Research Repository UCD

Title	Benchmarking latency effects on mobility tracking in WSN
Authors(s)	Tynan, Richard, O'Grady, Michael J., O'Hare, G. M. P. (Greg M. P.), Muldoon, Conor
Publication date	2009-05
Publication information	Tynan, Richard, Michael J. O'Grady, G. M. P. (Greg M. P.) O'Hare, and Conor Muldoon. "Benchmarking Latency Effects on Mobility Tracking in WSN." IEEE, 2009.
Conference details	Paper presented at the Second International Workshop on Applications of Ad hoc and Sensor Networks (AASNET), held at the 23rd IEEE International Conference on Advanced Information Networking and Applications AINA 2009, 26-29 May 2009, Bradford, United Kingdom
Publisher	IEEE
Item record/more information	http://hdl.handle.net/10197/1250
Publisher's version (DOI)	10.1109/WAINA.2009.120

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Benchmarking Latency Effects on Mobility Tracking in WSNs

Richard Tynan M.J. O'Grady G.M.P. O'Hare Conor Muldoon

CLARITY: The Centre for Sensor Web Technologies

University College Dublin, Belfield, Dublin 4, Ireland

{richard.tynan, michael.j.ograde, gregory.ohare, conor.muldoon}@ucd.ie

Abstract

The number of active nodes in a WSN deployment governs both the longevity of the network and the accuracy of applications using the network's data. As node hibernation techniques become more sophisticated, it is important that an accurate evaluation methodology is employed to ensure fair comparisons across different techniques. Examining both energy and accuracy ensures a claim of increased longevity for a particular technique can be contrasted against its associated drop, if any, in application accuracy. This change can also be as a result of increased latency and the accuracy encapsulates many aspects of WSN performance in one metric. In this work, we detail the first in a series of experiments designed to demonstrate the trade-offs for a WSN and we employ mobility tracking as the application to benchmark accuracy. Additionally, we demonstrate experimental evidence for a potential adaptive mobility tracking protocol.

1 Introduction

For a typical WSN, the accuracy of an application and the longevity of the network will be inversely proportional to each other. This is due to the finite power reserves of the nodes and the desire for applications to have large volumes of fresh data to perform their calculations. A number of techniques have been proposed that can opportunistically hibernate sensors, such as CCP [9], but no work to date has focussed on a complete analysis of such approaches from both the accuracy and longevity perspectives. For instance, consider a claim that a particular hibernation technique can double the life of a WSN. Will its error double also?

In this work, we detail the first of a suite of experiments designed to analyse the performance of a WSN when nodes are hibernated. This first benchmark demonstrates the tradeoffs that exist for a WSN

where no node hibernation's occur in the network. In this case the WSN topology will be static and serves as a baseline in WSN performance for future experimentation. Hibernation techniques can then be evaluated by how much they increase longevity or decrease accuracy from this baseline. Various parameters of the individual techniques can be tweaked in order to achieve better accuracy or longevity. Crucially, by examining both QoS metrics neither one can be artificially increased to distort performance. The relationships identified here are not unique to mobility tracking applications, many potential WSN applications requiring a balance between energy consumption, density, latency and accuracy may also be able to exploit the results and trends identified here.

In addition, this work also demonstrates experimental evidence for the operation of an adaptive mobility tracking application. We examine the localisation accuracy of a typical WSN configuration with varying degrees of message latency, target speed and application deadline selection. When values for some of the parameters can be estimated i.e. target speed, it is envisaged that an adaptive protocol may be able to tune other parameters to optimise accuracy.

In the next section we look at mobility tracking applications that are used in our experimental evaluation. Following this, the system architecture is detailed, including the protocol stack used on each node. The experimental setup is then provided in section 4, with the results given in section 5. We close with conclusions drawn from this experimentation as well as a discussion of how this work can form the basis for benchmarking future WSN performance under the EDLA [4][6] trade-offs in the presence of hibernating nodes.

2 Target Localisation

The task of target localisation, is to transform the streams of sensed data from the WSN into co-ordinates that pinpoint the location of a target in the sensed area

[1]. Two basic target localisation techniques are chosen for the application in this work, since they specifically do not require any prior characterisation of the target, making them generally applicable for many environments and many targets. They are the Weighted Average Localisation (WL) and the Maximum Signal Strength Localisation (ML). For the Weighted Average technique each sensor that is active will be able to sample the signal at its location and the larger the value, the greater its influence will be on the estimated location.

$$x_{target} = \frac{\sum_{i=1}^N signal_i^2 x_i}{\sum_{i=1}^N signal_i^2} \quad (1)$$

A similar equation is applied to the y co-ordinate to locate the source in the 2-dimensional area, and this technique has been adopted previously in [4] and [3]. The second localisation technique adopted here is a simple, but effective method, which assigns the location of the maximum signal value sensed at an active sensor to the location of the target and has been used in [3]. Two techniques are applied here so that a broader sense of how the application performance is affected by latency, deadline and target speed can be presented.

In standard target localisation application [3], the nodes sample their sensors and forward their data to the base station either at a given time, according to a schedule or in response to a command. The base station can wait for all the data to reach it before calculating the targets position, however, in that time the target will moved a certain distance. The longer the time it takes for all the messages to reach the base-station, the greater the distance travelled by the target, leading to an increase in localisation error. For this reason, we analyse the effect of message latency on localisation accuracy.

In many cases all the data will not reach the base-station, due to failed or exhausted nodes, in which case a deadline must be chosen when the available data is used to make a decision. Selecting the appropriate deadline by which the decision is made, means messages received afterwards are discarded. This approach, in effect, limits the density of messages observed by the base-station. This deadline can be determined experimentally on a case by case basis or it may be possible to derive this value from a theoretical analysis of various hardware and software parameters of the WSN.

3 System Architecture

In order to deliver the aforementioned target tracking application we adopted a standard protocol stack,

whereby equivalent layers communicate with each other on neighbouring nodes, figure 1 (a). When multiple nodes wish to communicate, they cannot do so at the same time due to interference on the channel, so a MAC layer is required in order to mediate the use of the channel and to retransmit failed packets. As such, the first layer on the WSN device for this system architecture will be the MAC layer, with direct control over the transceiver. For WSNs numerous approaches to this have been developed, including B-MAC [5], but we have opted for the in built 802.11 implementation provided with J-Sim [7]. J-Sim is a port of the successful NS-2 simulation environment to Java. It provides many of the protocols required to assemble a complete WSN application and is the simulation environment used for our experimentation. A MAC layer will typically use an RTS/CTS mechanism to manage communication between nodes but these control messages can be lost due to collisions in the channel, so perfect reliability is rarely achieved for an ad-hoc network in practice.

Figure 1. (a) Protocol stack used in experimentation, where NT represents the nodes neighbour table. (b) Node and environment configurations.

The next component of the stack provides the multi-hop communication for nodes out of direct transmission range of the base station. Greedy Perimeter Stateless Routing (GPSR) [2] is a multi-sink protocol, which uses the geographic location of the source, intermediate forwarding nodes and the destination to route the packet. The GPSR protocol [2], provided with J-Sim, will deliver the required forwarding for our experimental purposes. While the choice of routing protocols can impact the performance of the WSN in terms of latency, we are not examining the latency characteristics of individual protocols and we leave an analysis of the impact of other protocols on our results to future work. The next component is the application resident on the nodes. This samples the sensed data at the node and relays it to the base station every ten seconds, for this set of experiments. A number of alternative configurations could harvest the data from the network, for example, the base station could flood the network with a command packet. Active nodes will respond

with their data through the multi-hop topology. Another possibility could be for the base station to send a unicast message to specific nodes and only they reply. These variations are not considered here however.

At the base station a corresponding application layer receives data and calculates the location of the target based on the sensing information. With the fixed density deployed, the base station must use a timeout in order to balance the message latency with the number of readings received, and so a timer is started every ten seconds. After the timer expires, data which has reached the base station at that point is used to evaluate the location of the target in the environment. The longer the timer, the more data for the calculation, but the greater the subsequent distance the target will have travelled, potentially increasing error. For a density maintenance technique, increasing the number of active nodes will increase contention for the channel and therefore increase latency. The two protocols used to decide on the targets' location, ML and WL, operate with identical data, i.e. the time for the first protocol executing does not impact the timeout used for the second.

4 Experimental Setup

The experimental methodology adopted in this work is designed to illustrate the relationships between latency, target speed, optimal timeout and tracking accuracy. The density of the nodes remains fixed, to remove its effect on performance and to focus attention specifically on latency. The message latency is varied by reducing the transmission radius of the nodes, which in turn causes messages to traverse a longer path to the base-station. Various different target speeds are also used to demonstrate how choosing the correct balance between density and latency, through careful selection of the application timeout, is vital in determining the performance of the WSN.

The simulated area for this set of experiments is defined as 100 meters x 100 meters with a deployed node density of one node every 6m, figure 1 (b). The result of this is that a fixed density of 256 nodes are used to cover the region of interest. One of the primary reasons for selecting this setup is allow the results to be generalised to large areas by concatenation of networks similar to this. For example, a 500m x 500m region could be configured using 25 instances of the setup used here in a 5 x 5 grid formation.

The target in the environment is allocated a magnitude of 1000 units and signal received at each sensor decays according to the inverse square law of distance to the target. This model is applicable in many instances,

including thermal radiation, light, sound and magnetic and gravitation fields, and has been used previously for similar experiments in [3]. It is initially located in the centre of the sensed area and takes a random walk around the area at the specified speed. It is assumed that no prior information is available about the targets' characteristics, however, for these experiments its maximum speed is limited to 5 m/s or 18 km/h. Under this setup, all of the nodes remain active and no hibernation of any nodes takes place.

5 Results

The goal of the first of these experiments, is to demonstrate the effect of hop count and latency on the competence of a WSN, measured through application performance. The localisation accuracy of the target tracking application is measured for various possible target speeds. Following this, the influence of target speed is removed by averaging the results for all possible speeds, which produces an average QoS metric. This also focuses attention on the influence hop count has on latency, which in turn affects localisation accuracy. The results are finally aggregated to demonstrate changes in accuracy and optimal timeout selection as the transmission radius decreases. The results are obtained through the averaging of five individual executions of the simulation.

5.1 Target Speed, Hop Count and Latency

In order to investigate the effect of latency, hop count and target speed on WSN performance, a series of experiments were derived such that the transmission range of each node becomes progressively smaller, starting at 150m, decreasing to 75m and finally to 25m. This restricts the nodes ability to communicate with the base station directly and necessitates the use of multiple hops for the messages to be delivered. The more constrained the range, the more hops are required for messages to traverse, which increases latency. In this first experiment, the transmission range for each node is set to 150m, so that every node is within range of the base station. This removes the requirement and latency involved in multi-hopping messages. The characteristic EDLA curves can be seen in figure 2 for both the ML and WL techniques. As the timeout increases, more data reaches the base station, thus increasing the accuracy of the localisation. After a critical point, which is dependent on the speed of the target, the effect of latency outweighs the increased density and the accuracy decreases.

Figure 2. Error for ML and WL - transmission radius for each node is 150m.

A number of interesting points can be observed here, firstly the minimisation point of the ML error is dependent on the targets' speed, for 1 m/s the minimisation point occurs at 0.4 seconds, however, when the target moves at 5 m/s the optimal point shifts to the earlier deadline of 0.3 seconds. From this it can be concluded that latency becomes more important as the speed of the target increases. This conclusion is reinforced as the graphs continue past the optimisation point, the gradient of the error for faster moving targets is steeper than that of slower moving targets. We will return to this important result subsequently in this paper.

Similar trends are observed for the WL technique, figure 2, however the difference in accuracy between the WL and ML approach is clearly visible from the graphs, with the WL technique outperforming ML at all locations. An interesting observation can be made in the characterisation of the WL and ML graphs. Firstly the WL curves are flatter before the optimal point indicating that the WL approach copes better with less data. Secondly more divergence of the graphs can be seen after the optimal point. This implies that in the earlier portion of the graphs the effect of density is magnified by the targets' speed and this effect continues after the optimal point is passed. The conclusion to be drawn

Figure 3. Error for ML and WL - transmission radius for each node is 75m.

from this is that latency and target speed are more significant for the WL technique than for the ML approach. This is supported by the new optimal timeout location, which is 0.3 seconds for WL, regardless of speed.

The transmission radius is now reduced to 75m for the nodes, which essentially means that most of the area will remain within direct range of the base station. A small portion of the network, however, will have to multi-hop data through a single hop since the dimensions of the area are 100m x 100m. Consequently latency will be increased slightly as some packets will need to be transmitted twice.

The accuracy of the ML approach for this decreased transmission radius is depicted in figure 3. With the increased latency, the effect of target velocity is magnified, as can be seen by the greater separation between the curves. The range of optimal timeout values has now increased, with a value of 0.4 seconds suitable for a target of 5 m/s but a value of 0.6 seconds should be selected for the slowest target. The increased latency has also resulted in a consistent increase in the localisation error, even at the optimal timeout values for the different target speeds. The added latency, due to such multi-hopping, not only widens the gap of the

Figure 4. Error for ML and WL - transmission radius for each node is 25m.

optimal values but also shifts them to the right. This implies that the slight increase in latency has little impact on the optimal value for fast targets but has a significant effect on the optimal timeout for slower targets. This supports the conclusion that faster moving targets favour fresh data over large volumes of old data.

As seen previously, in figure 2, the effects of this latency are magnified for the WL approach, figure 3. The relatively large separation between the curves again demonstrates the sensitivity of this approach to target speed. In a similar fashion to the ML technique, the optimal timeout values have once again shifted to the right, however now there is a clear distinction between optimal timeouts for differing target speeds under the WL technique. Once more, the WL approach outperforms the ML technique, but in relation to the previous radius, figure 2, the error is consistently higher across all values. This demonstrates the effect that increasing latency, through hop count, can have on application accuracy and optimal timeout selection.

The transmission radius, initially set at 150m and subsequently reduced to 75m, is now decreased further to 25m for the final experiment in this section, which produces a maximum hop count of six for the most distant sensors from the base station. The effect on

Figure 5. Average localisation error for varying target speed when the transmission radius is 150m, 75m and 25m.

the accuracy of the ML technique can be seen in figure 4. From this graph the interplay between latency and target speed can once again be seen clearly, in fact the effect is so great that the error for the slower moving sources has not begun to level off even after 1 second. The faster moving target, however, has an optimal timeout value of 0.6 seconds, thus quite a large gap now exists for the different target speeds, thus latency magnifies the effect of target speed, for this application. The increased latency here has introduced another possible QoS metric, how quickly can the optimal accuracy of the application be reached, and we will examine this later in this section.

The WL approach experiences similar trends, in terms of increased separation of optimal timeouts, based on target speed and increased error due to the added latency associated with multi-hopping. Contrasting the ML and WL approaches for this hop count, the WL approach again performs considerably better. In addition, for slower moving targets, waiting longer for more data does not improve the accuracy of the WL technique to the same extent it does in the ML technique. This can be seen in the gradients of the error curves in figure 4. This would suggest that in some cases, not only is target speed relevant for selecting an appropriate timeout value, but possibly the localisation technique to be used should be considered also. Such factors are implicitly encoded in the EDLA trade-offs through the accuracy metric.

The results here would suggest that the selection of a timeout based on the target speed could improve the performance of the localisation considerably, particularly given the effect the added latency multi-hopping has on WSN competency. Potentially, if the target speed could be estimated, then an adaptive protocol could vary the timeout to an appropriate value based

	150	75	25
ML	2.5	2.8	3.3
WL	1.6	1.9	2.7
timeout	0.3	0.5	0.8

Figure 6. Optimal timeout positions.

on the approximated speed and localisation technique in use. Such an adaptive protocol is part of our ongoing work and we aim to benchmark the adaptive protocol's performance with that outlined here. However, in this work, it is assumed that no such prior information about the target is known. With this in mind the results from the previous experiments are now averaged for the different target speeds, in order to focus on the effect the increased hop count has on the timeout to be selected and the optimal application accuracy.

5.2 Hop Count and Latency

Averaging the error values of figure 2, produces the graph in figure 5. Here the characteristic U-shaped EDLA tradeoff curve can be clearly seen. In addition, the consistently superior performance of the WL technique is also visible for all timeout values. In fact, if using the WL technique to localise the target, the solution obtained after 0.2 seconds will be more accurate than the optimal solution of the ML technique achieved after 0.3 seconds. In real time applications, where decision deadlines must be met, opportunistically switching the algorithm in use may prove beneficial.

The effect of the additional latency introduced through the reduction in transmission range from 150m to 75m, can be seen in this figure also. Firstly, for the equivalent timeout, 0.3 - 0.5 seconds, the graphs are sloped in opposite directions; the trends of 150m are decreasing in accuracy for this domain, whereas the 75m trends are increasing. The minimisation point has also shifted to the right and the errors have increased, similar to the previous results that also distinguish target speeds. The timeout domain is shifted to the right in order to focus on the area of the graph where the trade-off exists, this approach must also be followed as the radius decreases further to 25m, illustrating the sensitivity of the optimal point to the latency of messages. The trends of the previous two radii are also repeated when the range is set to 25m; accuracies are decreased, optimal timeout has increased and the WL technique consistently outperforms the ML approach.

Combining the previous three graphs yields the table in figure 6, where the changes in optimal timeout and accuracy for both localisation techniques as the hop count increases, are presented. This summarises the previous experimentation concisely, in that the trends of increasing error with latency, WL outperforming ML

Figure 7. Optimal localisation errors as target speed increases.

and the optimal timeout increasing as hop count does are all present. One interesting point is that when the effect of target speed is removed, the optimal time out for both ML and WL coincide at the same location. A final trend to be observed, is the relatively large jump in the WL error as the transmission range reduces from 75m to 25m. A corresponding increase in the ML approach is not observed, indicating the possibility that the performance of the WL technique would degrade considerably with a further increase in hop count and latency, and there may come a point that the ML and WL techniques perform the same for a high latency.

5.3 Hop Count and Target Speed

In this section, we examine the optimal points of the preceding graphs to illustrate the precise effect latency has on WSN accuracy. Firstly, from figure 7 it is clear that as target speed increases, localisation error also increases. Secondly, the added latency of the additional hops decreases the accuracy of the system for both of the techniques ML and WL. Finally, the target speed has a greater impact on the accuracy of the system as the hop count and subsequent latency increases. With regard to the temporal performance of the system for different hop counts and target speeds, it is clear that faster targets favour shorter timeouts. From the results obtained here, the optimal timeout increases as the number of hops does, due to the additional latency of messages reaching the base station.

6 Conclusions

In order to characterise the performance of a WSN, two primary metrics are typically of interest - longevity and accuracy. The range of experiments outlined here, evaluates a WSN through the performance of a target tracking application, which is particularly suited

to the task due to the continual motion of the target while messages are en-route to the base station. Specifically, the effect of latency, introduced through multi-hop communication, on the ability of the application to carry out its task is examined. The trade-off between latency and density is managed through the selection of an appropriate application timeout in order to achieve optimal system performance. Such tradeoffs are characteristic of an entire class of WSN applications that must receive an appropriate amount of data by a certain deadline in order to function correctly. A number of interesting results have been obtained, including the fact that latency reduces WSN accuracy considerably, shifts the optimal position to a later value and that target speed is a factor in determining the optimal location of the timeout.

In keeping with the EDLA tradeoffs, the average node lifetime in our simulations is 120 seconds on 1000 Joules of energy. We have presented the baseline in WSN performance under the accuracy and longevity metrics. Experiments conducted using such metrics would mean that neither can be altered without impacting the other. When node hibernation techniques are introduced, we are now in a position to contrast their performance with this benchmark to see what degradation in application accuracy results from the increase in longevity. Many of these node hibernation techniques require additional, redundant nodes to be deployed and our experimental approach caters for this also. If more nodes are introduced into the deployment, latency will increase due to contention for the channel. This, as is shown here, decreases the accuracy of the tracking application and so a fair comparison will result.

We propose to develop a framework for such an evaluation of power management protocols for WSNs e.g. CCP [9] or those based on interpolation [8]. This would allow new techniques to be plugged into the framework and evaluated according to the longevity and accuracy metrics. Additionally, this work also demonstrated the effect both message latency and target speed has on the selection of the optimal timeout value for the localisation. If these values can be estimated for instance using timestamping or a crude velocity calculation, then the base station deadline could be tuned in order to deliver optimal accuracy across many latency values and target speeds. This is part of our ongoing experimentation and evaluation.

References

[1] Richard R. Brooks, Parameswaran Ramanathan, and Akbar M. Sayeed. Distributed target classifi-

cation and tracking in sensor networks. *Proceedings of IEEE*, 91(8):1163 – 1171, August 2003.

- [2] Brad Karp and H. T. Kung. Gpsr: greedy perimeter stateless routing for wireless networks. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 243–254, New York, NY, USA, 2000. ACM Press.
- [3] Jeongkeun Lee, Kideok Cho, Seungjae Lee, Taekyoung Kwon, and Yanghee Choi. Distributed and energy-efficient target localization and tracking in wireless sensor networks. *Computer Communications*, 29(13-14):2494–2505, 2006.
- [4] Sundeep Patten, Sameera Poduri, and Bhaskar Krishnamachari. Energy-quality tradeoffs for target tracking in wireless sensor networks. In *IPSN*, pages 32–46, 2003.
- [5] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 95–107, New York, NY, USA, 2004. ACM Press.
- [6] Curt Schurgers, Vlasios Tsiatsis, Saurabh Ganerival, and Mani Srivastava. Optimizing sensor networks in the energy-latency-density design space. *IEEE Transactions on Mobile Computing*, 1(1):70–80, 2002.
- [7] Ahmed Sobeih, Wei-Peng Chen, Jennifer C. Hou, Lu-Chuan Kung, Ning Li, Hyuk Lim, Hung-Ying Tyan, and Honghai Zhang. J-sim: A simulation environment for wireless sensor networks. *anss*, 00:175–187, 2005.
- [8] Richard Tynan, G. M. P. O'Hare, and Antonio Ruzzelli. Signal based node activation in wireless sensor networks. In *AVSS '06: Proceedings of the IEEE International Conference on Video and Signal Based Surveillance*, page 47, Sydney, Australia, 2006. IEEE Computer Society.
- [9] Guoliang Xing, Xiaorui Wang, Yuanfang Zhang, Chenyang Lu, Robert Pless, and Christopher Gill. Integrated coverage and connectivity configuration for energy conservation in sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 1(1):36–72, August 2005.