distance-to-level point map in Figure 2 to obtain a distance estimate.

## 3.2 In a Room

Any localization algorithm in the plane needs (approximate) distance measurements from at least three non-colinear anchors. Therefore, the next step in our experiments is to expose a single node in a room to several anchors and test the obtained measurements for their usability.

The experimental setup is similar to the one before. We place four anchor nodes on the corner of a rectangle in a room. A test node is then placed within that rectangle. An anchor sends out a packet at each power level from 1 up to 100, then the next anchor does the same, and so on, in a round robin fashion. Each time, the test node reports which packets it receives. Figure 3 shows how often a given power level is received by the test node from Anchor 1 over the course of the experiment. The results for the other anchors are similar and thus omitted.

This already goes to show that the link between the test

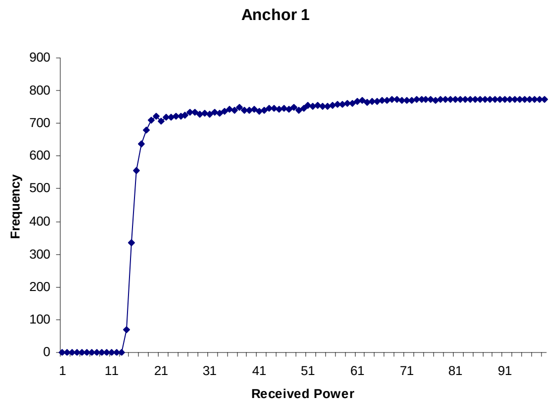


**Figure 3: The number of times a given power level was received. Anchor 1 was 1.39 meters from the test node.**

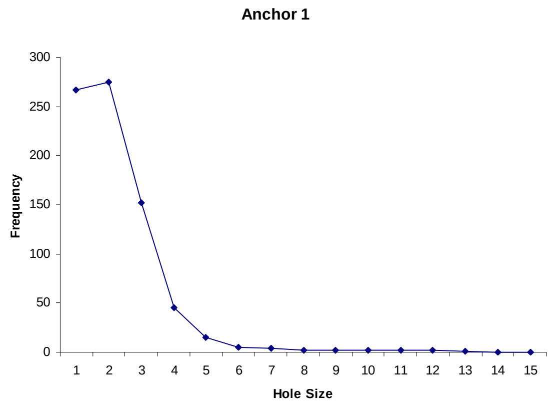


**Figure 4: The size of the "holes" in the received power level progression and how often they occured.**

node and the anchors is reasonably stable over time. As we already see from the earlier experiments in the lab, the *minimum* power level received, perhaps averaged over time, gives a good indication for the distance to the sending node. This is advantageous since it saves memory compared to storing all received power levels. To further strengthen this hypothesis, we also examine the size of the "holes" in the received power levels for each iteration. In other words, how big is the gap between one received power level and the next and how often does it occur. For example, if the first heard packet has level 15 and the next one heard level 18, then this results in a hole of size 3. Figure 4 shows that there are only small holes in most cases. Meaning that the minimum power level is a good estimate if the requirements are not too stringent.

In a further step, we add obstacles to our room by placing various everyday objects in the area of the rectangle. The result is that the general behavior of the link quality does not appear to be affected, seen in Figure 5. The curve has the same "shape" as before (less data points being the reason for the height difference). Astonishingly, the peak as resulting from the experiment with obstacles is shifted to the left compared with the non-obstacles experiment. Yet, this result only completes the picture of the unpredictability of real-world sensor node behavior.

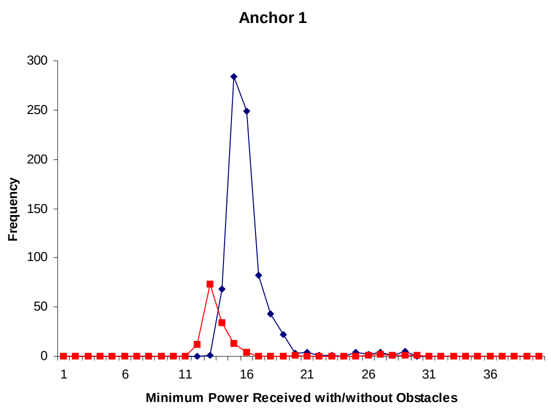


**Figure 5: Number of times that a given power level was the minimum frequency received. The higher peak (blue) is from the setup without obstacles, the other one (red) with obstacles. The reason for the lower peak is that the experiment was run for less time due to external contraints.**

Concluding this subsection, we can say that in an environment with several stable anchors, each sending out packets one at a time to avoid collisions, the links appear to be stable over time and exhibit a sharp peak.

## 3.3 Network

The results of the previous sections suggest that distance measurements based on radio power levels can be used if the accuracy constraints are not too tight. This leads us to set up a positioning experiment with four anchors at the corners of a rectangle (a table) and several nodes on the inside of the rectangle. We use a spring-based algorithm in which each node iteratively computes its new position as a function of its neighbors' positions. This approach stems from the graph drawing context and was first adapted to positioning in [14]. Since the heuristics proposed in [14] are far too power and resource consuming, we implement a simplified version which meets our purposes. For the power-to-distance conversion, the data gathered in the experiments above is used. Surprisingly, between most of the computed positions and the corresponding actual positions there is seemingly no correlation. The errors are in the order of the magnitude of the sensor field. A closer examination reveals that already the powers received do not correlate with the distances at all in the sense that a node close by needed significantly more power to communicate than some nodes far apart, even without any obstacles in the room. This finding is closer examined in the next experimental setup.

We perform the following experiment, the result of which can be seen in Figure 6. We take nine nodes placed arbitrarily in a large room. Iteratively each node sends out a series of packets, starting from level 1 to 100. In each iteration all non-sending nodes record the minimum power level

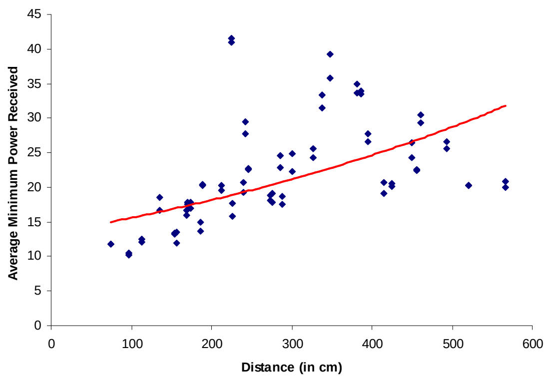


**Figure 6: Measured power levels at various distances. The x-axis is distance (in cm), the y-axis is average minimum power level. A point is the measurement of a node to some other node.**

received from the sending node in that iteration. As can be seen, the graph looks very different from the one in the first phase of the experiment, Figure 1. Whereas in the first phase of the experiments a curve connecting the data points is discernable, the data points measured in the current experimental setup are much more scattered. Observe also that the scale of the $y$-axis in both figures is different and that the deviation of the data set to the best-fit parabola in this experiments is about 30 units.

The conclusion to be drawn here is that in a two-node setting with a constant environment, the power expectedly increases approximately quadratically with distance. In a network, however, there are several nodes, unpredictable environmental conditions such as people walking around, and different positions of walls, cables, other computers and the like. So even without the effect of node transmission interference, which we excluded in this setup, the power-to-distance correlation becomes utterly useless. The data points are scattered across the graph. While *two nodes* in the network can have a quadratically increasing (theoretical) distance-to-power function, this function between *different pairs* of nodes is not necessarily related in a predictable manner.

On the positive side, this experiment allows us to conclude on the symmetry of the links between two nodes. To that end, we compare the minimum power levels between two nodes in the network. In Figure 7 the number of occurrences a power level was the lowest one received for an arbitrary pair of nodes is plotted. The peaks are sharp and coincide.

We furthermore compute the average minimum power level received for each node to all other nodes and take the difference between the node pairs. This value is rounded to the nearest integer and Figure 8 shows the number of times a difference occurs among the $\binom{9}{2} = 36$ node pairs.

# 4. DISCUSSION AND FUTURE WORK

In this paper, we presented an experimental study of the link quality in real-world sensor networks. If one expects the sensor network to be in place for an extended period, then we

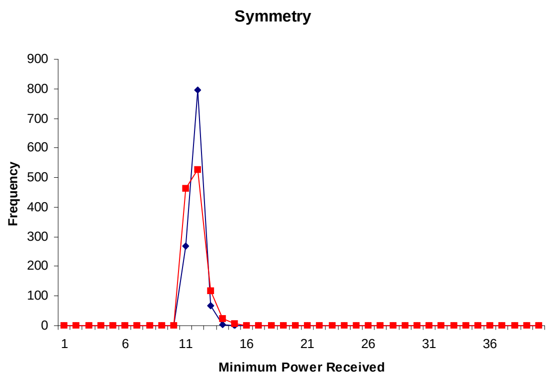


**Figure 7: Number of times that a given power level was the minimum frequency received. The blue curve is from node 75 to node 65 and the red curve the other way around.**

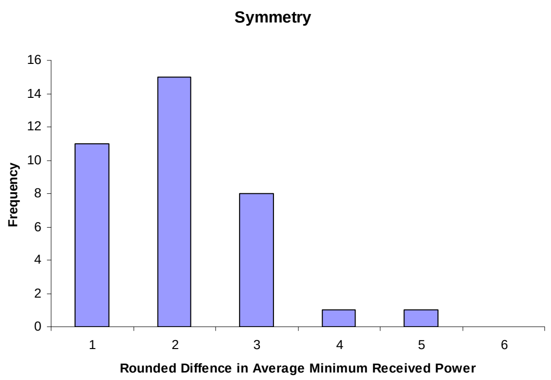


**Figure 8: Rounded differences between average minimum received power.**

do not necessarily require an unvarying connection between two nodes, but that the link characteristics should be *steady over time*. The minimum power at which a node receives a packet from a sender node was shown to be stable, sharp and symmetric over time in all experimental setups. Whereas this power level is a good indication for the distance between nodes if there are only two nodes in a static experimental setup, this correlation does not apply to all nodes pairs in a larger arbitrary network. Our experiments illustrate that in larger networks the power levels measured in lab conditions are not correlated with distances in the real world. Thus drastically new models are required for sensor networks if theoretical constructs are ever to be applied with a realistic chance.

Despite the pessimistic results, the work in this paper opens up a number of interesting directions meriting further investigation. First of all, these tests should be performed on different hardware to see whether this is a general problem. Second, with the new version of the ESBs, we can perform actual received signal strength measurements which would refine the results observed in this paper. Third, it would be interesting to investigate whether there is a qualitative difference between short and long range measurements.

We would like to thank the anonymous reviewers for their helpful suggestions.

# 5. REFERENCES

[1] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-based User Location and Tracking System. In *Proc. of Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM)*, 2000.

[2] R. Bischoff and R. Wattenhofer. Analyzing Connectivity-based, Multi-hop Ad hoc Positioning. In *Proc. of the IEEE Intl. Conf. on Pervasive Computing and Communications (PerCom)*, 2004.

[3] P. Biswas and Y. Ye. Semidefinite Programming for Ad Hoc Wireless Sensor Network Localization. In *Proc. of Intl. Symp. on Information Processing in Sensor Networks (IPSN)*, 2004.

[4] E. Dijk, K. van Berkel, R. Aarts, and E. van Loenen. A 3-D Indoor Positioning Method using a Single Compact Base Station. In *Proc. of the IEEE Intl. Conf. on Pervasive Computing and Communications (PerCom)*, 2004.

[5] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex Behavior at Scale: An Experimental Study of Low-Power Wireless Sensor Networks. Technical report, UCLA Computer Science Technical Report UCLA/CSD-TR 02-0013, 2002.

[6] ScatterWeb GmbH. `http://www.scatterweb.net`.

[7] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher. Range-Free Localization Schemes in Large Scale Sensor Networks. In *Proc. of Mobile Computing and Networking (MobiCom)*, 2003.

[8] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins. *Global Positioning Systems: Theory and Practice.* Springer, 5th edition, 2001.

[9] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. Geometric Ad-Hoc Routing: Of Theory and Practice. In *Proc. of Symp. on Principles of Distributed Computing (PODC)*, 2003.

[10] F. Kuhn, R. Wattenhofer, and A. Zollinger. Ad-Hoc Networks Beyond Unit Disk Graphs. In *Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, 2003.

[11] T. Moscibroda, R. O'Dell, M. Wattenhofer, and R. Wattenhofer. Virtual Coordinates for Ad hoc and Sensor Networks. In *Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, 2004.

[12] D. Niculescu and B. Nath. Ad Hoc Positioning System (APS). In *Proc. of IEEE Global Communications (GLOBECOM)*, 2001.

[13] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket Location-Support System. In *Proc. of Mobile Computing and Networking (MobiCom)*, 2000.

[14] A. Rao, C. Papadimitriou, S. Ratnasamy, S. Shenker, and I. Stoica. Geographic Routing without Location Information. In *Proc. of Mobile Computing and Networking (MobiCom)*, 2003.

[15] K. Römer. The Lighthouse Location System for Smart Dust. In *Proc. of ACM/USENIX International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2003.

[16] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *Proc. of the First IEEE Conference on Sensor and Adhoc Communication and Networks*, 2004.