Research Repository UCD

Title	Evaluating interpolation-based power management
Authors(s)	Tynan, Richard, O'Hare, G. M. P. (Greg M. P.)
Publication date	2008-12
Publication information	Tynan, Richard, and G. M. P. (Greg M. P.) O'Hare. "Evaluating Interpolation-Based Power Management." IEEE Press, 2008.
Conference details	International Workshop on Adaptation in Wireless Sensor Networks (AWSN-08), held in conjunction with the 2008 IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA-08), December 10th -12th , 2008, Sydney
Publisher	IEEE Press
Item record/more information	http://hdl.handle.net/10197/1199
Publisher's version (DOI)	10.1109/ISPA.2008.71

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

Evaluating Interpolation-based Power Management

Richard Tynan & G.M.P. O'Hare

CLARITY: The Centre for Sensor Web Technologies

School of Computer Science and Informatics

University College Dublin, Belfield, Dublin 4, Ireland

Email: richard.tynan@ucd.ie gregory.ohare@ucd.ie

Abstract—Power management for WSNs can take many forms, from adaptively tuning the power consumption of some of the components of a node to hibernating it completely. In the later case, the competence of the WSN must not be compromised. In general, the competence of a WSN is its ability to perform its function in an accurate and timely fashion. These two, related, Quality of Service (QoS) metrics are primarily affected by the density and latency of data from the environment, respectively. Without adequate density, interesting events may not be adequately observed or missed completely by the application, while stale data could result in event detection occurring too late. In opposition to this is the fact that the energy consumed by the network is related to the number of active nodes in the deployment. Therefore, given that the nodes have finite power resources, a trade-off exists between the longevity and QoS provided by the network and it is crucial that both aspects are considered when evaluating a power management protocol. In this paper, a novel node hibernation technique based on interpolated sensor readings is analysed according to these four metrics: *energy consumption*, *density*, *message latency* and the *accuracy* of an application utilising the data from the WSN. A comparison with a standard WSN that does not engage in power management is also presented, in order to show the overhead in the protocols operation.

I. INTRODUCTION

The single biggest factor in determining the performance of a WSN is the number of nodes in the deployment, as the density of nodes increases the *aggregate* energy consumption also increases. When considering a WSN using a density maintenance scheme, such as CCP [16], the number of *active* nodes ultimately govern both the longevity and competence of the network. As the active node density increases, more power is consumed on average by the nodes, and thus the longevity of the network is reduced. On the other hand, the increased node activity results in a higher resolution of the area to be observed which, theoretically, should lead to an increase in the competence of the network and therefore improve the accuracy of the application using the data.

This increase however, is limited by a secondary constraint caused by the upper bound on data communication rates, which contributes to the latency of messages reaching the data sinks or decision engines. When too few nodes are active, entire subsets of the network may become disconnected, resulting in blind spots within the region of interest. As the number of nodes increases, bottlenecks in the sparse topology reduce and multiple message paths to the data sinks can emerge, which ultimately reduces latency. After an optimal density is reached, however, further increases result in more contention for the

channel and the latency increases once again. The challenge, therefore, is to select a *density*, which produces optimal *energy* consumption, *message latency* and application *accuracy*. This gives rise to the Energy-Density-Latency-Accuracy (EDLA) trade-offs [9] [8], and it is within this evaluation framework that the performance of an Interpolation-based Redundancy Identification and Sensor Hibernation (IRISH) [13] [12] technique is considered in this paper.

IRISH differs significantly from traditional coverage techniques in that its decisions about redundancy and hibernation are based on sensed data as opposed to node density [16]. Existing techniques tend to maintain a uniform node density as sensors expire, but this fails to take into consideration the fact that more sensed data may be required from a specific area of the environment where interesting events are occurring. IRISH is an attempt to solve this in a general by allowing the network to hibernate a node according to what is being perceived in the environment. By specifying the interpolation error instead of the sensing radius for example, the node activity can be altered to balance the EDLA trade-offs.

In the next section, these EDLA relationships are examined in greater detail. This is followed in section III by details of the IRISH power management protocol and the experimental setup is provided in section IV. The results of this experimentation are presented in section V and the paper closes with conclusions drawn from this experimentation as well as a discussion of how this work can be improved in the future.

II. ENERGY-DENSITY-LATENCY-ACCURACY TRADE-OFFS

The EDLA trade-offs refer to the four distinct performance characteristics *Energy*, *Density*, *Latency* and *Accuracy* of a WSN. As more *Energy* is consumed by the nodes, the operational lifetime of the network reduces. The *Density* refers to the number of nodes deployed in a given region to be monitored. When considering a WSN using a node hibernation technique, such as CCP [16], the density is defined by the number of *active* nodes in the area. *Latency* of messages refers to the delay between the time a message is sent and the time it is received at its destination. Finally, *Accuracy* refers to the precision that an application can perform its task, using the WSN and its data.

It is crucial that all of these attributes are considered when evaluating power management techniques for WSNs. Consider the case where a node will last 1 day on average and a naive hibernation scheme that keeps only 1 node active for

a day and when it expires another node activates. Clearly, the network will last for n days, where n is the number of nodes in the network. However, such longevity will probably lead to a considerable reduction in an applications ability to make accurate inferences about the environment because its data only ever comes from a single sensor. This illustrates the requirement for using all of the EDLA trade-offs in analysing a hibernation technique and each one of these is now examined in further detail.

A. Energy

The energy consumed by a node is due to the activity of a number of its components, such as transmission of data via the transceiver, executing instructions using the processor, storing program state in memory and sampling the sensors. Therefore, the amount of time a given node is active, the shorter its lifetime will be due to its finite power supply. This implies that energy and node lifetime are inversely proportional to each other. In this paper, both node lifetime and activity are utilised for the evaluation of the IRISH protocol. Hibernating a node consists of switching some, or all [10], of its hardware components into a low power sleep mode, where it consumes a fraction of the energy and thus can operate for an increased period of time. For a network as a whole, therefore, the energy consumed will be proportional to the number of active nodes or node density. In addition to the energy saved at the hibernating node, further energy is conserved along the path its packets are forwarded to the base station, through the absence of its data. The drawback of this is that while in hibernation, the nodes sensory data is not available, resulting in a blind spot in the network. In addition, the node is not available to forward messages from its neighbours while hibernating and so an alternate path must be found.

B. Latency

For a WSN, a fixed bandwidth is available for the transfer of data between entities. As the number of nodes increases, more paths become available to route packets to their destination simultaneously, which in theory increases performance by reducing latency. Crucially, the routing component must be able to take advantage of these additional routes by continuously attempting to improve its QoS. There will come a subsequent point, however, where the latency of transmission begins to increase proportionally to the number of nodes. This is due to additional contention for the wireless channel as the node density increases. Such a trade-off has been observed and preliminary experiments have verified such a relationship [3]. This implies that adding some nodes to the deployment can result in improved performance, in terms of latency, however, with too many nodes the network becomes saturated. In addition to contention for the channel, an increased number of nodes increases the probability for both collisions and also instances of the hidden terminal problem for local broadcasts, which together can not only increase latency through retransmissions but can also lead to lost packets [3]. In this work, the latency effect of the IRISH protocol is demonstrated by observing the % of messages received at the base station sent from active nodes for a given timeout duration.

C. Accuracy

The accuracy of an application is directly related to the networks ability to provide timely delivery of a sufficient density of data to it. When this is the case, it is assumed that the application will be able to perform its task to an appropriate standard, and therefore, accuracy is directly related to the competence of the WSN. A deficiency in either density or latency will result in sub-optimal application performance, through either inaccurate inferences about the state of the environment or perhaps excessive delays in decision making. Application accuracy may not degrade uniformly with the competence of the WSN i.e. sharp declines or increases may result from slight variations in density or latency and so it may be necessary to finely tune the parameters of the network in order to achieve the desired accuracy.

It is crucial that in selecting an application to benchmark a networks competence that it requires both aspects of WSN competence to operate effectively. To examine this in more detail, consider a target tracking application that uses the sensed data to localise a target in the environment [6]. With excessive *latency*, the target will have moved a greater distance by the time sufficient data has reached the base station, and thus contributing to inaccuracies in the applications attempt to pinpoint its location. Without adequate *density* of information about the target, it may not be possible to localise the source accurately, as it may be in an area that is not close enough to a sensor for it to be detected. This is symptomatic of many applications where decisions need to be made at a certain time in order to strike a balance between density and latency for optimal performance, so simply increasing density will not typically achieve the desired increase in accuracy. In this work, target tracking will benchmark the accuracy of both the standard WSN and the network using the IRISH protocol to hibernate redundant nodes.

D. Density

From the previous three discussions, it is apparent that the single biggest factor in determining the performance of a WSN is the number of active nodes in the deployment. Ultimately, density drives the power consumption, network longevity, latency and resolution of data perceived by the application. However, the overarching responsibility of the WSN is to ensure the performance of the application. Therefore, the density must be chosen with a corresponding accuracy in mind. This density must also be able to sustain the operation of the network for the desired lifetime. Therefore, the challenge in selecting an appropriate density is to ensure that both accuracy and longevity are maximised in tandem. Within this paper, the interpolation error of IRISH will govern the density of the active network of nodes and this is then evaluated according to the metrics previously outlined. One additional density factor has been mentioned in the previous section. For a practical application a timeout must be selected at which point the forces of density and latency balance. Therefore, there are two means to control the density of messages, the decision timeout and the active node density. For comparison purposes, these

metrics are evaluated for a standard WSN of fixed density and a WSN using the IRISH protocol of variable density.

III. IRISH PROTOCOL

Interpolation is a technique used to predict the value of an unknown function, at a specified location, using known values of that function. From a WSN perspective, the function can be either the temporal or spatial variation in the sensed data. The known points are at the sensors locations and at the instant samples are taken. The Interpolation-based Redundancy Identification and Sensor Hibernation (IRISH) uses the predicted value and the actual sensed value to decide on a nodes redundancy. Essentially, when the difference between the interpolated and the actual values are less than a given threshold then the node is deemed redundant and can be hibernated [13] [12].

A. Interpolation Domain and Model

In order to successfully interpolate a sensed value, an interpolation domain and model must be selected. The interpolation domain is a subset of the network that is used to predict the sensed value at the desired location. A number of such partitions have been identified for WSNs: k Nearest Neighbours, Neighbours within a given radius and also a predefined set [15]. Similarly, there are many interpolation models available for WSNs: Kriging, B-Splines and Weighted Averages [5]. In general, the choice of interpolation domain and model will largely be dependent on the environment and sensory modality respectively. Obstacles in the environment also have to be considered as these can dramatically affect the accuracy of interpolated values.

B. State Transitions

The behaviour of the IRISH approach can be summarised as a set of state transitions within the nodes, figure 1. A node begins by beaconing a request to its neighbours for their sensed data. After a timeout period, the responses are aggregated on the requesting node, to calculate the interpolated value for this node. This is then compared to the actual sensed reading and if the node is redundant it requests to hibernate, otherwise it activates. The reason for the protocol requesting to hibernate is that the IRISH layer may not be able to hibernate if other layers such as the connectivity maintenance layer deem the node critical and this will be discussed further in section IV. After a specified period in the sleep or active states the node enters the listen state and the process repeats.

Fig. 1. State transition diagram for the IRISH protocol.

The neighbour table information maintained by the node is populated in the listen state. It is then consulted to see if the node is sensing redundant or not. Crucially under this approach, as distinct from density based approaches, it is not an option to assume neighbour information is the same before and after a sleep period since under the IRISH regime the neighbour information will include sensed data. Other techniques may or may not purge their neighbour data before sleeping, under the assumption that the topology is relatively static and the cost in repopulating the table would not justify the potential performance degradation by using stale neighbour information for decisions. It is assumed that some form of permanent storage exists in this case on the node so that the data will survive the power down process.

1) *Request Beaconing*: The request beaconing process of the the IRISH protocol is similar in many respects to the beaconing in CCP. Firstly, beacons are conducted using local broadcast, however, not all receiving nodes may respond given a specific interpolation domain. Secondly, the beacons may contain location information of the requesting node so that neighbours could compute their distance to the node. This can reduce the computation required of the requesting node. The IRISH beaconing, however, is significantly different from that of a traditional density maintenance approach such as CCP. Firstly, the beacons serve as a request for data rather than simply informing neighbours of the node's existence. The beaconing process drives the redundancy calculation on the node using the IRISH protocol, however the beaconing of a node under CCP typically occurs after a node has decided to become or remain active. An additional distinction between the two approaches is the purpose of the beacons, within IRISH the beacons serve to identify the redundancy of the requesting node, while under CCP beacons are utilised by neighbours to calculate their own redundancy.

2) *Timing Considerations*: Timeout selection for the IRISH protocol is significantly more complex than for CCP. Under CCP, a node can only be redundant due to variations in topology i.e. a node will only become critical due to the failure or hibernation of a neighbour. While this is true to some extent within IRISH, another factor must also be considered - the environment. Consider the case where the activity state of a node's neighbours remains constant while it is considering its redundancy. If an anomaly is placed over the sensor in such a way that it experiences a large interpolation error then it will activate even though topologically the network is the same. This illustrates the requirement to select timing parameters not just to allow the network to respond to changes in topology but to allow a response to changes in sensing data [13]. With a long sleep time not only can the topology have changed considerably but so too can the sensing information. Unlike CCP, where beaconing mid active period can improve the neighbour tables for a node's neighbours, there is no need to beacon within the IRISH protocol while a node is active until it needs to re-evaluate it's redundancy.

C. Relationship to CBR Editing Techniques

A similar technique to this has been used successfully in the field of Case-Based Reasoning (CBR). CBR is an AI technique

used to classify instances of a target problem from a set of training data with corresponding classifications. Essentially a case has features and these can be used to locate a case in an n -dimensional region, where n is the number of features. A new cases class can be predicted by examining the class of its k nearest neighbours within the case base. Speed of classification is inversely proportional to the case base size, however accuracy of the classification is typically proportional to size. Notice the correlation between this and WSN energy consumption and competence.

CBR editing techniques attempt to reduce the case base size without compromising the competence of the resulting case base in terms of classification accuracy. One candidate solution for achieving this is called Condensed Nearest Neighbour CNN [7], where a case is removed from the case base if it can be correctly classified by its k nearest neighbours. The reasoning behind this is that such a case is redundant due to the fact that a similar case, to the one being removed, should be classified correctly by the remaining cases. This category of technique typically maintains cases close to boundaries between class regions in the case base.

The analogy with a wireless sensor network is through viewing the nodes and their data as cases in a case base. Nodes have locations and samples of the environment are available at its location through its sensors. Using the CNN technique, the QoS of the case base is maintained while reducing case base size. For WSNs a case being removed from the case base is equivalent to it being hibernated for a period of time. A number of other CBR editing techniques also exist such as RNN, ENN, DROP and ICF but the basic CNN algorithm is considered here primarily because it requires the least amount of communication between nodes in order to operate. In the next section, the IRISH protocol is realised within a protocol stack in order to deliver its hibernation capabilities and to evaluate its performance under the EDLA metrics.

D. Related Work

There are a number of other techniques that can be used for power management within WSNs, however, they tend to focus primarily on maintaining a uniform node density. Many of such techniques are surveyed by Wang and Xiao [14]. Borgne and Bontemi [1] have explored work similar to this, however their system is unable to cope with dynamic data at all as they have to recompute the activation patterns for the nodes. This work demonstrates that the IRISH technique can adapt to changing sensed data introduced to the environment by a dynamic target.

IV. EXPERIMENTAL SETUP

The simulation environment used for the experimentation is J-Sim [11]. The simulated area is defined as 100 meters x 100 meters with two deployed node densities for the standard and IRISH setup. Under the standard setup, there is one node every 6m in the environment. The result of this is that a fixed density of 256 nodes are used to cover the region of interest. For the IRISH setup, additional nodes are deployed so that there is a density of one node every 5m resulting in 400 nodes to

cover the area. Recall that node hibernation techniques require redundant nodes to operate, therefore, we must have additional nodes in order for it to operate. The transmission range for each node is fixed at 25m resulting in a maximum hop count of 6 for the most outlying nodes.

The target in the environment is allocated a power of 1000 units and decays according to the inverse square law of distance. 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 [6]. 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 target's characteristics, however, for these experiments its maximum speed is limited to 5 m/s or 18 km/h.

A. System Architecture

In general, few power management techniques for a WSN operates in isolation and therefore nodes must co-operate. When multiple nodes wish to communicate using their transceivers, 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 node for this system architecture will be the MAC layer, with direct control over the transceiver, figure 2. For this experimentation, the 802.11 MAC layer in J-Sim has been used in both configurations.

Fig. 2. Protocol stack to demonstrate standard WSN performance using EDLA trade-offs (left). Protocol stack to demonstrate IRISH performance using EDLA trade-offs (right).

For the standard WSN, GPSR [4] is used to provide the multi-hop communication for nodes out of direct transmission range of the base station. Finally, the application samples the sensed data at the node and relays it to the base station every ten seconds. Above the MAC layer is the connectivity maintenance protocol in the IRISH setup. This layer is necessary because the IRISH protocol is unable to maintain a connected network when hibernating nodes. The CCP protocol is adopted to maintain connectivity using the property that a sensing coverage technique will maintain connectivity when $2 * R_{sense} = R_{transmission}$ [16]. Above this layer is the GPSR routing layer upon which the IRISH protocol operates. Finally the same application is used in the IRISH protocol stack as the standard WSN. The values chosen for the power consumption in the different states of the transceiver are based on those used in the experiments of CCP [16]: Transmit power 1.4W, Receive power 1W, Idle power 0.830W, sleeping power

0.130W. Nodes are initially given 100 Joules of energy, which roughly translates to 1/100 of the potential energy for a node powered using current battery technology.

B. Hibernation Broker

One important aspect of this architecture is the explicit separation of the decision to hibernate from the actual mechanics of performing the hibernation. Within this design the redundancy identification layers do not interact, except to transmit messages between nodes. Another point to note here is that the application must be able to deal with the hibernation and activation of the node. The base station must expect nodes to hibernate and thus receive no data from them on a temporary basis. Not only does the hibernation broker decide on whether or not to hibernate, it also governs various temporal aspects such as how long a decision by a layer persists or whether it ever expires. It decides how long the node should hibernate for and the time between the hibernate signal to layers and the actual shutdown. Various shut down procedures can also be implemented, for instance some layers may be given priority when shutting down i.e. a layer may be called first over another layer, thus giving it more time to inform its neighbours before the transceiver is powered off. In addition, the exact components to hibernate can also be controlled from the mediator [10] and substituting one policy for another can be achieved without disruption to the layers in many cases.

C. EDLA metrics

The strategy adopted here is to use the EDLA trade-offs in the experimentation in order to analyse the performance of the IRISH protocol. With this in mind, target localisation is used to observe the change in application accuracy. 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 [2]. 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 applicable for many environments. They are the Weighted Average Localisation (WL) and the Maximum Signal Strength Localisation (ML) [8] [6] [13]. Two techniques are applied here so that a broader sense of how the application performance is affected by latency can be presented.

At the base station a corresponding application layer receives data and calculates the location of the target based on the sensing information. 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 deliberation, but the greater the subsequent distance the target will have travelled, potentially increasing error. The energy consumed by the network is observed in two ways, the average lifetime of a node and the % of the nodes that remain active. Finally, density or node activity is controlled by specifying the error threshold to the IRISH protocol.

V. EXPERIMENTAL RESULTS

The results of these experiments demonstrate the performance of the IRISH protocol under the EDLA trade-offs. The primary variables under consideration are the base station timeout and interpolation error. The sleep time is also considered as it can have a significant effect of IRISH performance [13]. Specifically, three active/sleep values are used in this experimentation: 5, 10 and 15 seconds. The duration of time allowed for nodes to reply to the interpolation request beacon is set at 0.2 seconds. In each case, the first set of data points obtained are averages of 5 distinct iterations of the simulation. This produces graphs that are dependent on the five target speeds 1m/s - 5m/s and the average of these makes up the results presented here.

A. Energy(Density)

The effect of the choice of error value on node activity is depicted in figure 3. A number of interesting points can be seen from this graph, firstly, node activity decreases as the interpolation error increases. A second point of interest is that this decrease is bounded as can be seen from the small reduction in node activity given the relatively large increase in interpolation error from 5 to 10. The diminishing node activity from figure 3 is mirrored in an increasing average node lifetime as the threshold increases, figure 4. Additionally, the trend levels off at around 160 seconds per node, indicating that an asymptotic value exists for the average node lifetime. The reason for the levelling off of both graphs is the underlying connectivity maintenance protocol, which maintains the core set of nodes for the network to route data to the base station.

Fig. 3. Node activity as the error threshold increases.

The energy consumed under the different timing regimes is also illustrated in figure 3. The first point of interest is that the shapes of the three curves are similar. In all cases, as the threshold increases the number of active nodes decreases. The best performer, in terms of activity, is when the period is set to 5 seconds, giving quick evaluation of both the IRISH and CCP connectivity layers. Increasing the value beyond 10 seconds does not impact on the percentage of active nodes in the network, implying that there is a certain limit to the effect timing parameters can have on activity for the IRISH protocol coupled with the connectivity layer.

Similar trends are also observed in the average lifetime of a node, figure 4. In all cases the average lifetime of a node increases as the error threshold does, however, an interesting relationship between the trends exists here. A reduction in the % active metric will typically translates into an increase in the

Fig. 4. Average lifetime of a node as the error threshold increases.

average lifetime of a node, however, the opposite relationship exists here. When an active/sleep period of 5 seconds is used, a lower node lifetime is observed. One of the reasons for this is due to the additional energy consumed from the increased frequency in beaconing and redundancy evaluation. This result is significant in that it shows that node activity is not always a good indicator of the longevity of a node or indeed the network. In a similar fashion to the average % active, the average node lifetime appears to have an upper bound, and this can be seen from the identical behaviour of the 10 and 15 second trends.

B. Latency(Density)

With a large density, corresponding to a small error threshold, a large number of nodes are active and compete for use of the channel. Therefore it takes longer for the same percentage of data to reach the base station. This can be seen from the straightness of the 0.1 threshold trend in figure 5. As the threshold increases, the number of active nodes and the latency of the messages reduces, thus a greater percentage reaches the base station within a given timeout period. This illustrates the tradeoff that must be selected for the application - whether a small amount of fresh data is better than a large volume of potentially stale data. Only the 10 second trends are included here because these graphs outperform both the 5 and 15 second graphs.

Fig. 5. % of active nodes whose messages reach the base station.

The inevitable reduction of this effect is also observed as the error threshold becomes greater than 5, due to the removal of alternate message routing paths and the possibility of sub-graphs becoming temporarily disconnected. This is observed in the graph through the increasing proximity of the trends as the error increases, demonstrating the reduced effect of the error on latency. Additionally, as the threshold increases, the curvature of the trends appears to increase also, implying that larger percentages of data reach the base stations

sooner and thus waiting longer will not significantly increase the additional number of messages received. Evaluating the precise effect this wait has on performance will be achieved through the application accuracy in the next section.

C. Accuracy(Density, Latency)

In order to comprehensively evaluate the IRISH protocol, two localisation techniques are employed which have been discussed previously in section IV-C. The accuracies are evaluated for different interpolation error thresholds and latencies and in order to make the experimental data clearer, polynomial trend lines have been used here rather than the raw line graphs. As in the previous latency experiments, only the 10 second sleep time graphs are included here because they outperform the trends when the sleep time is set to 5 and 15 seconds.

1) *Maximum Localisation Accuracy*: The standard latency effect on accuracy can be seen in figure 6, as the timeout increases, more data reaches the base station thus improving localisation accuracy. After a critical point, however, this curve rises again due to the continual motion of the target while the sensing information is en-route. The effect of density, or node activity maintained by IRISH, can also be seen here through less dense deployments producing better accuracy due to the reduced latency. This can be seen from the optimal error of the trend line around 10.8m for an error threshold of 0.1 but an error of 8m when a threshold of 5 is chosen.

Fig. 6. Localisation error for the ML technique.

The location of the minimisation point for the trend lines is also of interest, it appears that for the denser deployments the optimal point occurs at slightly later timeout values. The shape of the curves also indicates that more dense deployments are less sensitive to latency, based on the flatness of the curve with a threshold of 0.1 as compared to the curve for an error threshold of 10. A diminishing return in accuracy can also be seen from the trends with error 5 and 10, the accuracy difference between these two configurations is negligible even though a considerable change in error is present. The optimal localisation error for this technique is 7.9m for an error threshold of 5 and a timeout of 1.2 seconds.

2) *Weighted Localisation Accuracy*: Similar trend characteristics to those of figure 6, can be seen with the WL technique in figure 7. The minimisation point also occurring at 1.2 seconds and in addition, a similar relationship between error threshold and accuracy is also visible here. The improvement in localisation accuracy for WL, as compared to ML, is minimal with an optimal 6.8m error when the threshold is 5 and the timeout is 1.2 seconds.

Fig. 7. Localisation error for the WL technique.

D. EDLA Comparison

The EDLA trade-offs for the IRISH protocol and the standard WSN are presented in table I. Evaluation of the techniques occurs based on the location of their respective optimal performances. It is crucial to remember that using the EDLA metrics together, it is not possible to artificially improve the performance of one of the metrics without negatively impacting one of the other metrics and this ensures a fair comparison between techniques. In terms of optimal timeout value, the standard approach performs fastest, with its optimal position occurring after 0.8 seconds. From an accuracy perspective, the standard WSN is significantly better at localising the target in the environment but the IRISH protocol achieves a 30% increase in the average lifetime of a node.

	timeout	ML	WL	% active	lifetime	density
Standard	0.8	3.3m	2.7m	100%	120s	256
IRISH	0.9	7.9m	7.8m	50%	157s	400

TABLE I
EDLA TRADE-OFFS FOR THE IRISH PROTOCOL.

Another important result can be seen here, in that deploying nearly double the amount of nodes does not result in a corresponding increase in longevity without a drop off in accuracy. This is due to the additional communications required to maintain the soft deployment. Interestingly, on average 50% of nodes are active under the IRISH protocol, roughly corresponding to the same number of nodes in the basic configuration. Even though a similar number are active, a substantial reduction in application accuracy can be seen. This leads to the conclusion that considerable overhead exists for the IRISH protocol, which increases contention for the channel and contributes to poor performance.

VI. CONCLUSIONS AND FUTURE WORK

When hibernating sensor nodes to conserve energy in the network, care must be taken so that the QoS provided by the remaining resolution is not significantly reduced. This gives rise to the EDLA trade-offs, where a balance between energy consumption and application accuracy is achieved through judicious selection of the set of active nodes. By analysing all aspects of the performance of the IRISH protocol, accurate and fair comparisons can be made between experimental results. In this work, we can see how the IRISH protocol does not perform as well in terms of accuracy as a standard WSN not engaged in power management. It can however, increase the average lifetime of a node. Interestingly, despite

the reduction in accuracy, the increase in average lifetime of a node is not equivalent to the additional redundant sensors deployed indicating that the increase in redundant nodes may not directly translate into an equivalent longevity increase.

One possible reason for the under performance of IRISH is the requirement of homogeneous timing parameters in order for the hibernation broker to make its decision to hibernate. Two conflicting timing regimes are present here - the connectivity maintenance protocol and routing layers that favour a stable topology resulting from long sleep times. On the other hand the IRISH protocol favours short sleep times particularly for a dynamic environment [13]. Our future work will focus on mechanisms to allow individual node hibernation techniques to operate autonomously using independent timing regimes that optimise their performance.

REFERENCES

- [1] Y.-A. L. Borgne and G. Bontempi. Round robin cycle for predictions in wireless sensor networks. In *Second International Conference on Intelligent Sensors, Sensor networks and Information Processing*, Melbourne, Australia, December 2005. IEEE.
- [2] R. R. Brooks, P. Ramanathan, and A. M. Sayeed. Distributed target classification and tracking in sensor networks. *Proceedings of IEEE*, 91(8):1163 – 1171, August 2003.
- [3] Z. Hu and B. Li. On the fundamental capacity and lifetime limits of energy-constrained wireless sensor networks. In *10th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2004)*, pages 38–47, Toronto, Canada, May 25-28 2004.
- [4] B. 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.
- [5] N. S.-N. Lam. Spatial interpolation methods: A review. *The American Cartographer*, 10(2):129–149, 1983.
- [6] J. Lee, K. Cho, S. Lee, T. Kwon, and Y. Choi. Distributed and energy-efficient target localization and tracking in wireless sensor networks. *Computer Communications*, 29(13-14):2494–2505, 2006.
- [7] E. McKenna. *Competence Models for Case-Based Reasoning Systems and their Applications*. PhD thesis, Department of Computer Science, University College Dublin, National University of Ireland, Dublin, Ireland, January 2001.
- [8] S. Pattem, S. Poduri, and B. Krishnamachari. Energy-quality tradeoffs for target tracking in wireless sensor networks. In *IPSN*, pages 32–46, 2003.
- [9] C. Schurgers, V. Tsitsis, S. Ganeriwal, and M. Srivastava. Optimizing sensor networks in the energy-latency-density design space. *IEEE Transactions on Mobile Computing*, 1(1):70–80, 2002.
- [10] A. Sinha and A. Chandrakasan. Dynamic power management in wireless sensor networks. *IEEE Design and Test of Computers*, 18(2):62–74, March/April 2001.
- [11] A. Sobeih, W.-P. Chen, J. C. Hou, L.-C. Kung, N. Li, H. Lim, H.-Y. Tyan, and H. Zhang. J-sim: A simulation environment for wireless sensor networks. *anss*, 00:175–187, 2005.
- [12] R. Tynan, G. O'Hare, and A. Ruzzelli. Autonomic wireless sensor network topology control. In *IEEE International Conference on Networking, Sensing and Control*, London, April 2007. IEEE.
- [13] R. Tynan, G. M. P. O'Hare, and A. 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.
- [14] L. Wang and Y. Xiao. A survey of energy-efficient scheduling mechanisms in sensor networks. *Mobile Networks and Applications*, 11(5):723–740, 2006.
- [15] M. Welsh and G. Mainland. Programming sensor networks using abstract regions. In *First USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '04)*, 2004.
- [16] G. Xing, X. Wang, Y. Zhang, C. Lu, R. Pless, and C. Gill. Integrated coverage and connectivity configuration for energy conservation in sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 1(1):36–72, August 2005.