EXPLORING LANDMARK PLACEMENT STRATEGIES FOR SELF-LOCALIZATION IN WIRELESS SENSOR NETWORKS

Farid Benbadis[1], Katia Obraczka[2], Jorge Cortés[2], Alexandre Brandwajn[2]

[1]Laboratoire LIP6/CNRS, UMR 7606
Université Pierre et Marie Curie – Paris VI
Paris, France
{benbadis}@rp.lip6.fr *

[2]Baskin School of Engineering
University of California, Santa Cruz
Santa Cruz, CA, USA
{katia,jcortes,alexb}@soe.ucsc.edu

ABSTRACT

In this paper, we explore the impact of reference node, or "landmark", placement on the accuracy of the coordinate systems built using topology-based localization techniques. Such techniques employ landmarks to which each node computes its hop-count distance. A node's coordinates is given by the hop-count distance to all landmarks. To our knowledge, our paper is the first to study the impact of landmark placement on the accuracy of the resulting coordinate system. We show that placing landmarks on the periphery of the topology yields more accurate coordinate systems when compared to placing landmarks in the interior of the topology. Nevertheless, our simulation results also show that, in general, if enough landmarks are used, random landmark placement yields comparative performance to placing landmarks on the boundary randomly or equally spaced. This is an important result since boundary placement (especially at equal distances) may turn out to be infeasible and/or prohibitively expensive (in terms of power consumption as well as processing– and communication overhead). This is also the first study to consider not only uniform–, synthetic topologies, but also, non-uniform topologies resembling more concrete deployments.

I INTRODUCTION

Sensor networks typically refer to a collection of nodes that have sensing–, processing–, storage–, and (wireless) communication capabilities. In general, because of their small form factor and low cost, sensor network nodes often have limited capabilities; furthermore, as they are frequently battery powered, energy is a premium resource that needs to be conserved in order to maximize the lifetime of nodes and the sensor network as a whole.

Because of their ability to embed themselves in the real world, sensor networks have a wide range of applications with significant scientific and societal relevance [1]. Example applications [2] include environmental monitoring, object tracking, surveillance, and emergency response and rescue operations. While some scenarios allow for manual placement of sensor network nodes in the field, others require "random" deployment where nodes are simply "dropped" (e.g., from an airplane), and once they land they need to self-organize into a network and start performing the task at hand.

One important step in self-organization is positioning, which refers to having nodes find their physical location. Node positioning is required by sensor network core functions such as topology control, data aggregation, and routing [3, 4, 5], and may also be needed by a number of applications. For instance, the sensor network could be tasked to report the air temperature's running average by geographic region.

One clear solution to the positioning problem is provided by satellite-based systems [6, 7, 8], among which GPS (Global Positioning System) is probably the most widely utilized. However, in some scenarios, the use of satellite-based localization is not possible. This is the case of indoor, underwater, and underground deployments. Furthermore, equipping sensor nodes with GPS receivers might be prohibitive for reasons related to cost, form factor, energy consumption, or a combination thereof. A possible alternative is to equip only a subset of the nodes with GPS receivers and have all other nodes compute their position relative to the GPS-capable nodes. For instance, in a multi-tiered heterogeneous deployment, nodes that have extended life batteries and/or have higher processing power could have GPS capabilities. However, this may still be infeasible in some deployments.

To address this problem, numerous GPS-less methods have been proposed. In general, these methods may be classified as (1) using physical measurements or (2) using topological information. Examples of measurement-based GPS-less techniques include mechanisms that use propagation laws [9] to approximate Euclidean distance using received signal strength (RSS). The RSS can be converted into distance either directly, if the propagation law is uniform and known, or using multiple signals and time difference of arrival (TDoA) [10]. Then, trilateration techniques allow node coordinates to be inferred. The use of directional antennas to triangulate positions has also been proposed [11]. One main drawback of measurement-based mechanisms is that they typically require specialized equipment/capabilities to perform the measurements.

Topology-information based positioning, on the other hand, relies solely on topological information. For example, NO-GEO [12] first discovers border nodes, then computes their relative coordinates, and finally infer, through a relaxation method, non-border node coordinates relative to border nodes.[1] Alternatively, in GPS-FREE-FREE [14], JUMPS [15], VCAP [16], and BVR [5], the hop distances to reference nodes, or "landmarks," are transformed into "virtual coordinates".[2] GPS-FREE-FREE uses trilateration to ob-

*This work has been partially supported by the RNRT project SVP under contract 01504.

[1]The correctness of this algorithm is analyzed in [13].

[2]The hop distance from a node to a landmark is given by the minimum

1-4244-1144-0/07/$25.00

tain virtual coordinates from corresponding hop distances, while JUMPS, VCAP, and BVR use the hop distances directly as a nodes' coordinates. The denser the network, the more accurate it is to approximate Euclidean distance using hop distance.

However, existing hop-count based positioning systems make the strong assumption that, for better performance (e.g., accuracy), landmarks need to be placed along the perimeter of the topology at equal distances from one another. To the best of our knowledge, this assumption is purely intuitive, and has never been justified either empirically, experimentally, or analytically.

Thus, the focus of this paper is to explore the effect of landmark placement on the accuracy of the resulting coordinate system. Our paper, to our knowledge, is the first to show that, indeed, placing landmarks on the periphery of the topology yields more accurate coordinate systems than when landmarks are placed anywhere in the interior [3]. This is also the first study to consider not only uniform topologies, but also, non-uniform ones resembling more concrete deployments. In our study, we evaluate different landmark placement strategies, namely: 1) "uniform boundary placement" as inJUMPS and VCAP, where landmarks are placed at the boundary of the topology at equal distances from one another; 2) "random boundary placement", where landmarks are placed on the boundary but at random intervals; and 3) "random placement" which places landmarks anywhere in the topology completely at random. As performance metrics, we consider the ability to uniquely identify a node and how well position-based routing performs over the resulting coordinate system (when compared against routing with real coordinates).

In summary, the contributions of this paper reside in answering two main questions: "How landmark placement affects the accuracy of the resulting hop-count coordinate system?" and "Can landmark placement be avoided altogether?" In answering the first question, our simulation results confirm that placing landmarks on the topology periphery yields more accurate coordinates. The answer to the second question is critical when designing self-organizing networks, since border node selection/placement may be too expensive or even infeasible in some deployments. Our results also show that, in general, landmark placement strategies only have significant performance impact when the number of landmarks is low. In other words, if enough landmarks are used, random landmark placement yields comparative performance to placing landmarks on the boundary (randomly– or equally spaced). We contend that the work here is a first step towards the development of reliable and efficient methods for landmark placement in virtual positioning systems.

The remainder of this paper is organized as follows. In the next section, we describe existing hop-count positioning systems for sensor networks in more detail. The methodology and results of the simulation study we conducted on the impact of landmark placement considering both uniform– and non-uniform topologies are described in Section III. Finally, Section IV presents our concluding remarks and identifies directions for future work.

II HOP-COUNT BASED POSITIONING SYSTEMS

The use of topological or hop-count based localization methods in wireless sensor networks is advantageous because they are simple and do not require additional equipment or devices. To our knowledge, only three algorithms to-date use hop count to build coordinate systems. We describe them below.

GPS-FREE-FREE [14] constructs a two-dimensional coordinate system based on hop-count distances using three landmarks. Landmarks in GPS-FREE-FREE are nodes chosen from the interior of the topology in such a way that they form an equilateral triangle. Each landmark broadcasts a packet in order to allow other nodes to discover their hop-count distance to it. This packet also contains virtual position of the landmark. Thus, each landmark knows its hop-distances to landmarks and their virtual coordinates. Based on this knowledge, and using the hop-distance as a metric, each node calculates its virtual coordinates through trilateration.

VCAP [16] is another hop-count positioning algorithm very similar to GPS-FREE-FREE. VCAP also uses three landmarks at equal distances from each other but instead of a two-dimensional system, VCAP builds a three-dimensional one. In other words, the hop-count distances to the landmarks are directly used as the three coordinates of a node. The advantage of VCAP when compared to GPS-FREE-FREE is that (1) it requires less computation, since the trilateration phase is avoided and (2) it provides better accuracy, since the hop count to the third landmark is used as a real coordinate.

Another difference between GPS-FREE-FREE and VCAP is in how they place the landmarks. While both algorithms form an equilateral triangle with the landmarks, VCAP positions them on the boundary of the topology, while GPS-FREE-FREE place them in the interior.

JUMPS [15] is another positioning system based on hop-distances. The JUMPS algorithm is very similar to VCAP. It also places landmarks on the border of the network at equal distances of one another and uses, as coordinates, hop-count distances to landmarks. JUMPS utilizes, however, up to ten landmarks instead of the three used in VCAP. It has been shown [15] that adding landmarks increases the accuracy of the resulting coordinate system. BVR [5] proposes a routing mechanism that works over a hop-distance based positioning system. As in JUMPS, a BVR node's virtual position is given by its distance to up to ninety landmarks that are randomly placed.

The common point shared by most of these positioning methods (i.e., GPS-FREE-FREE, VCAP, and JUMPS) is that they assume that landmarks can be manually placed at specific locations. For that to happen, either manual deployment (i.e., manually placing nodes when deploying a sensor network) or landmark election mechanisms are required. Many scenarios make manual deployment infeasible (e.g., dropping sensors from a plane in hostile, hard to access regions). In such cases, election algorithms are required to select bor-

number of hops from that node to the landmark.

[3] Besides the simulation results presented here, we have also proven analytically that placing landmarks on the boundary yields more accuracy. We do not include our analytical results in this paper due to space limitations.

der nodes with specific placement. The fact that these algorithms may be prohibitively expensive (as they require additional computational and several rounds of communication among nodes), highlights the importance of avoiding landmark placement and election as in BVR. However, in BVR, they do not explicitly justify the choice of random landmark placement as well as the reason for using larger numbers of landmarks. The results from our work provide an explanation for their design choices.

Motivated by the state-of-the-art in hop-count based positioning systems, this paper aims at evaluating the effect of landmarks placement strategies on the quality of the resulting coordinate system. Our principal goal is to investigate if landmark placement/election can be either simplified by designating as landmarks any border node, or, better, avoided by assigning the role of landmarks to any node in the topology.

III SIMULATION ANALYSIS

For the simulation experiments, we have written our own simulator since existing network simulators work at the packet-level and are too fine-grained for our purpose.

This simulator only 1) places nodes according to the distributions described in Section A, 2) determines hop-distances to landmarks by successive neighborhood discoveries, these hop-distances are then used as coordinates, and 3) discovers paths, based on the hop-count coordinate system, between randomly selected source and destination.

A Parameters

The environment considered is a disc of radius 1,000 meters, and the radio coverage range of the nodes is 60 meters. We assume that nodes are homogeneous, i.e., they all have the same capabilities, and that neighborhood discovery is provided by the MAC layer.[4] We simulated topologies where the number of landmarks ranges from 3 to 10. Thus, we can evaluate the performances of both VCAP and JUMPS. The different landmark placement strategies we use are the following: **Strategy 1** places landmarks on the boundary of the topology, at equal distances from each other. In **Strategy 2**, landmarks are randomly placed on the boundary of the topology. **Strategy 3** randomly places landmarks anywhere in the topology. Their location might be on the boundary or inside the disc area.

The overall number of nodes, including landmarks is set to 3,000, and as previously pointed out, distributed in two kinds of topologies: uniform topologies where nodes are uniformly distributed over the field, and non-uniform topologies where nodes are placed around "concentration points" according to a normal distribution.

We should point out that, unlike the studies conducted in VCAP and JUMPS, we consider the case of disconnected networks. This means that nodes with no direct neighbors may exist. Such nodes can obtain coordinates from a subset of landmarks only, or do not obtain any coordinate at all.

For every scenario (i.e., combination of node distribution, number of landmarks, number of nodes, and landmark placement strategy), every data point we obtain is the average over 50 runs.

B Performance Metrics

Zones: Similarly to VCAP [16], our performance metrics are based on the concept of `zones`. A `zone` is the set of nodes sharing the same virtual coordinates. The `zone size` is thus the maximum Euclidean distance, measured using the real coordinates, between two nodes within the same zone. Thus the `zone size`provides a measure the coordinate system's ambiguity. In other words, the smaller the zone size, the more accurate the coordinates.

In this paper, we consider three zone-related metrics. First, we evaluate, the *average zone size* for each scenario. We also measure the *maximum zone size*, i.e., the largest zone in a scenario. Ideally, the maximum zone size should be smaller than the node's transmission range so that nodes sharing the same coordinates are physically neighbors, and can communicate directly. Finally, we report the *number of nodes per zone*. The lower this number, the more accurate the coordinate system. Ideally, we obtain one node per zone, which means that no coordinate ambiguity exists.

Route computation: Another important criterion we use in our experimental evaluation is how well routing performs over the resulting virtual coordinate system when compared to using real coordinates. To evaluate routing performance, we consider the rate of successfully delivering packets. We run our routing experiments as follows. For every run, we pick 1,000 random source-destination pairs and perform greedy routing, i.e., the next hop decision is solely based on the position of the node and its neighbors. Routing then tries to select as next hop the closest neighbor to the destination. It cannot, however, guarantee success due to local minima situations where no neighbor is closer to the destination than the current node. In such situations, the route computation procedure is considered as failed.

C Results

In this section, we present results from our simulation experiments. Every data point is obtained as the average over fifty simulation runs.[5]

Average Zone Size: Figure 1 shows the average zone size as a function of number of landmarks for the different strategies. We can observe that the shape of the curves is similar irrespective of the strategy, showing that as the number of landmarks increase, the benefits of placing landmarks in the boundary of the topology (equally spaced or randomly) decrease. For this particular experiment, for example, while there are clear performance differences between the three strategies for five or less landmarks, the average zone size does not change significantly when seven landmarks or more landmarks are used even under different placement strategies. This observation remains valid for both uniform and non-uniform topologies.

[4] We simulated a perfect MAC layer, which means that 1) two nodes are neighbors if the distance between them is less than r, the radio coverage range described in Section A, and that 2) there is no packet loss during transmissions.

[5] Because the confidence interval is negligible, compared to the average value, we do not represent it on these figures.

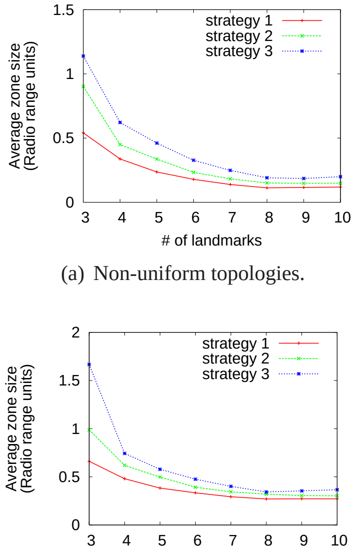

(a) Non-uniform topologies.

(b) Uniform topologies.

Figure 1: Average zone size in radio range units as a function of number of landmarks for different landmark placement strategies.

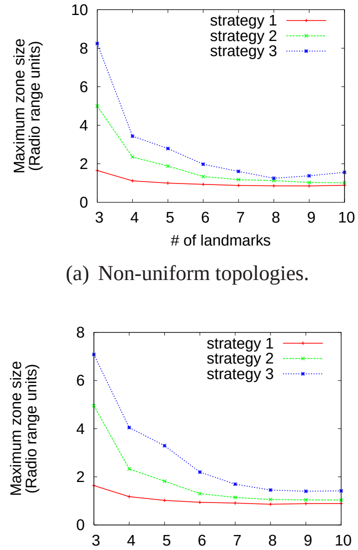

(a) Non-uniform topologies.

(b) Uniform topologies.

Figure 2: Maximum zone size in radio range units as a function of number of landmarks.

Maximum Zone Size: In VCAP, the authors propose to combine position-based and proactive routing. Indeed, VCAP generates zones with size of up to two radio ranges. Therefore, a packet can reach a node 2-hops distant from the intended destination. Adding 2-hop neighborhood knowledge is then required so that, when a node receives a message intended to another node with the same virtual coordinates, it uses proactive routing within the 2-hop neighborhood to forward the packet to its right destination. Thus, the maximum zone size is an important metric, since it determines what kind (and how expensive) of proactive forwarding method to be used in addition to the position-based one.

In Figure 2, we show the maximum zone size (in radio range units) as a function of the number of landmarks and their placement strategy. We observe that, placing landmarks on the boundary improves the worse case, i.e., reduces the average maximum zone size, independent of the number of landmarks. For our simulation scenarios, for example, lower numbers of landmarks randomly placed generate zones of up to eight radio range units. This requires a 8-hop proactive routing protocol, which can be quite expensive in terms of overhead. As before, the difference between landmark placement strategies, however, becomes less significant when topologies are more uniform and/or the number of landmarks increases.

Number of nodes per zone: A single zone for the whole topology is the worst possible case one can obtain – it means that all nodes have the same coordinates. On the other hand, the ideal case is when there are as many zones as nodes. Thus, the lower the number of nodes per zone, the more accurate the coordinate system.

We show in Figure 3 the average number of nodes per zone. We observe that the difference between the strategies becomes less important when the number of landmarks increases. This agrees with the trend shown by Figures 1 and 2.

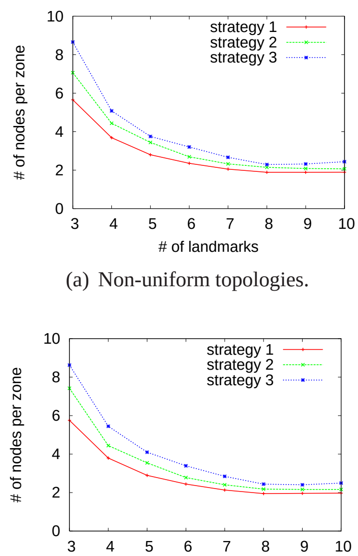

(a) Non-uniform topologies.

(b) Uniform topologies.

Figure 3: Average number of nodes per zone, as a function of number of landmarks.

Route computation: Figure 4 shows that different landmark placement strategies have significant impact on routing performance. We observe that placing landmarks on the boundary yields the best results, especially when they are at equal distances from one another.

This behavior is closely related to the number of nodes per zone represented in Figure 3. Indeed, when a node receives a packet to forward, it chooses, depending on the virtual coordinates, which neighbor is the more appropriate to be the

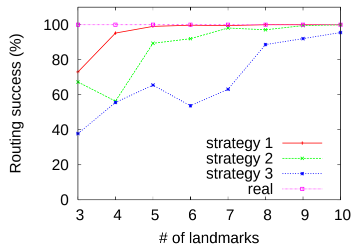

(a) Non-uniform topologies.

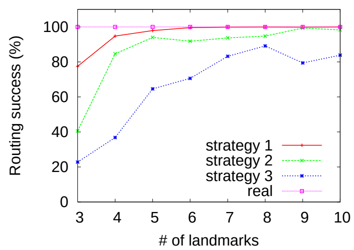

(b) Uniform topologies.

Figure 4: Route computation success rate as a function of number of landmarks.

next hop. If two nodes or more share the same coordinates, the forwarding node chooses one of them randomly. If the average number of nodes among a zone is high, then the probability of choosing the right next hop is lower. Thus, routing is more efficient in scenarios where the average number of nodes sharing the same coordinates is lower.

Routing over coordinates obtained using strategies 1 or 2, however, leads to similar performance when compared to routing over real coordinates, provided that sufficient landmarks are employed. This is an important observation as it shows that strategy 2, i.e., placing landmarks on the periphery, is enough to achieve adequate routing performance, avoiding the need of equally distant landmark placement.

We also notice again that as the number of landmarks increases up to a certain threshold, considerable performance gains are achieved. However, beyond the threshold, the gains are not very significant. For the scenarios we ran, seven landmarks seem to be the threshold for achieving adequate packet delivery.

IV Conclusion

In this paper, we tackle the problem of landmark placement for hop-count based positioning systems. While previous studies choose as landmarks nodes consistently distributed on the periphery of the topology, we show here that such a criterion does not necessarily yield sufficient performance benefits that warrant its cost.

The experimental results we obtained show that, indeed, placing the landmarks on the topology boundary (randomly or equally spaced) improves the performance of the coordinate system when compared to random landmark placement (i.e., anywhere in the topology). We also confirm the results obtained in JUMPS, showing that increasing the number of landmarks increases the accuracy of the underlying coordinate system. However, we go beyond that result and show that, if enough landmarks are used, random landmark placement yields comparative accuracy to placing landmarks on the topology boundary (equally spaced or randomly). This is an important result for energy-constrained network designers, planners, and providers, since boundary placement can be prohibitively resource consuming.

We also evaluate the performance of routing over the resulting topology-based positioning system against routing using real coordinates and show a similar trend, i.e., that the benefits of boundary placement decreases as the number of landmark increases. This trends hold for both uniform and non-uniform topologies. This result also supports the choice of random landmark placement using a large number of landmarks used by BVR [5].

As future work, using the insight gained in this work, we plan to propose mechanisms that dynamically determine the number of landmarks needed to obtain the most accurate coordinate system. This mechanism should also be able to identify, given a certain node distribution, the optimal landmark locations.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Magazine*, vol. 40, no. 8, pp. 102–114, Aug. 2002.

[2] L. Iannone, F. Benbadis, M. D. de Amorim, and S. Fdida, "Some applications of wireless sensor networks," in *Journées Scientifiques Techniques CetMet*, 2004.

[3] M. Mauve, J. Widmer, and H. Hartenstein, "A survey on position-based routing in mobile ad hoc networks," *IEEE Network Magazine*, vol. 15, no. 6, pp. 30–39, Nov. 2001.

[4] S. Giordano, I. Stojmenovic, and L. Blazevie, "Position based routing algorithms for ad hoc networks: a taxonomy," 2001. [Online]. Available: citeseer.ifi.unizh.ch/giordano01position.html

[5] R. Fonseca, S. Ratnasamy, D. Culler, S. Shenker, and I. Stoica, "Beacon vector routing: Scalable point-to-point in wireless sensornets," in *In proceedings of NSDI'05, Boston, MA*, 2005.

[6] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins, *Global Positioning System: Theory and Practice, Fourth Edition*. New York: Springer-Verlag, 1997.

[7] "Glonass," http://www.glonass-center.ru/.

[8] CNES, ESA, "The galileo project," http://www.cnes.fr/html/_455_471_1679_.php.

[9] A. Goldsmith, *Wireless Communications*. Cambridge University Press, 2005. [Online]. Available: http://wsl.stanford.edu/andrea/wireless/

[10] F. Gustafsson and F. Gunnarsson, "Positioning using time-difference of arrival measurements," in *Proceedings of ICASSP*, Hongkong, Apr 2003.

[11] D. Niculescu and B. Nath, "Ad hoc positioning system APS using AoA," in *Proc. of IEEE INFOCOM'03*, San Francisco, CA, USA, 2003.

[12] A. Rao, S. Ratnasamy, C. Papadimitriou, S. Shenker, and I. Stoica, "Geographic routing without location information." in *Proc. of ACM MOBICOM'03*, San Diego, CA, USA, 2003.

[13] A. Jadbabaie, "On geographic routing without location information," in *Proceedings of the IEEE Conference on Decision and Control*, Paradise Island, Bahamas, 2004, pp. 4764–4769.

[14] F. Benbadis, T. Friedman, M. Dias de Amorim, and S. Fdida, "GPS-Free-Free positioning system for sensor networks," in *Proc. of WOCN 2005 Conference*, Dubai, UAE, 2005.

[15] F. Benbadis, J.-J. Puig, T. Friedman, M. D. de Amorim, D. Symplot-Ryl, and S. Fdida, "JumPS: Enhancing hop-count positioning in sensor networks using multiple coordinates," LIP6, Tech. Rep., 2006, http://rp.lip6.fr/ benbadis/jumps/.

[16] A. Caruso, S. Chessa, S. De, and A. Urpi, "GPS-Free Coordinate Assignment and Routing in Wireless Sensor Networks," in *Proc. of IEEE INFOCOM'05*, Miami, FL, USA, 2005.