Research Repository UCD

Title	A stable routing framework for tree-based routing structures in wsns
Authors(s)	Delaney, Declan T., Higgs, Russell, O'Hare, G. M. P. (Greg M. P.)
Publication date	2014-10
Publication information	Delaney, Declan T., Russell Higgs, and G. M. P. (Greg M. P.) O'Hare. "A Stable Routing Framework for Tree-Based Routing Structures in Wsns" 14, no. 10 (October, 2014).
Publisher	IEEE
Item record/more information	http://hdl.handle.net/10197/7208
Publisher's statement	© © 2014 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Publisher's version (DOI)	10.1109/JSEN.2014.2329391

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

A Stable Routing Framework for Tree-Based Routing Structures in WSNs

Declan T. Delaney, Russell Higgs and Gregory M.P. O'Hare

Abstract—Achieving industry standards for transmission of data is an ultimate goal for Wireless Sensor Network (WSN) designers. Such standards are difficult to attain due to the challenging communications environment in which WSNs operate where best transmission routes may change rapidly. In such dynamic environments it is not only advantageous but commonly necessary to change the used route for better performance. However, with so many changing routes it is possible to introduce instability, which can hinder the provision of services. Route instability can be particularly detrimental in a tree-based routing structure. In these scenarios it is beneficial to maintain more stable routes in the tree to preserve high standards. The focus of this paper is the development of a route stability framework whereby currently used metrics are adapted to promote routes that achieve greater stability in highly dynamic network conditions. The central concept introduces *Neighbourhood Heuristics (NHs)*, a method of combining a sensor's routing metric with those of its neighbours to highlight both the quality of the current route and the quality of the routing options available to the sensor should its current route become unavailable. The additional information afforded by the new combined metric allows sensors to choose good quality routes that can better maintain quality despite the degradation of an upstream link. The NHs framework is implemented with the Routing Protocol for Low Power and Lossy Networks (RPL) routing protocol. Experiments are conducted both in simulation and on an open public testbed which compare routing stability using the Expected Transmissions (ETX) metric and ETX under the NHs stability framework (ETX-NH), showing a marked increase in route stability for ETX-NH over using ETX.

Index Terms—Wireless Sensor Networks, routing, metrics, heuristic

I. INTRODUCTION

DESPITE WSNs having far-reaching applications in a number of fields from military to medical [1], [2], wider industry has yet to adopt WSN solutions for production use. WSNs can provide cheap and flexible solutions but they still suffer from a functional issue in delivering the strict requirements needed by industry.

Building Energy Management Systems (BEMSs) provide a pertinent industrial use case for WSNs [3], [4], but also

D.T. Delaney and G.M.P. O'Hare are with the CLARITY: Centre for Sensor Web Technologies, School of Computer Science and Informatics, University College Dublin, Dublin 4, Ireland. e-mail: declan.delaney@ucd.ie; gregory.ohare@ucd.ie.

R. Higgs is with the School of Mathematical Sciences, University College Dublin, Dublin 4, Ireland. e-mail: russell.higgs@ucd.ie

This work is supported by Science Foundation Ireland (SFI) under grant 07/CE/I1147.

The authors would like to thank the reviewers and editor for their suggestions and comments to help make this a more complete and accurate document.

The authors would like to thank Daniel Bimschas at the University of Lübeck for his patient and prompt handling of queries with regard to the Wisebed testbed.

present a number of interesting challenges for WSN design. BEMSs are tasked with monitoring the environmental state of a building and controlling the building infrastructure to best support the user's needs while maintaining energy efficiency.

A number of challenges, however, must be addressed before a WSN can become a viable solution for the BEMS space. Much of this is directed at the underlying network, which carries the sensed data and instructions to the required destinations. Buildings provide a harsh environment for wireless communication. Difficult link conditions between sensors in a WSN are well-documented [5]. This effect is exasperated in dynamic environments such as office or industrial spaces with moving objects and doors frequently changing the communication links between sensors. Increasing link dynamics will introduce greater route instability which can hinder the networks capability to maintain high standards. Addressing the issues of instability can lend a platform from which WSNs can provide an elevated level of service.

Route instability presents a problem in many network structures, from the Internet's border gateway routers [6], *ad-hoc* networks [7], [8] to Mobile Ad-hoc NETWORKs (MANETs) [9]. Route stability is described as one of the current underlying challenges to communication in WSNs [10]. Excess route instability is detrimental in multihop networks, causing routing loops and affecting overall performance. Woo *et al.* [10] use stability as a metric to evaluate routing protocols in WSNs and suggest that the stability of a network affects the design and implementation of higher level functions, such as scheduling and aggregation. Intelligent agent systems [11] for WSNs will also benefit from a stable platform.

Tree routing has become a popular structure for use with BEMSs [12], [13]. All sensors in a routing tree send data to a single collection point, known as the root, with each sensor choosing a parent through which it will forward all data. Figure 1 shows how a link metric can be used for minimum cost path routing in a routing tree. The figure shows how both the advertised and link metric of a parent combine to make the sensor's metric score. Its also highlights the dependency between a sensor's metric value and each of the links between it and the root, exposing a particular frailty that exists in tree-based routing with regard to instability. Should a single link value in the route change, a knock-on effect occurs affecting each sensor on the route down from this point. This phenomenon can lead to numerous parent changes, as mentioned in [10], leading to greater instability in routes.

Other efforts for managing stability rely on route repair [14], [15]. The approach presented in this paper differs in that it focuses on reducing the impact of a parent change on the

Fig. 1. Using an arbitrary link quality metric for cumulative minimum cost path routing. This process is repeated until all sensors in the routing tree have acquired a routing metric.

routing tree. Should a link change occur forcing a parent change, the effect of this change is minimised and the knock-on effect is greatly reduced, suppressing further route changes.

Neighbourhood Heuristics (NHs) describe a process to combine a sensor's routing metric with those of its neighbours to highlight both the quality of the current route and the quality of alternative routing options available to the sensor. Sensors with good quality routes to additional neighbours are subject to a bias, which increases their chances of being used as a forwarding route. Each sensor then has a greater chance of choosing a parent with good quality alternate routes. In the event of a sensor's parent's primary route being disturbed, the parent has a better chance of routing through another neighbour with a similarly good route metric, minimising the effect on the sensor.

NHs provide a framework for stabilising routes chosen using any routing metric without imposing heavily on the characteristics of the metric itself.

The concept outlined in the paper is evaluated both in simulation and on an open public testbed. The framework is implemented under the Routing Protocol for Low power and Lossy Networks (RPL) protocol using the Expected Transmissions (ETX) metric. A BEMS application scenario is adopted to evaluate the technique.

The contributions of the paper can be summarised as follows:

- Providing a technique for combining both the current route quality and the quality of the alternate routes on a sensor. This achieves a routing score that portrays the current forwarding capacity of the sensor as defined by the metric and also the capacity of the sensor to handle a disruption to the current route gracefully.
- The introduction of NHs, a route stability framework for tree-based routing protocols. NHs minimises the domino effect of route changes in a routing tree leading to a more stable network, while maintaining the key characteristics of the metric used.
- Promoting the wider spread of traffic in a tree-based structure, reducing the load on sensors bearing the heaviest communications burden.

II. MOTIVATION

The nature of WSNs is fraught with difficult communications conditions with highly dynamic links. As a result much of the research in the field is directed at exploiting the temporal nature of these links and embracing the need for frequent route changes. This has led to the principle of ensuring stable routes within the network being neglected. While there are clear benefits seen from changing a route based on current conditions, there are a number of services that are severely disrupted by a route change. These higher level services are practices considered integral to WSNs such as data aggregation [16]–[18], scheduling [19]–[21] and Transmission Power Control (TPC) [22]–[24]. In each of the practices, route changes in the tree lead to reorganisation resulting in communications or performance inefficiencies.

Madden *et al.* [16] present a highly regarded data aggregation solution. The authors maintain that route instability has a negative effect on the aggregation process, stating that any parent switch can cause temporary disconnections that can lead to lost records. High error is reported in the aggregation process when a route is chosen through a poorly connected subtree. Such results act as a catalyst for the work presented in this paper. Route changes are also detrimental when performing predicted aggregation [17], [18] on learnt data. Predicted aggregation uses regression or estimation on previously learnt data to predict future aggregates. When a sensor changes from its current parent all learnt data is lost greatly reducing the aggregation efficiency.

A similar situation arises for both scheduling and TPC. In many scheduling algorithms, for example distributed TPC, the routing structure is decided upon in an a priori manner. Based on this routing structure a scheduling or power strategy is devised for efficient communication. As the routing structure changes the current strategy becomes less efficient to a point where it may require further setup phases leading to unwanted communication overhead.

The key to providing increased levels of QoS in WSNs is to continue to exploit the link conditions to choose routes that are currently performing well while maintaining route stability. This provides a platform for higher level network services to perform at their best. With this in mind the framework presented in this paper aims to satisfy two conditions:

- Demonstrate adherence to a routing tree that reflects the intentions of the routing metric chosen by the user.
- Reduce the number of changes to the tree in order to maintain higher levels of route stability.

III. RELATED WORK

Stability remains a burning question in wireless networks. The issue of route stability is noted as a key weakness in WSNs [10]. WSNs exhibit unique link dynamics that necessitate a unique solution for stability. We show how these dynamics are exploited in WSNs, and how different methods of dealing with these dynamics can effect stability in the network. The state of the art methods, protocol and metrics are then examined and we show how our technique can offer a valuable addition to the field.

Route stability is discussed in a number of closely related fields such as MANETs [9] and *ad hoc* networks [7], [8]. Chen *et al.* propose a stability structure based on the persistence of a sensor's neighbours. A sensor is more likely to route through a neighbour that has a greater tendency for constant connectivity. The approach proposed does not consider the quality of the link, leaning more towards route persistence and readability. With constant movement expected in MANETs this solution is justified. Route persistence is important for all reliable wireless communication, but for this method to be used effectively in a WSN context a link or route metric would also need to be considered to choose the best path.

Using static sensors [7] may better reflect a WSN suited for BEMs, as most sensors are likely to remain in a fixed position. Ramachandran *et al.* similarly use persistence, together with prevalence and route flap as metrics in defining the stability of a network. A number of interesting results are drawn from the experiment. It was determined that routes in a static mesh network are weakly dominated by a single prevalent route. Persistence analysis was performed on the dominant routes illustrating that such routes on average are short-lived. In fact the paper shows a high quantity of short-lived route flaps with many of these flaps made with minimal throughput advantage from the original route, indicating a level of instability for minimal gain. The paper recommends setting thresholds for route selection to further assert dominant routes and increase stability. The recommendations posed in the literature are useful but as the link dynamics in WSNs are more volatile further insights are necessary.

In [5], the authors provide a detailed analysis on the fundamental dynamics experienced by low power radio links such as those used in WSNs and provide statistical models for both short-term link dynamics and long-term link quality. This gives rise to many techniques for estimating link quality between sensors, which inform a sensor when choosing a neighbour to forward data to. There are two schools of thought about how to address these pernicious dynamics. One method has parallels to opportunistic routing and takes advantage of the broadcast nature of wireless networking and exploits short-term performance fluctuations on links [25]–[27]. The other is to look for and use links that show greater stability over a longer period [28]. The first technique represents the best route through a network at a given time, the second represents a more stable solution that prefers routes which show good performance over a longer time period. To strike a balance between route stability and maximising the use of the best link available is critical. In [25], the authors present a link quality estimation technique that can take advantage of short-term link dynamics, allowing a sensor to choose the best link over short spaces of time. While this can improve throughput over the short-term, both [26], [27] show that the value of opportunistic routing is heavily dependant on the link correlation in the network. In [28], a link metric, *competence*, is presented that measures the long-term quality of a link. Also proposed is a framework for routing that promotes more competent (or stable) routes increasing stability further. This technique is shown to give performance increases over ETX and Link Quality Indicator (LQI) based routing.

Fig. 2. Showing the domino effect of a single link change in the network of a tree routing structure.

This method does increase stability in a network, but only works with a single metric, which does not provide a construct where by a wide variety of different requirements can be achieved from the network. In [29], the author aims to model the variations on the link using the Hidden Markov Model (HMM) to determine in advance if a link will be capable of facilitating communication and can be used to determine times when a link proves consistent. This constitutes an opportunistic paradigm but can be tailored to fit a more reliable solution.

Stability can be greatly affected by the chosen route metric in wireless networks. Passos *et al.* [8] use the Optimised Link State Routing (OLSR) routing protocol to show the different characteristics for the ETX metric and the proposed Minimum Loss (ML) metric, which choose paths that exhibit minimum packet loss ratio. Among other measured differences the disparity shown between the two metrics with regard to stability was vast, with ML showing fewer route changes. This highlights the importance of metric make-up on route stability and points to a key role that metrics can play in maintaining stability in a WSN.

Much work has been conducted on route and link metrics to achieve stability and maintain quality routes in a network. The technique presented in this paper differs from other approaches in that it does not build a single metric that achieves greater route stability, but provides a framework whereby any metric can be adapted to better deal with instability in difficult link dynamics. This subtlety is important as this strives to maintain the variety of network responses delivered from the range of metrics available while maintaining network stability.

IV. PLAYING DOMINOES WITH TREES

The tree routing concept is simple and intuitive. A single (or multiple) data sink acts as a root to the tree. All data

in the network is directed towards the root for collection. Each sensor chooses a single neighbour that it will use to forward all data. This chosen neighbour is called the sensor's *parent*. Using a single parent to forward all data reduces the chance of routing loops occurring in the network. The parent is chosen as the neighbour that is in the best position to forward data to the root, or is the root itself. The routing metric used will determine how fit a neighbour is to forward data to the root. Figure 1 illustrates how a sensor chooses a parent by evaluating the neighbours, using an arbitrary link estimation metric as a minimum direct routing metric.

There are inherent stability issues present when using the tree structure. The tree sets up a routing environment whereby a sensor's metric value is wholly dependent on each of the links that form the route to the root. A sensor may in turn act as a forwarding sensor for other sensors down the branch. Fluctuations on a single link in the chain can end up affecting the metric value of a large number of sensors in the network. Furthermore, a large change in a single link instigating a change of route can propagate down the branch resulting in many more route changes. Figure 2 illustrates this in a simple network. A network forming a tree structure is presented in figure 2a. Many sensors in the tree form a route through *A*. Sensor *A* uses link 1 to forward all traffic from sensors down the branch. As link 1 becomes unavailable, sensor *A* must resort to link 2 as its primary link as no other routes are available to it as shown in figure 2b. Link 2 is of poor quality and results in a severely degraded metric value for sensor *A*. The immediate sensors, *B* and *C* learn of the degraded link as *A* updates its metric value, and react accordingly by switching to alternative routes as depicted in 2c. Sensors *B* and *C* metric values are also adversely affected as a result of the route changes. As their neighbours learn of the updated metrics they re-evaluate routing options. Figure 2d shows the resultant network after the changing metric values propagate through the network causing more sensors to change route. This phenomenon is known as a *domino effect* and is a fundamental issue in tree type routing. Reducing the domino effect is key to achieving increased stability.

Currently, instability is managed by applying thresholds to those metrics used. This dictates that a sensor will only change parent to a competing neighbour, if the competing neighbour advertises a metric that is better than the current parent by a predetermined amount. This amount is carefully chosen as it poses a trade-off between stability and dynamism in the network; two aspects we wish to maintain in the network. This solution is easy to implement but relies solely on a single threshold and therefore does not constitute a complete solution. NHs is unique as it aims to increase stability by asserting proactive damage control by routing through sensors offering good failover options. Route changes are inevitable, but if a used link goes down unexpectedly, another similar route is more likely to be available to carry traffic minimising the knock-on effect of the change preserving stability.

V. NEIGHBOURHOOD HEURISTICS

To address the issue of stability in tree-based routing structures, this paper proposes NHs, a novel framework to reduce

the knock-on effects of a route change occurring within the tree. NHs provides a framework for routing metrics providing a means to include additional contextual information into the metric regarding a sensor's ability to handle route failure. NHs facilitates the combination of a sensor's metric with those of its neighbours, creating a single metric which highlights sensors in the network with good failover options. NHs necessarily introduces additional terminology:

- *Neighbourhood Metric (NM)* is the value assigned to a sensor as a result of the NHs process.
- *Neighbourhood Effect (NE)* is the influence a single neighbour, or whole neighbourhood, lends to a sensor's NM.

NHs aims include:

- Delivering increased stability in the network by mitigating the domino effect in trees.
- Maintaining flexibility, accommodating a rich set of metrics.
- Preserving the communications goal of the network as prescribed by the underlying metric.

The process involves evaluating each of the neighbours, excluding the current parent, for possible failover capability, scoring the neighbours on both likelihood of becoming the failover option and the quality of the resulting route upon using the option. This score is the NE applied by that neighbour to the sensor. The NE of each neighbour is added to the sensor's own current routing metric to create the sensor's NM. The sensor's NM is used for all routing decisions. NHs achieve increased stability by identifying sensors that present both a high quality route and suitable failover options. Figure 3 highlights how NHs discriminates between good and bad failover options by using the NE from the neighbourhood. Figure 3a and 3b show how an arbitrary minimum cost path routing metric does not consider the failover option, using only the accumulated link values to the root as the sensor value. Using NHs the sensor adopts a NE from the additional available neighbour. Figure 3c presents a good failover option to the sensor and hence applies a strong NE, making the sensor value more attractive. A bad neighbour, however, will lend a lesser NE as depicted in figure 3d. The means by which the NE is calculated is described in section V-C. The underlying metric ultimately remains responsible for achieving the desired response from the network. By adopting a framework that accommodates a number of different routing metrics rather than presenting a single metric for stability, NHs can provide for a wide range of network requirements.

A. Concept

NHs combine additional neighbourhood information into a sensor's routing metric. This additional information contains insights into how a sensor may handle the failure of a currently used route. NHs increase stability in a network by defining *sensors*, rather than single routes, that can facilitate good quality communication using information, which is freely available yet unused in current metric computation implementations. The NHs mechanism combines all of a sensor's possible routing options into a single metric yielding a more holistic

Fig. 3. Using an arbitrary link metric as a Minimum Cost Routing Metric to show how the NH process uses the NE to promote sensors with good neighbours.

perspective of a sensor's ability to facilitate communication in the network. Sensors that have a larger number of good routing options are allowed to advertise an improved metric value to their neighbours, essentially up-voting sensors with good failover options. This increases the chances of sensors with multiple routing options being chosen for data forwarding and consequently reduces the chance of routing through a sensor with no, or very few, good neighbours. The benefits of this are illustrated in figure 4.

A simplified network is presented with data source and tree root. Links in the figure are represented in two ways. The quality of the route is represented using dashes on the link. All available links are represented on the figure. Links that are chosen to forward data are represented using a thicker line, with available but unused links using a thin line. The figure includes a data source and data sink (root). The source must use the network to transfer this data to the sink. The data source must route through either sensor A or B to forward data through the network. Sensor A presents two routes. One represents the best route to the root, the other a poor route. Sensor B maintains two good routes to the root. Figure 4a shows the routing structure as deemed by the current approach. The source chooses to route through A using the single best path presented in the network. Figure 4b sees the source choose a route through B implementing the NHs approach. In both cases the load-bearing link L is removed requiring a restructuring in the tree. In figure 4a, sensor A must choose a path through a bad link. This forces the source to change route from A to B, resulting in two route changes. The lost link L in figure 4b has a lesser effect as B changes to route through another good path. This does not greatly impact the metric value that the source sees through B allowing it to maintain a good route to the root using B. The removed link leads to a single route change in this case. It is important to note

however, that the source still has a route available through A that is better than B. The source does not change to this route due to the stability threshold mentioned in section IV. This threshold still plays an integral part in maintaining stability. NHs then presents a trade-off between choosing the single best route to the source at any given time and choosing a lesser quality route that proves more stable under route failure.

Using NHs a change in metric is less likely to propagate down the branch leading to route changes. Sensors that act in a forwarding capacity on the branch have a higher chance of having immediate neighbours of better quality. In the occurrence of a forwarding sensor failure or link disturbance, the next best neighbour is chosen to route through. Using NHs the next best neighbour is more likely to be of high quality. When another good quality neighbour is at hand, the change to the routing metric propagated down the branch is less significant leading to fewer route changes and increased stability.

B. Mechanism

The mechanism catering for neighbourhood metrics is designed with low communications cost of paramount importance. The process can be broken down into four phases: metric collection, parent selection, neighbourhood metric calculation and metric advertisement.

- *Metric collection:* The first phase requires the collection of metrics used in making routing decisions. The advertisement messages from each of a sensor's neighbours are received. Each advertisement message is inspected for both the routing metric and neighbourhood metric. Both metrics may be available from a single advertisement message from a neighbour. The sensor ID of the neighbouring sensor, advertising its metric, is also present in the advertisement message. All information is stored in a *neighbour table*. The sensor calculates the link metric for each neighbour and stores this metric in the neighbour table. The sensor updates the neighbour table on a continuous basis as it receives the advertisement messages from its neighbours. Each sensor maintains a list of neighbours with their respective metric information. When all relevant metrics are collected a sensor has all the information necessary to select a parent.
- *Parent selection:* Only the neighbourhood metric and link metric are considered when choosing a parent. Combining the link metric and neighbourhood metric will give an overall score for each neighbour. The neighbour with the best overall score is chosen as the sensor's parent. The chosen parent is indicated in the neighbour table. After a sensor has chosen its parent, the neighbourhood metric can be calculated.
- *Neighbourhood metric calculation:* The routing metric and link metric of each of the entries maintained in the neighbour table are used to calculate the sensor's neighbourhood metric. An overall score for each neighbour is calculated. The overall score of each neighbour reflects the routing metric the sensor would assume should it use that neighbour as parent. The NM is determined by

Fig. 4. Showing the benefits of NHs over current practice to route stability after a link disturbance.

combining each score using a weighted function. Scores are weighted depending on their relative distance from the chosen parent's score. The weighted function is described in V-C. A sensor with no neighbours will maintain an undefined neighbourhood metric precluding it from acting as a possible route to other sensors.

- **Metric advertisement:** The final phase entails advertising of metrics to the network. Both the routing metric and the neighbourhood metric must be advertised in order to facilitate the mechanism. The neighbourhood and routing metric may be sent in the same advertisement message. The regulation and timing of advertisement messages is defined by the routing protocol and is out of scope of this paper. The neighbourhood metric for each sensor is then advertised throughout the network in a similar manner as the routing metric until each sensor has advertised its neighbourhood metric to its neighbours.

The strength of this mechanism lies in utilising additional information without necessitating additional communications overhead. The neighbourhood metric is advertised and propagates through the network in the exact same manner as the routing metric advertisements. Each neighbourhood metric, like the routing metric, is composed of a single value. As these messages can be shared, the neighbourhood metric can essentially *piggyback*, causing minimal overhead.

C. Weighting Neighbours

The means by which neighbourhood metrics are calculated is pivotal in achieving the desired behaviour from the network. The neighbourhood metric should reflect the real quality of a given route chosen but also carry information regarding the ability to handle failure or disturbances to the route. This must be achieved using a single value metric. A single value

metric removes any ambiguity for a sensor when choosing a parent. If a neighbour in the neighbourhood has the single best neighbourhood metric value, this becomes the preferred parent with which to forward data. If two neighbours exhibit the same, best value, a simple choice for preferred parent can be made using the sensor ID. The NM is calculated by starting with the sensors routing metric and adding a NE from each of its neighbours. How each neighbour effects the sensors overall NM is described by a set of directives which are defined in this section.

Use of routing metric: The routing metric used in the network will determine the real quality of the route used. The real quality is important as it produces the response desired from the network. It is important for the neighbourhood metric to reflect the real quality of the route, thus the routing metric for the current forwarding route plays a primary role while calculating the NM.

Similar routing options: In order to reduce major fluctuations in the routing metric when a disturbance occurs, which requires a parent change, a sensor should change to a parent that presents a similar routing metric to the original route before the disturbance. The sensor will then maintain and advertise a similar routing metric that will not greatly affect the rest of the branch. This concept is key to maintaining stability in the routing tree. Routing through sensors which have multiple routes of similar quality is therefore advisable. This effect will also be considered when calculating the neighbourhood metric.

Positive effect: To further define the nature of the NM, examination of a scenario is helpful. Figure 5 highlights a scenario where a sensor can choose a routing path between two neighbours. The figure shows a network of sensors with data source, and a data sink marked as root. The source must choose

Fig. 5. NHs prefers routes that present any failover advantage. Despite the poor additional gains offered by sensor *C*, neighbour *A* is preferred in NHs routing.

a path to the root through either neighbour *A* or *B*. Neighbour *A* has two neighbours, one with good links to the root, which result in a good routing metric value. One with bad links to the root, resulting in a poor routing metric value. Neighbour *B* has a single neighbour with good links to the root. Despite the poor quality of the additional route which *A* provides, for the sake of stability it is assumed that route *A* still provides a marginal advantage over route *B*. Each additional neighbour will then have a positive effect on the neighbourhood metric. This is not to say that it is worth *changing* parent from sensor *B* to sensor *A* if this scenario were to occur. A stability bound as discussed in section V-A is important in this regard. The stability bound will determine at what point the sensor should change parent over to an improved route. A change in route should only occur if there are significant improvements to the real route value. Changing over to a route of equal or lower real value, even with superior failover capacity, undermines the purpose of the mechanism, namely to promote stability. The value of the neighbourhood metric must minimise influence on a parent change, but must still influence a sensor's ability to be chosen. The neighbourhood effect applied to a sensor by its neighbours must then be relative to the stability bound. The neighbourhood effect on the metric must never surpass the stability bound resulting in a parent switch.

Diminishing return: Each neighbouring sensor, which is part of the routing tree, will offer a route through which the sensor can send information. Each neighbour, not chosen as parent, can be considered to offer a failover route should the current route become unavailable or compromised in some way. Each neighbour will increase a sensor's probability of maintaining a route to the root. However, each additional neighbour that is added delivers a diminishing return on maintaining a route. If a sensor loses a route through its current parent, it will switch to route through the next best neighbour. If both the current parent and next best route become unavailable the sensor must route through the second best neighbour. This is less likely to occur. In order for the third best neighbour to be chosen as the new parent, both the best and second best neighbours must become unavailable making this scenario far less likely to occur. This effect must be given consideration when calculating the neighbourhood metric.

Tunable metric: Depending on the deployment and the needs of the network the reliance on failover information may change. A fluctuating network that experiences excess route flapping may want to rely more heavily on routes with good failover possibilities, while a stable network experiencing little network change will be better served using the routing metric

Fig. 6. Using a Gaussian weighting function to weight the neighbourhood effect.

without a neighbourhood bias. This gives rise to a need for a tunable metric.

From these observations a set of directives can be used when creating the neighbourhood metric for a given sensor:

- The *routing metric* of the chosen parent must remain the primary component of the resulting NM.
- The NM should promote sensors that have many *similar routing options*. The NE of any neighbour will be greater as the value tends towards the chosen parent metric value.
- Each neighbour ought to have a *positive effect* on the NM, with the effect considered as a fraction of the stability bound.
- Each additional neighbour has a *diminishing return* on the NM.
- The NM must have a capacity to be *tunable*, allowing for the different requirements of separate deployments.

Figure 6 shows how a Gaussian weighting function can be used in weighting neighbour metrics for use in calculations. The Gaussian function maintains a value of 1 at its peak. A sensor will use this function to give a weight to each of its neighbours. The routing metric value of each neighbour and the sensor itself are plotted on the *x*-axis. The Gaussian curve is centred on the sensor's own routing metric value. Centring the expected value of the curve on the routing metric value promotes neighbours with a similar metric, while penalising neighbours of dissimilar metric. The Gaussian fall-off graduation will determine how heavily each sensor is weighted. The figure shows the weighting applied to two neighbours: *N1* and *N2*. Here *N1* is considered a good neighbour as it has a similar metric value to the chosen parent's value, whereas *N2* is considered a bad neighbour because it has a metric value far removed from that of the parent's value. The steep fall-off presented by the Gaussian curve ensures that neighbours with similar values can be differentiated.

We define and implement these directives in a simple set of equations. Equation 1 shows the equation used in calculating the neighbourhood effect of a single neighbour on the sensor.

$$\text{Neighbourhood Effect} = \exp\left(-\frac{(\Delta v)^2}{2\delta^2}\right) \times \frac{1}{i^2} \cdot \frac{6}{\pi^2} \times \Theta \quad (1)$$

Algorithm 1 describes how the combination of this effect produces the neighbourhood metric for a sensor. The value

$\exp(-(\Delta v)^2/2\delta^2)$ determines the weighting applied to the neighbour as a function of its value and how close this value lies to the parent metric value, where Δv is the difference between the current parent's value and the neighbour's value and δ^2 is the Gaussian variance used to control the width of the function. The term $1/i^2$ provides a diminishing return in weight to each additional neighbour as a quadratic fall-off, with i the index of the neighbour at hand. Θ is the stability bound for the metric. The term $6/\pi^2$ is added to ensure the resultant effect is never greater than the stability bound. The term derives from the "Basel problem" of the summation of infinite series. This addresses an important facet of the directives stipulating that the cumulative *positive effect* of all neighbours to a sensor will never be greater than the stability bound.

The equations above are used in a simple algorithm to calculate each neighbour's neighbourhood effect and the resulting neighbourhood metric for the sensor. The algorithm first sorts

```

sensor.NM ← -1;
sensor.value ← -1;
if neighbour update then
  if new neighbour or neighbour state change then
    Ns ← update neighbour table;
    sensor.value ← parent.value + parent.link;
    //order neighbour set;
    Ns → List in order of best to worst value;
    N_effect ← 0;
    for i = 1 → Ns.length do
      diff ←
        sensor.value - (Ns[i].value + Ns[i].link);
      N_effect +=
        exp(-(diff)2/2δ2) × Θ/i2 × 6/π2;
    end
    if min metric desired then
      NM ← sensor.value - N_effect;
    else if max metric desired then
      NM ← sensor.value + N_effect;
  end

```

Algorithm 1: Calculate Neighbourhood Metric for a sensor.

the list of neighbours into highest value to lowest. This is an important step as the neighbour with the highest value will present the highest return as it is the most likely candidate to be used in a failure scenario. The neighbourhood effect of each neighbour is then calculated. The effect of each neighbour is accumulated and then added or subtracted to the sensor's routing metric to produce the sensor's neighbourhood metric. Using this method the routing metric of the sensor remains the most prominent factor in the neighbourhood metric, each additional neighbour provides a diminishing return with a quadratic fall-off effect and allows the neighbourhood metric to be tuned. The Gaussian variance (δ^2 in Equation 1) can be tuned to give greater or lesser weight to weaker sensors in the neighbourhood.

VI. EXPERIMENTATION AND RESULTS

In order to validate the proposed framework a number of experiments were undertaken. The experiments were conducted in both simulation and physical testbed environments. The goal of the experiments was to compare the NH method with the current RPL implementation and determine how each solution handles route stability in networks of varying dynamics. The simulation environment was used to evaluate the NHs solution at scale. The testbed is used to evaluate the NHs solution is a real environment exhibiting real world dynamics. Stability of the routing structure is closely monitored for both the current RPL implementation and RPL using NHs. The end-to-end evaluation metrics: reliability, latency and energy efficiency are also considered for comparison.

A. Experimentation

Experiments were conducted using both micaz and TelosB sensors [30]. TinyOS [31] is used as the sensor platform for conducting experiments. All sensors contributing to the experiments run the TinyOS architecture. The RPL routing protocol is used as the routing protocol for all experiments. RPL is fast becoming a new standard for WSN protocol and represents the state of the art for industry deployment. The RPL design shares a similar philosophy with NHs, particularly with regard to routing metrics. Like NHs, RPLs mechanism is not bound to a single routing metric and encapsulates within its headers a number of fields that allow a description of the metric used. In future releases this feature could be used to notify the network whether or not to apply the NH mechanism to the chosen metric. The standard "out of the box" RPL implementation using the minimum rank with hysteresis object function [32] is used to compare with the NHs solution. Both implementations use the ETX metric for routing. Minimum rank with hysteresis represents current routing practice. ETX is an extensively tested metric providing quality results. The experiments were conducted over one hour intervals, where the network is set up and maintained for this hour. The experimental scenario considered reflects a simple BEMS application, where each sensor senses the environment and sends this information to a collection point at regular intervals. The packet frequency is set to send a packet every 60 seconds per sensor. This represents a fine-grained monitoring system for a BEMS dealing with building environmental control. This kind of scenario is relevant to many monitoring type scenarios suggested for WSNs.

Scale and density of the network are factors that will affect how a routing solution handles stability. A variety of scale and densities are presented in the experimentation. The simulation environment is used to examine stability in networks of up to 500 sensors, while the testbed environment is used to examine stability in varying densities and real world dynamics.

B. Simulation Environment

Simulated experiments were conducted in Tossim [33]. Tossim is a platform for large scale networks simulating the micaz CC2420 radio. Each experiment was performed

TABLE I
RELATIONSHIP BETWEEN NUMBER OF SENSORS AND SCALE OF THE
SIMULATION NETWORK.

Number of sensors	Scale	
	ETX	ETX-NH
50	2.3	2.5
100	3.4	3.4
150	3.7	3.6
200	3.9	3.7
250	4.3	4.1
300	4.8	4.7
350	5.3	5.1
400	5.5	5.2
450	5.7	5.5
500	5.9	5.6

Uniform Density: 15

on a network with randomly generated sensor positions. The random nature of this assignment assures the results derived from experiments are not specific to a single topology. The only changes made between simulations are sensor positions (random nature of change) and number of sensors (step change of 50). There are no changes made to the topology during simulation other than those affected by the protocol. The network is set in an open field with free space and log normal path loss transmission between sensors. Multipath effects in such an environment are minimal. Each sensor maintains at least one link to another sensor in the network allowing full connectivity. A noise floor is established in the scenario. A dynamic communication environment is simulated by modelling noise using Close fit Pattern Matching (CPM) over a noise trace collected in a real network. The “Meyer heavy” noise trace was used for all simulated experiments. The Meyer traces are a set of real noise traces taken from the Meyer library at Stanford University. Experiments were conducted in networks varying in scale from 50 to 500 sensors while maintaining a uniform average density for all experiments. This ensures the scale increases in terms of number of sensors in the network and in terms of average number of hops necessary to reach the root. Table I details how scale increases with additional number of sensors during simulations.

C. Reducing Parent Changes

The number of parent changes in the network is critical to route stability. Fewer parent changes presents a more stable network. A key factor to reducing the number of parent changes in the network is minimising the impact of a parent change on the value of a metric, thereby limiting further changes. Figure 7 shows a comparison between ETX and ETX-NH with this in mind.

The figure shows a simulation experiment consisting of 50 sensors. Each sensor shows the average difference in routing metric experienced by a sensor after a parent change. The average difference for the network is also displayed. During this experiment the metric values on all sensors were continuously monitored. The difference is measured when the routing option through the sensors current parent degrades to a point where another neighbour becomes a better option. The difference displayed is the difference between the routing

Fig. 7. Bar chart showing a simulation experiment consisting of 50 numbered sensors. Each bar represents the average difference in metric value experienced by a sensor after a parent change.

Fig. 8. The direct effect of a single parent change on the immediate network is expressed here as a probability of the change promoting further changes. This figure shows the probability of the change effecting one, two or three further changes. Each point represents an experiment conducted over varying scale from 50 to 500 sensors.

metric seen by the sensor before the route through its previous parent degraded and the routing metric seen by the sensor using its new parent. A greater difference suggests the network is routing through sensors which have more dissimilar routing options at their disposal. When a parent change occurs the effect on the sensor’s metric is more profound, which may not

Fig. 9. Leaf sensors in the network as a percentage of the total number of sensors. Leaf sensors are not used as a forwarding route for any other sensors.

Fig. 10. Total number of parent changes occurring in the network over increasing scale for both ETX and ETX-NH.

only lead to a further immediate change, but also a greater effect as this new metric propagates down the branch. It is clear from figure 7 that the solution using NHs minimises the impact of a parent change on the sensor value. This result is expected and is achieved by virtue of the additional mechanics supported in NHs to promote routing through sensors that have many similar alternative routing options. While this experiment shows the benefits of the NHs solution on the stability of a single sensor's routing metric value, we are still curious to see the significance of this result in the wider network.

A large scale experiment was conducted to determine this and discover its relevance to further stability in the network. The focus of the experiment was to examine the direct effect of a single parent change on further parent changes in the immediate network. Figure 8 shows results from a set of experiments varying in size from 50 to 500 sensors as described in section VI-B. The figure shows the probability of a single parent change occurring as a result of a degraded route directly effecting further parent changes in the immediate network. The immediate network constitutes all neighbouring sensors to the sensor that experienced a parent change. The probability of a parent change effecting a further change, two further changes, and three further changes is displayed. This is an important metric to observe as it gives direct insight into the possibility and scale of a domino effect occurring as a result of a parent change. From the graph we notice a characteristic shape in the data. The characteristic shape in the curve seen in figure 8 is related to the curve seen in figure 9. Figure 9 shows the number of sensors that act as a leaf as a percentage of the overall number of sensors in the network. Leaf sensors collect and send data through the network but do not act as a forwarding route for other sensors. As a result, when a leaf changes parent the knock-on effects to the rest of the network is minimal. It is then expected that when a higher percentage of leaf sensors is presented in the network the average knock on effect of a parent change is reduced. This relationship is clearly seen in figures 8 and 9. Figure 8 shows that in such large networks the probability of a parent change causing another is quite high.

In some cases there is over a 30% likelihood of a single route change, with a 10% chance of two route changes occurring in the network as a result of a parent change. The solution using NHs exhibits a distinct advantage in minimising the likelihood of further route changes, drastically reducing the chances of a route change to below 5%. With lower probability of a route change effecting further changes the extent of the domino effect is reduced. This again presents a significant result that has far-reaching consequences to route stability in the network. The benefits of reducing the domino effect are seen as fewer overall parent changes in the network. These benefits are clearly identified in figure 10, where the total number of parent changes for both solutions are compared. Total number of route changes occurring in the network is an important metric to compare and has an intrinsic connection with network stability. Each additional route change in the network can be seen as increased instability. The figure shows marked improvement in stability for the NHs solution particularly at larger scales. The larger scale network allows for higher dynamics making route stability even harder to maintain. The figure highlights a valuable insight into the growing issue of instability with scale.

Reducing the number of parent changes in the network is critical but does not constitute the whole story with regard to network stability. Network stability is discussed further as we examine the results from experiments conducted on a real testbed.

D. Testbed Environment

Experiments were conducted over a public experimental testbed. Wisebed [34], [35] is a testbed spread over nine facilities across Europe. Experiments described in this paper were conducted on the testbed situated in the University of Lübeck, Germany. This testbed facilitates access to 54 sensors in a network deployed in a working office environment. Figure 11 shows the positioning of the sensors within the building. Using this testbed gives an insight into how each solution copes in a real dynamic office environment, remaining in line with the BEMS application scenario. The working office will create a dynamic communications environment with moving bodies and open and closing doors. The different spaces such as corridors with private and open-plan offices will provide a number of communication challenges not available in a simulation environment. Experiments were conducted over a range of densities on the testbed. As the geographical location of the sensors in the testbed environment are fixed, the physical density and scale of the sensors cannot be altered. Adjusting the density and scale of each network must be controlled using the transmitter power on the sensor's radio chip. Varying the power of the radio will change the communications range of each sensor determining, within the physical bounds of the testbed, the neighbour set. Lowering the transmission power will reduce the parent set and increase the number of hops necessary to reach the root. A single power level is fixed for all nodes for an entire experiment. Table II shows how each power level maps to scale and density of the network.

The test environment presents a deployment in an office type building. WSN Deployments can be defined by a number of

Fig. 11. Sensor positions in Wisebed testbed.

 TABLE II
RELATIONSHIP BETWEEN RADIO POWER AND PHYSICAL NETWORK CHARACTERISTICS.

Power (dBm)	Wisebed Scale		Density
	ETX	ETX-NH	
-15	2.9	2.8	9.5
-10	2.4	2.3	14.7
-7	1.8	2.0	19.3
-5	1.6	1.7	21.5
-3	1.6	1.6	23.5
0	1.3	1.6	25.2

Scale: Average number of hops to the root.

Density: Average number of neighbours to a sensor.

physical attributes of the network. These include the number and density of sensors, scale of the network in terms of “hops” and the channel dynamics experienced in the environment. The Wisebed testbed exhibits channel dynamics that would be considered typical for either office or domestic environments. The number of sensors participating in experiments is maintained at the maximum, as this represents a medium sized deployment from which we can infer meaningful information regarding smaller and larger deployments. To emulate different types of networks the transmission power on each sensor is varied. The power changes give rise to new densities and scales for the physical network. Given that experiments are performed over a range of densities and scale the results are generalisable for many different office and domestic deployment types.

E. Measuring Stability

Paxon [36] declares that stability of a network can be measured using two independent metrics: the *persistence* and *prevalence* of the routes used within the network. Persistence is defined as how long a route is likely to endure before changing. Prevalence pertains to how likely a particular route is observed compared to other chosen routes. This has significance to network predictability. A stable network is deemed to be one that maintains routes, which are both prevalent and persistent. In our experiments the time of each route change was monitored, so we can easily determine the persistence of routes. Let l_r be the duration of route r and let N be the number of observed routes. Persistence can be defined as the average route duration:

$$R = \frac{\sum l_r}{N}. \quad (2)$$

Fig. 12. Route persistence in the Wisebed testbed over range of densities and scale.

Fig. 13. Plot showing dominant route prevalence in Wisebed testbed.

Prevalence looks at how often a particular dominant route is chosen. Each delivered packet will constitute a single chosen path. Let $k_r(n)$ be the number of times the dominant route r is chosen for sensor n and let $N(n)$ be the total number of packets delivered for sensor n . Then the prevalence of the dominant route for sensor n becomes:

$$\pi_{dom}(n) = k_r(n)/N(n). \quad (3)$$

The prevalence for the network can be calculated as the average prevalence of the dominant routes for all sensors:

$$\pi_{dom} = \frac{\sum \pi_{dom}(n)}{NumSensors}. \quad (4)$$

F. Stability in the Network

When discussing results from the real testbed we will examine stability of the network in terms of persistence and prevalence, as these factors can effect the provision of services and overall routing performance. Figure 12 shows a comparison of route persistence in the network. Figure 13 shows the prevalence of dominant routes in the network. Route persistence is measured in minutes and the figure shows the average route persistence in the network for each experiment. For all experiments conducted, NHs present a solution with

improved route persistence, with improvements ranging from 19 to 40%. The figure shows an increasing trend in route persistence as the radio power increases and the scale of the network is reduced. This behaviour is expected as the average number of hops to the destination decreases allowing for less dynamics in the routing structure. As radio power increases more sensors maintain a direct connection to the root directly. Sensors will then choose this direct route to the root. This can account for the dramatic rise in persistence from -7dBm , where 26% of sensors in the network use a direct route to the root, to -3dBm , where 42% of the sensors use a direct route to the root. This increasing trend however is broken and persistence is reduced with radio power set to 0dBm . This effect is due to the additional number of sensors that now have access to the root sensor at 0dBm transmission power. Some of these additional sensors, however, have a less reliable connection with the root sensor owing to a "longer link". This can lead to sensors choosing the longer, direct link to the root but having periodic difficulties resulting in route changes. This effect is more acute for the NHs solution. This is due to increased neighbourhood effect on sensors close to the root at high densities. The root node remains the dominant route choice, but the close neighbouring sensors become a more attractive option and are chosen more often than that with the ETX solution. This "last hop" effect is only apparent for nodes with a connection to the root sensor. Further investigation is necessary to develop additional insight and indeed a solution to this stability nuance.

Prevalence is a measure of how often a dominant route is observed in the network. In a WSN a perspective from local routing knowledge is considered. In this case the prevalence is determined by how often a sensor uses a dominant parent to transfer data. Experiments were conducted on the Wisebed testbed such that the route prevalence of each solution can be established. A comparison of the route prevalence is presented in figure 13. Again, the figure shows a general trend of the dominant route becoming more prevalent as the power increases and scale decreases. This is in spite of increasing density, which presents greater choice of routes to a sensor. The increased power, however, creates strong reliable links between sensors that are favoured, increasing prevalence. A dip in route prevalence is observed for experiments with radio power set to 0dBm . This again is a resultant from the "last hop" effect as sensors choose longer links to the root that become disturbed more frequently. The figure indicates a significant improvement in prevalence by applying the NHs technique. The stability in terms of both persistence and prevalence are improved using NHs. This is a significant result as route stability is important for both the network performance and the provision of higher level network functions.

G. Maintaining Goals in the Network

Increasing stability is of limited use if it is at the detriment of the overall evaluation metrics. We investigate the effect of the NHs framework with regard to reliability, latency and energy efficiency. These evaluation metrics constitute the essential network needs of a BEMS type application. We

Fig. 14. Comparison of packet reliability under real word dynamics in the Wisebed testbed.

compare ETX-NH and ETX in order to confirm that the use of NHs does not compromise delivery of these metrics.

Reliability is measured as the number of packets received at the root over the number of packets sent to the destination, and is displayed as a percentage. Figure 14 shows how the NHs solution compares the current ETX solution in terms of reliability of packet delivery. The figure shows the reliability of application packets in the Wisebed testbed environment. Both solutions show high levels of reliability in the network with all experiments providing a 96% packet delivery rate or better on packets sent. High reliability is a result of the packet acknowledgement and retransmission policy in RPL. The choice of route and hence routing metric is still important in regard to the choice of reliable routes. This is reflected in figure 14 with the differing performance between the ETX and NHs solutions. If we consider the experiments at 0dBm as a special case, the figure shows the ETX solution outperforming the NHs solution at smaller scales, and the NHs solution outperforming ETX as the scale increases. With these experiments we prove that there are minimal negative effects on reliability at very low scale with improving and even clear benefits to reliability with the use of NHs as scale increases. While it is difficult to discern at the office scale, the benefits of increased stability on reliability should become more clear as the scale increases to larger building wide networks. Future experiments will consider larger testbed environments to gain an understanding of reliability at larger scale. At 0dBm we see the performance of the ETX solution degrading slightly. This is due to the sensors choosing longer links to the root at higher power. The NHs solution exhibits a tighter bound on experiment results in general, over the whole range of experiments. This lends itself to a more predictable network and a greater definition of achievable QoS in the network. A similar result is seen when comparing latency. Both solutions remain competitive for all experiments. Latency performance increases as the scale of the network is reduced and more sensors have access to the root sensor. As the scale is reduced the current ETX solution becomes more competitive and achieves better latency. This is particularly noticeable in

Fig. 15. Comparison of packet latency for the Wisebed testbed.

Fig. 16. Showing the percentage of traffic carried by the most heavily burdened sensors in the network for ETX and ETX-NH.

the experiments using the 0dBm radio power, owing to more direct links to the root being chosen. It is clear from figure 15 that ETX-NH maintains a similar standard to that shown by ETX with regard to network latency.

Power use in WSNs remains a large factor for any successful communications method. Radio use on a WSN constitutes the majority of the energy spend [37], with active receiving and particularly transmitting consuming large amounts of energy. Excessive energy spend on sensors become an issue as these sensors may exhaust power reserves prematurely compromising the network. This can occur due to a single or small set of sensors carrying the majority of the traffic through the network. The use of NHs, while maintaining a network wide load consistent with current practices, can achieve a greater spread of traffic through the network reducing the communication spend on the most heavily burdened sensors [38]. NHs identifies and uses more forwarding routes than ETX by itself, allowing sensors to choose a more diverse set of routes. Figure 16 shows a comparison of the traffic carried by the most heavily burdened sensors over the network for the Wisebed testbed with radio power set to -10 dBm. The figure shows the load carried over the ten most heavily burdened sensors in the network as a percentage of the overall load carried in the network. Each point represents the accumulated traffic over the

stated number of sensors with the right most point representing the load carried over all ten of the most heavily burdened sensors. Using the NHs solution the communication load seen by these sensors is reduced. The sensor carrying the highest percentage of traffic benefits greatly with the proportion of traffic carried reduced from 46.8 to 28.7%. NHs also shows a greater diversity of routes with an average of 17 nodes carrying a load in the network compared to 12 when using the current ETX implementation.

NHs can show marginal improvements in a real world testbed with regard to reliability and latency at greater scales. Reliability appears more stable when using NHs for a given network allowing for a more predictable network. A more tangible improvement, however, can be seen with regard to power management and maintaining better power distribution over the network as a whole. This is important for WSNs as a complete network can be maintained for longer without damaging other network requirements.

H. Link Correlation

Shrinivasan *et al.* [27] discuss the concept of correlated link effects. Correlated link effects is a facet of wireless communications which states that links in similar geographical locations will experience similar disturbances. This has consequences for the NH framework. If a sensor experiences a disturbance on the link between itself and its parent, there is a chance the same disturbance will effect the links between the sensor and all of its neighbours damaging the quality of the intended failover route. If we revisit figure 4 and consider full link correlation, i.e. if the link from sensor A to its parent is disturbed, so too is the link to its neighbour and equally so for sensor B. Comparing current practice with NHs, we can conclude that both solutions result in the same number of route changes and eventual route structure. For this small example NHs performs at least on par with current practice, but has the capacity to provide improved stability with reduced link correlation. However, further experimentation is necessary to determine the full extent of link correlation on the framework in a wider network.

VII. CONCLUSIONS

In this paper we introduce a novel framework based on NHs. NHs will utilise all metric information available to a sensor in order to make better routing decisions. These routing decisions are not only based on the quality of the route used, but also the ability of the route to handle failover should the current route become unavailable. The aim of this method is to present a more stable platform to better deliver higher level services within the network, whilst maintaining the network goals. This approach is novel in a number of important regards:

- The issue of instability is addressed without compromising the underlying goals of the network.
- It shifts the paradigm from route choice based on current best path to highlighting good routes with better failover options.
- It amalgamates both route quality and failover capabilities into a single metric, which can be easily used for routing decisions.

The NHs framework is implemented using the RPL protocol, but is extensible to any tree type routing protocol. We prove through experiment that the two conditions for the framework described in section II are met. We test the neighbourhood metric framework against the current RPL implementation using ETX as a metric. We use both simulated and real communication environments to evaluate the proposed solution. For the test cases shown, the NHs technique shows increasing improvement in terms of network stability over the current ETX solution in simulation as network scale increases. Improvements in both route persistence and route prevalence are also seen in a real testbed environment. The framework provides a trade-off between choosing routes that better maintain network stability and routes that consist of a single best path. This does not, however, manifest itself in reduced network performance with the NHs solution remaining competitive with the current state of the art. Positive knock-on effects on the performance of the network from using NHs are particularly noticeable for energy use with a more even spread of energy use achieved over the network. The advantages of NHs come with little communication overhead as the framework is based on previously unused information already available to a sensor, negating the need to generate extra traffic.

REFERENCES

- [1] S. H. Lee, S. Lee, H. Song, and H. S. Lee, "Wireless sensor network design for tactical military applications: remote large-scale environments," in *Proceedings of the 28th IEEE Conference on Military Communications*, ser. MILCOM'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 911–917.
- [2] H. Furtado and R. Trobec, "Applications of wireless sensors in medicine," in *MIPRO, 2011 Proceedings of the 34th International Convention*, May 2011, pp. 257–261.
- [3] V. L. Erickson, M. A. Carreira-Perpiñán, and A. E. Cerpa, "Observe: Occupancy-based system for efficient reduction of hvac energy," in *IPSN*, X. D. Koutsoukos, K. Langendoen, G. J. Pottie, and V. Raghunathan, Eds. IEEE, pp. 258–269.
- [4] D. T. Delaney, G. M. P. O'Hare, and A. G. Ruzzelli, "Evaluation of energy-efficiency in lighting systems using sensor networks," ser. BuildSys '09. New York, NY, USA: ACM, 2009, pp. 61–66.
- [5] A. Cerpa, J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin, "Statistical model of lossy links in wireless sensor networks," in *IPSN*, 2005.
- [6] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet routing instability," in *Proceedings of the ACM SIGCOMM '97 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '97. New York, NY, USA: ACM, 1997, pp. 115–126.
- [7] K. Ramachandran, I. Sheriff, E. Belding, and K. Almeroth, "Routing stability in static wireless mesh networks," in *Proceedings of the 8th International Conference on Passive and Active Network Measurement*, ser. PAM'07. Berlin, Heidelberg: Springer, 2007, pp. 73–83.
- [8] D. Passos, D. V. Teixeira, D. C. Muchaluat-saade, L. C. S. Magalhães, and C. a. V. N. Albuquerque, *Mesh Network Performance Measurements*. Citeseer, 2006, pp. 1–8.
- [9] L. Chen and C. wei Lee, "Neighbor stability routing in MANETs," in *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 4, March 2005, pp. 1964–1969.
- [10] A. Woo, T. Tong, and D. Culler, "Taming the underlying challenges of reliable multihop routing in sensor networks," in *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems*, ser. SenSys '03. New York, NY, USA: ACM, 2003, pp. 14–27.
- [11] R. Tynan, D. Marsh, D. O'Kane, and G. M. P. O'Hare, "Intelligent agents for wireless sensor networks," in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '05. New York, NY, USA: ACM, 2005, pp. 1179–1180.
- [12] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection tree protocol," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New York, NY, USA: ACM, 2009, pp. 1–14.
- [13] Winter and *et al.*, "RPL: Ipv6 routing protocol for low power and lossy networks," IETF-ROLL, Tech. Rep., 2012.
- [14] D. Tian and N. D. Georganas, "Energy efficient routing with guaranteed delivery in wireless sensor networks," in *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, vol. 3, 2003, pp. 1923–1929.
- [15] T. A. Babbitt, C. Morrell, B. K. Szymanski, and J. W. Branch, "Self-selecting reliable paths for wireless sensor network routing," *Computer Communications*, vol. 31, no. 16, pp. 3799–3809, 2008, performance Evaluation of Communication Networks (SPECTS 2007).
- [16] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong, "TAG: A tiny aggregation service for ad-hoc sensor networks," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 131–146, Dec. 2002.
- [17] P. Edara, A. Limaye, and K. Ramamritham, "Asynchronous in-network prediction: Efficient aggregation in sensor networks," *ACM Trans. Sen. Netw.*, vol. 4, no. 4, pp. 25:1–25:34, Sep. 2008.
- [18] G. Wei, Y. Ling, B. Guo, B. Xiao, and A. V. Vasilakos, "Prediction-based data aggregation in wireless sensor networks: Combining grey model and kalman filter," *Computer Communications*, vol. 34, no. 6, pp. 793–802, 2011.
- [19] S. Kumar and S. Chauhan, "Article: A survey on scheduling algorithms for wireless sensor networks," *International Journal of Computer Applications*, vol. 20, no. 5, pp. 7–13, April 2011, published by Foundation of Computer Science.
- [20] S. Gandham, M. Dawande, and R. Prakash, "Link scheduling in wireless sensor networks: Distributed edge-coloring revisited," *Journal of Parallel and Distributed Computing*, vol. 68, no. 8, pp. 1122–1134, 2008.
- [21] G. Lu, N. Sadagopan, B. Krishnamachari, and A. Goel, "Delay efficient sleep scheduling in wireless sensor networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 4, March 2005, pp. 2470–2481 vol. 4.
- [22] L. H. Correia, D. F. Macedo, A. L. dos Santos, A. A. Loureiro, and J. M. S. Nogueira, "Transmission power control techniques for wireless sensor networks," *Computer Networks*, vol. 51, no. 17, pp. 4765–4779, 2007.
- [23] S. Lin, J. Zhang, G. Zhou, L. Gu, J. A. Stankovic, and T. He, "ATPC: Adaptive transmission power control for wireless sensor networks," in *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*, ser. SenSys '06. New York, NY, USA: ACM, 2006, pp. 223–236.
- [24] J. Kim, S. Chang, and Y. Kwon, "Odtpc: On-demand transmission power control for wireless sensor networks," in *Information Networking, 2008. ICOIN 2008. International Conference on*, Jan 2008, pp. 1–5.
- [25] A. Becher, O. Landsiedel, G. Kunz, and K. Wehrle, "Towards Short-Term Wireless Link Quality Estimation," in *HotEmNets 2008*, Charlottesville, Virginia, USA, Jun. 2008.
- [26] A. Basalamah, S. M. Kim, S. Guo, T. He, and Y. Tobe, "Link correlation aware opportunistic routing," in *INFOCOM, 2012 Proceedings IEEE*, 2012, pp. 3036–3040.
- [27] K. Srinivasan, M. Jain, J. I. Choi, T. Azim, E. S. Kim, P. Levis, and B. Krishnamachari, "The κ factor: inferring protocol performance using inter-link reception correlation," in *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '10. New York, NY, USA: ACM, 2010, pp. 317–328.
- [28] S. Lin, G. Zhou, K. Whitehouse, Y. Wu, J. A. Stankovic, and T. He, "Towards stable network performance in wireless sensor networks," in *Proceedings of the 2009 30th IEEE Real-Time Systems Symposium*, ser. RTSS '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 227–237.
- [29] A. Kamthe, M. A. Carreira-Perpiñán, and A. E. Cerpa, "M&M: multi-level Markov model for wireless link simulations," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New York, NY, USA: ACM, 2009, pp. 57–70.
- [30] "XBOW sensor motes specifications," 2008, <http://www.xbow.com>.
- [31] P. Levis, S. Madden, J. Polastre, R. Szewczyk, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, and D. Culler, "TinyOS: An operating system for sensor networks," in *Ambient Intelligence*. Springer, 2004.
- [32] O. Gnawali and P. Levis, "The minimum rank with hysteresis objective function," IETF-ROLL, Tech. Rep., 2012.
- [33] P. Levis, N. Lee, M. Welsh, and D. Culler, "Tossim: Accurate and scalable simulation of entire tinyos applications," in *Proceedings of the*

1st International Conference on Embedded Networked Sensor Systems, ser. SenSys '03. New York, NY, USA: ACM, 2003, pp. 126–137.

- [34] G. Coulson, B. Porter, I. Chatzigiannakis, C. Koninis, S. Fischer, D. Pfisterer, D. Bimschas, T. Braun, P. Hurni, M. Anwender, G. Wagenknecht, S. P. Fekete, A. Kröller, and T. Baumgartner, “Flexible experimentation in wireless sensor networks,” *Commun. ACM*, vol. 55, pp. 82–90, Jan. 2012.
- [35] T. Baumgartner, I. Chatzigiannakis, M. Danckwardt, C. Koninis, A. Kröller, G. Mylonas, D. Pfisterer, and B. Porter, “Virtualising testbeds to support large-scale reconfigurable experimental facilities,” in *Wireless Sensor Networks*, ser. Lecture Notes in Computer Science, J. Silva, B. Krishnamachari, and F. Boavida, Eds. Springer Berlin Heidelberg, 2010, vol. 5970, pp. 210–223.
- [36] V. Paxson, “End-to-end routing behavior in the internet,” *Networking, IEEE/ACM Transactions on*, vol. 5, no. 5, pp. 601–615, oct 1997.
- [37] O. Landsiedel, K. Wehrle, and S. Gotz, “Accurate prediction of power consumption in sensor networks,” in *Proceedings of the 2nd IEEE Workshop on Embedded Networked Sensors*, ser. EmNets '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 37–44.
- [38] D. T. Delaney, L. Xu, and G. M. P. O'Hare, “Spreading the load in a tree type routing structure,” in *Proceedings of the IEEE 22nd International Conference on Computer Communications and Networks (ICCCN 2013)*. IEEE, 2013.

Declan T. Delaney received his B.A. in computer engineering and communications from the National University of Ireland, Maynooth, Ireland, in 2006. After working with Erickson telecommunications as a network design consultant Declan chose to begin research in WSN design at the School of Computer Science and Informatics, University College Dublin, Ireland, where he remains at present as a PhD candidate.

Russell John Higgs was born in Essex, U.K., in 1958. He received the B.A. and Ph.D. degrees in pure mathematics from the University of Liverpool, U.K., in 1980 and 1985, respectively. He was an Associate Professor with the Department of Mathematics, University of Pittsburgh, Pittsburgh, PA, from 1983 to 1985. He was then an Assistant Lecturer (1985/1989), College Lecturer (1989/2002), and is currently a Senior Lecturer (2002/present) with the School of Mathematical Sciences, University College Dublin, Ireland. He is also a member of the Claude Shannon Institute, Ireland, as well as being a past President of the Irish Mathematical Society. His research interests include projective representations, coding theory, network coding, and quantum computing. Dr. Higgs received the 2003 Presidents Teaching Award from University College Dublin.

Gregory O'Hare completed his studies at the University of Ulster graduating with a B.Sc, M.Sc and Ph.D. He held the position of Head of the Department of Computer Science at University College Dublin (UCD) 2001-2004. Prior to joining UCD he has been on the Faculty of the University of Central Lancashire (1984-86) and the University of Manchester (1986-1996). He is a Associate Professor within the School of Computer Science and Informatics at UCD. He has published over 415 refereed publications in Journals and International Conferences and has won significant grant income (ca 28.00M). He is one of the Principal Investigators and founders of the Science Foundation Ireland funded (16.4M) Centre for Science and Engineering Technologies (CSET) entitled CLARITY: The Centre for Sensor Web Technologies (2008-2013). He is currently the Director of the Earth Institute at University College Dublin.