

Adaptive Beacon Placement*

Nirupama Bulusu

University of California, Los Angeles and USC/Information Sciences Institute
bulusu@cs.ucla.edu

John Heidemann

USC/Information Sciences Institute
johnh@isi.edu

Deborah Estrin

University of California, Los Angeles and USC/Information Sciences Institute
destrin@cs.ucla.edu

Abstract

Beacon placement strongly affects the quality of spatial localization, a critical service for context-aware applications in wireless sensor networks; yet this aspect of localization has received little attention. Fixed beacon placement approaches such as uniform and very dense placement are not always viable and will be inadequate in very noisy environments in which sensor networks may be expected to operate (with high terrain and propagation uncertainties). In this paper, we motivate the need for empirically adaptive beacon placement and outline a general approach based on exploration and instrumentation of the terrain conditions by a mobile human or robot agent. We design, evaluate and analyze three novel adaptive beacon placement algorithms using this approach for localization based on RF-proximity. In our evaluation, we find that beacon density rather than noise level has a more significant impact on beacon placement algorithms. Our beacon placement algorithms are applicable to a low (beacon) density regime of operation. Noise makes moderate density regimes more improvable.

1. Introduction

In the near future, advances in processor, memory and radio technology will enable small and cheap nodes capable of wireless communication and significant computation. The availability of micro-sensors and low power wireless communications will enable the deployment of very dense, fully distributed sensor/actuator networks for a wide range

of environmental monitoring applications, from marine to soil and atmospheric contexts. We can imagine ad hoc sensor networks deployed for various kinds of applications, providing continuous and spatially dense observation of biological, environmental and artificial systems [5]. Moreover these systems will eventually incorporate actuation, as well as sensing: allowing these systems to influence and interact with their environment. These ad hoc sensor networks will consist of nodes strewn arbitrarily in the environment and will be largely unattended. Such unattended networks must self configure and reconfigure to adapt to the particulars of their environmental setting and the availability of other nodes within the system.

Spatial localization *i.e.*, determining physical location of a sensor node is a critical service in these wireless sensor networks [2]. Because many of these systems are embedded to monitor or control the behavior of physical systems (as compared with strictly virtual information systems), nodes often need to label data or events with location information. Similarly, nodes often need to determine their action based on their physical location (am I the right sensor to monitor a particular object?). Localization is indispensable for context-aware applications that select services based on location, efficient routing and communication, and power conservation by combining data from multiple sensors. Traditional information systems have not had to have such a location focus and as a result our support for localization is relatively weak. Existing geolocation systems such as GPS [9] do not always meet the operational (*e.g.* low power), environmental (*e.g.* indoors) or cost constraints. For example, the cellular phone community has chosen RF-based approaches to meet localization requirements [11]. Consequently, localization has been

*To appear in the Proceedings of the 21st International Conference on Distributed Computing Systems, Phoenix, Arizona, April 2001.

the subject of a burgeoning amount of research in the past two years. Various localization approaches have been proposed and they fall into two broad categories: proximity based approaches [16, 2, 14] and multilateration¹ based approaches [1, 17, 7, 18, 9]. Many of these approaches are RF-based including RF localization [2], using signal strength [12], and using a database of signal strength signatures [1].




Figure 1. Beacon density vs. granularity of localization regions

Beacon nodes which know their position and serve as a reference are a vital aspect of nearly every localization system. Beacon placement strongly affects the quality of localization. For example, consider the simple proximity based localization system shown in figure 1 (described in the next section). In such proximity based approaches, beacon density and placement control the granularity of localization. Even in multilateration based approaches, the number of visible beacons and their placement geometry is crucial. Each node may need to hear from a certain minimum number of beacons to be able to localize itself, and the beacon nodes heard must be non-collinear. Additionally, in RF-based approaches (especially [2, 12]), the quality of localization is strongly affected by the placement of these beacons because of the environmental conditions (as we explore later in this section).

Much of the research thus far has focused on localization techniques, and beacon placement issues (with the possible exceptions of [14, 2]) have not been sufficiently explored. This paper is a first step towards addressing specific algorithms for beacon placement.

Intuitively, approaches such as uniform and very dense beacon placement should suffice to improve the quality of localization.

- Uniform placement is good, but insufficient due to the following reasons:

¹In multilateration, position is estimated from distances to three or more known points.

- Beacons may be perturbed during deployment. Consider for instance, a terrain comprising of a hilltop. Air dropped beacon nodes will roll over the hill, while lighter sensor nodes may stay atop the hill.
- Even when beacon placement is uniform, noise (in the form of terrain and propagation uncertainties) may affect the visibility of beacons that should be in range. Radio signal propagation in general is significantly affected by multi-path effects, fading, shadowing *etc.* Uneven terrains and obstacles bring in an additional dimension of uncertainty [15].
- Very dense placement may not be a viable solution due to several reasons:
 - Cost or Power: The cost of the beacons may preclude very dense beacon placement. Power considerations may require that only a restricted smaller subset of beacon nodes be active at any given time so as to prolong system lifetime [5, 19].
 - Terrain Commonality: Even when cost is not a concern, the environmental or terrain conditions may be such that merely increasing the density uniformly will not overcome the problem. For instance, if the number of air-dropped beacons were doubled, the same situation would persist. Also, the terrain may already have a very high density of beacons (enough to achieve the maximum possible quality of localization under ideal conditions) and hence the new beacons must be added in particular places to cope with noise.
 - Self-interference: At very high densities, the probability of collisions among signals transmitted by the beacons increases. Therefore even if we had unlimited numbers of beacons, we would like to limit their use.

The fundamental limitation of these two approaches is that they are basically fixed strategies, that do not take into account environmental conditions that cannot be predicted a priori. It is virtually impossible to preconfigure to such terrain and propagation uncertainties and compute an ideal (or even satisfying) beacon placement that uniformly achieves a desired quality of localization across the region. Clearly, the beacon placement needs to adapt to the noisy and unpredictable environmental conditions. This paper addresses this problem of *adaptive beacon placement*: given an existing field of beacons, how should additional beacons be placed for best advantage.

The paper makes the following contributions:

- We develop novel algorithms that permit a person or mobile robot to place additional beacons to incrementally extend an initial beacon field. This allows for empirical adaptation to terrain conditions.²
- We evaluate the gains from incrementally improving an RF-based location field.
- We introduce the notion of solution space density. The efficacy of algorithms (such as our beacon placement algorithms) designed to work in noisy environments is predicated on the assumption that the solution space for the problem must be dense in number of satisfying solutions. For instance, if the only way to improve the quality of localization in a region by adding an additional beacon is to place it at a single point in the region, then it is difficult to design algorithms that can identify that point in the presence of so much noise.

Although we evaluate our algorithms in the context of beacon placement for RF-based localization they may generalize to other problem domains where node placement is rather critical: global coverage or universal connectivity in wireless sensor networks, measurement based repositioning of seismic sensor nodes (surface conditions, coupling with the ground *etc.* are significant influences on the quality of sensing attainable in these nodes). In traditional Internet web caching, the placement of web caches *etc.*, may be done based on analyses of web traffic, web server requests *etc.*

2. Connectivity based RF localization

In this section, we review the connectivity based RF localization approach that was introduced and described in detail in an earlier paper [2]. The localization algorithm uses an idealized radio model, discussed next.

2.1. Idealized Radio Model

An idealized radio model is useful for predicting bounds on the quality of connectivity based localization. The model is simple and easy to mathematically reason about. It makes two rather unrealistic assumptions: perfect spherical radio propagation and identical transmission range (power) for all radios. In section 4 we relax the first constraint by introducing noise. The second constraint is appropriate for nodes with identical, fixed power radios.

²In general, the SCOWR project focuses on incorporating robotic motion and communication into distributed sensing applications (e.g see [13]).

2.2. Localization Algorithm

Beacons situated at known positions, (X_i, Y_i) , transmit periodically with a time period T . Clients listen for a period $t \gg T$ to evaluate connectivity. If the percentage of messages received from a beacon in a time interval t exceeds a threshold CM_{thresh} , that beacon is considered connected. A client then estimates its position (X_{est}, Y_{est}) as the centroid of the positions of all connected beacons.³

Given the actual position of the client (X_a, Y_a) , we can compute the accuracy of the localization estimate or the *localization error* LE , which is the distance between the client's estimated and actual positions.

$$LE = \sqrt{(X_{est} - X_a)^2 + (Y_{est} - Y_a)^2}$$

By increasing the density of the beacons that populate the grid, the granularity of the localization regions becomes finer, and hence the accuracy of the location estimate improves. This is illustrated in figure 1.

However, the approach of using the centroid implies that there is often some level of error. We have previously analyzed this error [2] and found that it is bound by the nominal beacon transmission range R and the separation distance between adjacent beacons d (under uniform beacon placement). For a range overlap ratio (R/d) of 1, the maximum error is bound by $0.5d$. This factor falls off considerably (to $0.25d$), when the range overlap ratio increases (to 4).

Alternative approaches would consider additional information of time-of-flight [18] or signal strength [12]. Although we did not choose to use these approaches because of very short times-of-flight and limited resolution of signal strength, these approaches too are influenced by beacon placement.

3. Adaptive Beacon Placement

Given a localization algorithm, one must deploy a field of beacons as infrastructure, and then extend this field if it proves insufficient.

Our approach to incremental improvement of localization through beacon placement is based on *empirical adaptation*. By *adaptation*, we mean we are improving the quality of localization by adjusting beacon placement or adding a few beacons rather than by completely re-deploying all beacons. By *empirical*, we mean the deployment of additional beacons is influenced by measurements of the operating localization system rather than by careful or complete off-line analysis of a complete system model.

³Under our idealized radio model, the client lies within the locus of points described by the intersection of a set of circles with centers corresponding to the positions of connected beacons and radii R . The centroid summarizes the locus. An alternative representation of the localization estimate is the full locus information.

Our general approach is to use a GPS-equipped mobile robot or human to explore the terrain. We assume that the robot (or human) can determine its geographic position using a high precision differential GPS receiver and map it to the local coordinate system. The robot has a short range radio similar to the one used by the sensor nodes, and can thus compute its localization estimate using the connectivity based localization algorithm. Thus it has a means of computing the localization error at any point on the terrain. It also has a capability to carry a certain number of beacons that it can deploy as additional beacons wherever it deems necessary. Therefore, based on its measurements of localization error at different points in the region, it must compute good places to deploy additional beacons. We define this problem as *adaptive beacon placement*.

3.1. Assumptions

Completely exploring the design space of possible robot-based beacon placement algorithms is a large problem. We have begun with a simple choice: an off-line algorithm with complete terrain exploration and no measurement noise. We use this simple problem to define the problem and preliminary solutions. We are currently working on ways to generalize these solutions.

3.2. Beacon Placement Algorithms

We have three simple off-line beacon placement algorithms. The goal of these algorithms is to determine candidate points for placement of an additional beacon, so as to maximize the gains obtained. These three algorithms differ in the amount of global knowledge and processing used to make their decision.

We assume the terrain to be a square of $Side$ meters and each robot will take measurements $step$ meters apart ($step < Side$). As defined earlier, beacons have a nominal transmission range R .

3.2.1 Random

This is the simplest algorithm, which pays no attention to the quality of localization at different areas of the region and simply selects a random point in the region as a candidate point for adding an additional beacon.

Step 1 Select a random point (X_r, Y_r) in the terrain.

Step 2 Add a new beacon at (X_r, Y_r) .

We investigate this primarily for comparison with the other algorithms, but also because it is similar in character to uncontrolled airdrop of additional nodes. The complexity of this algorithm is $O(1)$.

3.2.2 Max

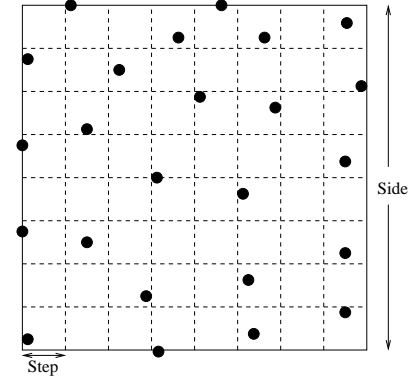


Figure 2. Illustration of the Max algorithm

The Max algorithm (illustrated in figure 2) can be described in three steps:

Step 1 Divide the terrain into $step \times step$ squares.

Step 2 Measure localization error at each point $(i \times step, j \times step)$ in the terrain that corresponds to a square corner. ($0 \leq i, j \leq \frac{Side}{step}$)
Number of data points in the terrain, $P_T = (\frac{Side}{step} + 1)^2$.

Step 3 Add new beacon at the point (X_m, Y_m) that has the highest measured localization error among all points.

This algorithm is predicated on the assumption that points with high localization error are spatially correlated. The advantage of this algorithm is that it can be computed in a very straightforward way. However, it may be overly influenced by propagation effects or random noise that may cause very high localization error at one point while the localization error at points very close to it remains low; *i.e.*, it is sensitive to local maxima.

The complexity of the Max algorithm is linear in P_T , the number of data points at which the localization error is measured *i.e.*, $O(P_T)$.

3.2.3 Grid

The Grid approach to determining a candidate point is to compute the cumulative localization error over each grid, for several overlapping grids in the terrain. This is based on the observation that adding a new beacon affects its nearby area, not just the point where it is placed.

The Grid algorithm (illustrated in figure 3) consists of the following steps:

Steps 1 and 2 are the same as the Max algorithm.




Figure 3. Illustration of the Grid algorithm

Step 3 Divide the terrain into N_G partially overlapping grids as follows:

Step 3.1 Each grid has a side, $gridSide = 2R$. Thus each grid encloses the radio reachability region of its center.

Step 3.2 For $1 \leq i, j \leq \sqrt{N_G}$, the grid $G(i, j)$ is defined by its center $G_c(i, j) = (X_c(i, j), Y_c(i, j))$ where

$$X_c(i, j) = \frac{gridSide}{2} + \frac{(i-1) \cdot (Side - gridSide)}{\sqrt{N_G} - 1}$$

$$\text{and } Y_c(i, j) = \frac{gridSide}{2} + \frac{(j-1) \cdot (Side - gridSide)}{\sqrt{N_G} - 1}.$$

Step 4 For each grid $G(i, j)$, compute the cumulative localization error $S(i, j)$ at all the points measured in Step 2 that lie in the grid $G(i, j)$. Number of data points per grid, $P_G = P_T \times \frac{Side^2}{(2 \cdot R)^2}$.

Step 5 Add the new beacon at the center $G_c(i, j)$ of the grid $G(i, j)$ with the maximum cumulative localization error.

While the Grid algorithm has the advantage that it can improve many points at once, it is computationally far more expensive than the Max and Random algorithms because it additionally divides the terrain into several grids and computes the cumulative localization error in each grid.

The complexity of the Grid algorithm is linear in the product of N_G , the number of grids considered and P_G , the number of data points per grid at which the localization error is measured., i.e $O(N_G \cdot P_G)$.

Section 4 provides a performance comparison of these three algorithms. We note that these are by no means the only possible algorithms, but these are representative of the effectiveness attainable with different degrees of processing.

4. Performance Evaluation

In this section, we report on some results from a preliminary performance evaluation of our beacon placement algorithms. We use numeric simulations to explore, in some detail, the implications of several design choices.

4.1. Goals, Metrics and Methodology

Our goals in conducting this evaluation study were three-fold:

- Place the performance of Grid and Max algorithms in the context of the Random algorithm. This serves as a sanity check for the intuition behind the Grid and Max algorithms, as also to explore the influence of the level of knowledge on algorithm performance.
- Understand role of beacon density on algorithm performance.
- Understand impact of noise such as propagation losses and terrain features on the beacon placement algorithms.

We choose two metrics to analyze and compare the performance of our algorithms. These metrics are statistics evaluated for the observed localization error at all $step \times step$ square corners obtained by subdividing the region.

Improvement in Mean Error computes the difference between mean localization error at all measured points in the terrain before and after the beacon is added. This metric indicates the overall impact of adding a beacon to quality of localization in the entire terrain.

Improvement in Median Error computes the difference between the median localization error at all the measured points in the terrain before and after the beacon is added. This metric indicates the improvement due to adding a beacon on the quality of localization at the top 50% of the points with the highest localization error in the terrain.

We study these metrics as a function of beacon density. We consider a square terrain of side 100m. Each node has a nominal radio range of 15m. To study the performance of our algorithms as a function of beacon density, we generate a variety of beacon fields of different densities.

In each of our experiments, we vary the number of beacon nodes from 20 to 240 in increments of 10 beacon nodes. The corresponding beacon density varies from 0.002 beacons per square m to 0.024 beacons per square m. To put these density values in context, the corresponding number of beacons per nominal radio coverage area ($\pi \cdot R^2$) varies

Table 1. Simulation Parameters

| Parameter | Value |
|-------------|-------|
| <i>Side</i> | 100m |
| <i>R</i> | 15m |
| <i>step</i> | 1m |
| N_G | 400 |

from 1.41 to 17. For each density, we generate 1000 different beacon fields. Each beacon field is generated by randomly placing the beacons in the $100m \times 100m$ square terrain. The performance metrics, for each algorithm and beacon density, are averaged over the 1000 beacon fields. To characterize the stability of our results, all graphs include 95% confidence intervals. The simulation parameters are listed in Table 1.

4.2. Simulation Results

As observed earlier, beacon density has a considerable impact on quality of localization. To quantify this effect, we evaluate the relationship between mean localization error and beacon density. Figure 4 graphs mean localization error for varying beacon densities under idealized radio propagation conditions. We see that the mean localization error falls sharply with increasing beacon density, until it reaches a density of 0.01 beacons per square m (approximately 7 beacons per nominal radio coverage area) and saturates at around 4m ($0.3R$). We refer to this density as the *saturation beacon density*. There is little to be gained from deploying beacon nodes at more than this density.




Figure 4. Mean localization error vs. beacon density (Ideal)

Our first experiment compares the three algorithms under idealized radio propagation conditions (*i.e.*, perfect connectivity for distances $\leq R$, no connectivity otherwise). The aim of this experiment is to isolate and study the im-

part of beacon density on the Random, Max and Grid beacon placement algorithms.

Figures 5 graphs the improvements in mean and median localization error for various beacon densities. As expected the Random algorithm has the least improvement.

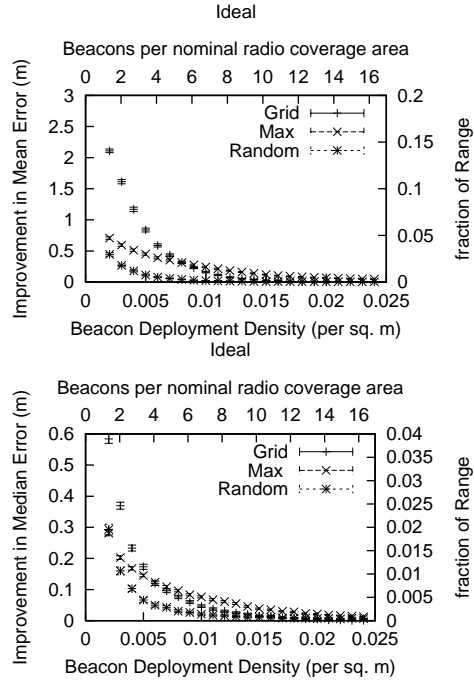


Figure 5. Improvement in mean and median errors vs. beacon density (Ideal)

At low densities (≤ 0.005 , much below saturation density), the Grid algorithm clearly performs best, with improvements in mean localization error at least twice that of the Max algorithm. Grid achieves such performance because it considers the quality of localization over a grid, and can improve many points at once. The performance of Max is slightly better than Grid for regions of moderate density (0.008 to 0.02 per square m). At these densities, the points with maximum localization error are very loud, and Max suppresses them better. At very high beacon densities (≥ 0.02 beacons per square m), the quality of localization is saturated, and the performance of the three algorithms is about the same.

A similar trend with respect to beacon density is observed for the median localization error, although the improvements in median localization error are relatively more modest (roughly 25% of the improvements in the average localization error with Grid). This is because the algorithms are effective in fixing a few hot spots with high localization error with the addition of a single beacon rather than in improving the localization throughout the terrain.

From our analysis we infer that at least under idealized conditions, our beacon placement algorithms (Grid and Max) are applicable only to a regime corresponding to low or insufficient beacon density.

4.2.1 Impact of Noise

As stated earlier, idealized radio propagation conditions are rather unrealistic. Random noise can severely affect radio connectivity [15], and thereby degrade the quality of localization. Since this noise cannot be predicted, beacon placement algorithms must adapt to it through measurements. To study the impact of noise on our beacon placement algorithms, we model random propagation noise as follows. For each beacon field, connectivity to any beacon B at any given point P is determined based on a noise model. In our noise model, connectivity to a beacon B exists at a point P , if $distance(P, B) \leq R(1 + u.nf(B))$. $nf(B)$ is the noise factor of the beacon B , and is chosen uniformly between 0 and $Noise$, the maximum noise factor for the field. u is chosen uniformly at random between -1 and 1 . The intent was to create non-uniform propagation noise for the beacons, and to create random regions with higher propagation noise than the rest of the location field. We did this because the impact of noise is less evident when each beacon has an identical propagation field. Note that this noise model is location based and static with respect to time *i.e.*, not time varying. We use 4 different settings of $Noise$, 0 (corresponding to Ideal propagation), 0.1, 0.3 and 0.5.




Figure 6. Mean localization error vs. beacon density (Noise)

To quantitatively characterize the impact of noise, we evaluate variation in mean localization error with beacon density in the presence of noise. Figure 6 graphs mean localization error for various beacon densities and noise levels. We observe a steady increase in mean localization error for each beacon density (*e.g.*, from 18m to 23m for 0.02 beacons per square m) and the saturation beacon density (from 0.01 to 0.015 beacons per square m) as the level of

$Noise$ increases from 0 to 0.5. The mean localization error follows the same general trend with increasing beacon density with noise as with idealized radio propagation.

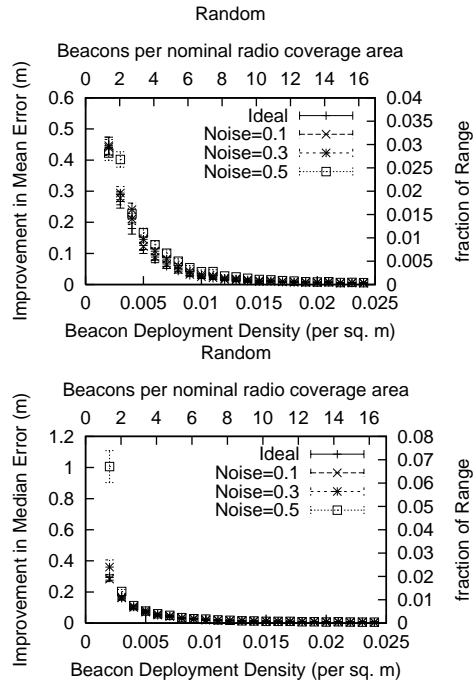


Figure 7. Performance of the Random algorithm with Noise

Figure 7 graphs the improvement in the mean and median localization error with the Random algorithm, for various beacon deployment densities and noise levels. The gains in both metrics with the Random algorithm are somewhat unchanged with noise. This is as expected, because noise is not an input in the Random algorithm, which does not make any measurements.

Figures 8 and 9 graph the improvement in the mean and median localization error with the Max and Grid algorithms respectively, for various beacon deployment densities and noise levels. We observe that noise makes regions of moderate beacon densities (0.005 to 0.01 beacons per square m) more improvable (improvements of 0.5m to 1m in mean error for corresponding increases in mean error of 1m to 3m) with the Grid algorithm, and to a lesser extent with the Max algorithm. The improvements to the median error are relatively unchanged with noise, because as we noted earlier, the focus of the algorithms is on improving a few hot spots.

4.3. Summary of Results

There are several lessons that we can draw from this preliminary evaluation of our beacon placement algorithms:

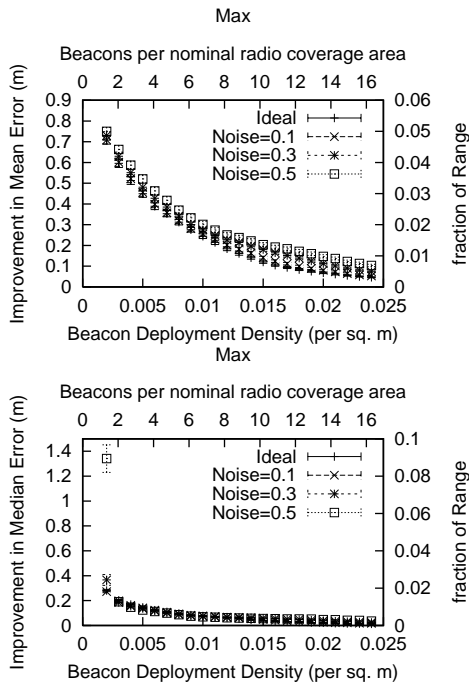


Figure 8. Performance of the Max algorithm with Noise

- Our beacon placement algorithms are applicable to a regime corresponding to low or insufficient beacon density deployment (≤ 0.01 beacons per square m or 7 beacons per nominal radio coverage area in the ideal case).
- At low densities, the Grid algorithm has the potential for significant improvements to the mean and median errors compared to the Max or Random algorithms.
- When noise level is increased from 0 to 0.5, there is a steady increase in both the mean localization error (up to 33%) and saturation beacon density (up to 50%).
- The Grid algorithm is clearly the best and superior to Max and Random algorithms even in the presence of noise. The performance of the Random algorithm is unchanged with noise, whereas noise makes even moderate density regions more improvable with the Grid algorithm.

5. Related Work

To our knowledge, empirically adaptive algorithms for beacon placement have not been previously studied in the networking literature. However, our work has been informed and influenced by a variety of other research efforts in several fields, which we now describe.

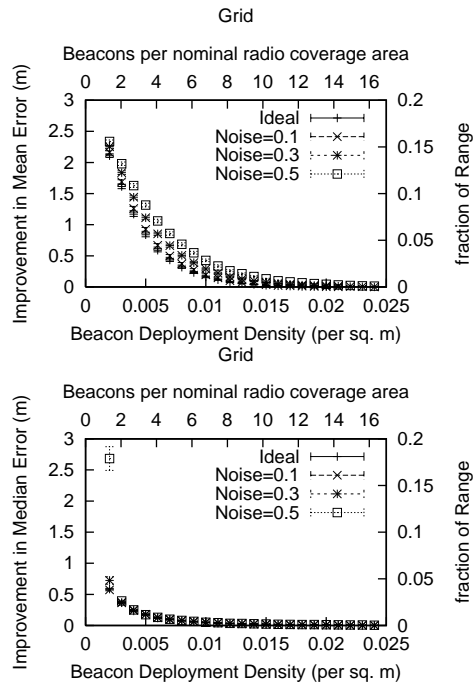


Figure 9. Performance of the Grid algorithm with Noise

Empirical *i.e.*, measurement based adaptation has served as a powerful design principle for various networking protocols, including TCP, SRM [6], and measurement based admission control [10]. TCP adaptively sets its timers or congestion control windows based on round trip time measurements in order to adapt to a wide range of link bandwidths while maintaining high performance. Algorithms in the Scalable Reliable Multicast framework (SRM) [6] dynamically adjust their control parameters based on observed performance within a multicast session. This allows applications using the SRM framework to adapt to a wide range of group sizes, topologies and link bandwidth while maintaining robust and high performance. The measurement based admission control algorithm described in [10] uses ongoing measurements rather than apriori characterization to determine behavior of existing flows, which enables it to provide predictive service with fairly reliable delay bounds at network utilization significantly higher than those achievable under guaranteed service.

Within the context of unattended ad hoc sensor networks, the design of adaptive algorithms as a self-configuring mechanism is a burgeoning area of research. For instance, the AFECA algorithm proposed in [19], exploits node deployment density and demonstrates adaptive fidelity. It adapts sleep times based on node density, scaling back node duty cycles (and so reducing routing “fidelity”) when many

interchangeable nodes are present. This allows it to substantially increase the network lifetime.

The significance of beacon placement in determining the overall quality of a service such as localization or coverage, has been stressed by others as well. The Cricket Location Support System [14], which is also proximity based, proposes deployment guidelines for beacons in indoor environments based on practical considerations. Whenever a beacon is placed to demarcate a physical or virtual boundary corresponding to a different space, it must be placed at a fixed distance away from the boundary demarcating the two spaces. Such placement ensures that a receiver rarely makes a wrong choice, unless caught within a small distance from the boundary between two beacons advertising different spaces.

In robotics, art gallery and pursuit evasion [8] problems have been well studied. In the “art-gallery” analogy, the robot’s goal is to move from one position to another to maximize visual coverage of its surroundings, as a human might try to do in a gallery. A complementary set of approaches addresses the pursuit-evasion problem in which a robot tries to move so as to evade observation or capture by mobile trackers. However these approaches are based on modeling the environment as a polygon and account for neither the noise nor the wide variety of terrain conditions one would expect to encounter for ad hoc sensor networks.

Facility Location [3, 4] problems are a well known class of theoretical computer science problems and have been the subject of extensive research over the past thirty-five years. In these facility location problems, there is a set of locations, where the cost of building a facility at location i is $f(i)$; furthermore, there is a set of client locations (such as stores) that require to be serviced by a facility, and if a client at location j is assigned to a facility at location i , a cost of $c(i, j)$ is incurred. The objective is to determine a set of locations at which to open facilities, so as to minimize the total facility and assignment costs. Since these problems are NP-hard, it is unlikely that there exist efficient algorithms to find optimal solutions. Instead, the focus has been on designing algorithms that are guaranteed to find solutions within a particular factor of the optimum. Solutions are based on linear relaxations to the natural integer programming formulations that yield extremely good lower bounds.

The novel aspect of the work described in this paper is applying the concept of *empirical adaptation* to beacon placement, in order to cope with a wide variety of noisy environments.

6. Future Work

We have identified a number of areas for future work. In our simulations, we used a probabilistic propagation model

to stress our algorithms under high levels of random noise and examine their behavior for a broad spectrum of conditions. We plan to do further simulations with a more sophisticated terrain map and propagation model (incorporating time varying propagation loss) to analyze the effects of terrain commonality. We also plan to evaluate the algorithms with respect to the gains obtained when several beacons are added at once (instead of just one beacon).

As stated in 2, our localization estimate is represented by the centroid of connected beacons rather than full locus information of the corresponding overlap region. Knowledge of loci enables a new perspective on adaptive beacon placement, such as adding new beacons to break down the loci with the largest area into smaller loci. To some extent, the Grid algorithm incorporates this strategy. Although the locus information is not reliable under non ideal radio propagation conditions (as in the real world), such algorithms are worth pursuing from a theoretical standpoint.

Our approach to adaptive beacon placement was based on exploration and instrumentation of the terrain by a person or mobile robot. While this approach has the advantage of instrumenting any point in the terrain through human or robot agent enabled mobility, the scope for designing truly distributed algorithms is limited. An alternative approach, which we plan to explore is beacon based; wherein, a reasonably dense beacon deployment is assumed, and the beacon nodes themselves instrument the terrain conditions based on interactions with other (beacon) nodes, and decide whether to turn themselves on *i.e.*, be active or be passive.

We addressed beacon placement in the context of proximity based localization approaches. An interesting point of comparison are beacon placement algorithms for multilateration based localization approaches, as the error characteristics of the two are significantly different. In the former approach, localization error is governed by beacon placement and density, whereas in the latter approach, it is influenced by the geometry of the beacon nodes. We plan to recast our existing beacon placement algorithms for multilateration based localization approaches.

7. Conclusions

In this paper we emphasized the importance of beacon placement in localization approaches and motivated the need for empirically adaptive beacon placement algorithms. We described the notion of solution space density, which forms the basis for our algorithms.

We outlined a general approach for adaptive beacon placement based on exploration and instrumentation of the terrain by a mobile human or robot agent. We designed and evaluated three algorithms based on this approach: Grid, Max and Random. Our algorithms are applicable to a regime of low beacon density deployment, wherein, Grid

clearly outperforms the Max and Random algorithms. In our simulations, beacon density rather than the noise level has a higher impact on the performance of beacon placement algorithms. When the noise level is increased from 0 to 0.5, there is a steady increase in both mean localization error (up to 33%) and saturation beacon density (up to 50%). The algorithms exhibited the same relative trend in the presence of noise as in an ideal scenario, although noise makes regions of moderate beacon density more improvable.

This paper is an initial foray in addressing the problem of beacon placement. The novel aspect of our approach is the emphasis on empirical adaptation.

Acknowledgments

The authors wish to thank Lewis Girod, Ashish Goel, Mani Srivastava and Gaurav Sukhatme for their valuable suggestions and feedback. This research is supported by the SCOWR project at USC through NSF grant ANI-9979457.

References

- [1] P. Bahl and V. N. Padmanabhan. Radar: An in-building user location and tracking system. In *Proceedings of the IEEE Infocom 2000*, volume 2, pages 775–84, Tel-Aviv, Israel, March 2000.
- [2] N. Bulusu, J. Heidemann, and D. Estrin. Gps-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, October 2000.
- [3] M. Charikar, S. Guha, D. Shmoys, and E. Tardos. A constant-factor approximation algorithm for the k median problem. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, pages 1–10, Atlanta, Georgia, USA, May 1–4 1999.
- [4] F. A. Chudak and D. Shmoys. Improved approximation algorithms for capacitated facility location problems. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages S875–S876, Baltimore, Maryland, USA, January 17-19 1999.
- [5] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99)*, pages 263–270, Seattle, Washington, USA, August 15–20 1999.
- [6] S. Floyd, V. Jacobson, C.-G. Liu, S. McCanne, and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *Proceedings of ACM SIGCOMM '95*, pages 342–256, Cambridge, Massachusetts, USA, August 1995.
- [7] L. Girod. Development and characterization of an acoustic rangefinder. Technical Report USC-CS-TR-00-728, University of Southern California, April 2000.
- [8] L. Guibas, , J. C. Latombe, S. M. LaValle, D. Lin, and R. Motwani. A visibility-based pursuit evasion problem. *International Journal of Computational Geometry Applications*, 9(5):471–494, October 1999.
- [9] B. Hofmann-Wellenhoff, H. Lichtenegger, and J. Collins. *Global Positioning System: Theory and Practice, Fourth Edition*. Springer Verlag, 1997.
- [10] S. Jamin, P. B. Danzig, S. Shenker, and L. Zhang. A measurement-based admission control algorithm for integrated services packet networks. In *Proceedings of ACM SIGCOMM '95*, pages 2–13, Cambridge, Massachusetts, USA, August 1995.
- [11] H. Koshima and J. Hoshen. Personal locator services emerge. *IEEE Spectrum*, 37(2):41–48, February 2000.
- [12] F. Koushanfar, A. Savvides, G. Veltri, M. Potkonjak, and M. B. Srivastava. Iterative location discovery in ad hoc sensor networks. *Slides from UCLA CS review*, April 2000. http://www.cs.ucla.edu/farinaz/location2_1.ppt.
- [13] M. Mataric and G. Sukhatme. Embedding robots into the internet. *Communications of the ACM*, 43(5):67–73, May 2000.
- [14] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location support system. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, pages 32–43, Boston, Massachusetts, USA, August 6-11 2000.
- [15] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall PTR, 1996.
- [16] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.
- [17] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications Magazine*, 4(5):42–47, October 1997.
- [18] J. Werb and C. Lanzl. Designing a positioning system for finding things and people indoors. *IEEE Spectrum*, 35(9):71–78, September 1998.
- [19] Y. Xu, J. Heidemann, and D. Estrin. Adaptive energy-conserving routing for multihop ad hoc networks. Research Report 527, USC/Information Sciences Institute, October 2000.