

Perfect Link Routing for Energy Efficient Forwarding in Geographic Routing

A. Kleerekoper^{a,*}, N. P. Filer^a

^a*School of Computer Science, The University of Manchester, Oxford Road, M13 9PL, Manchester, UK*

Abstract

Geographic routing has been widely advocated for use with multihop ad hoc and sensor networks because of its scalability and use of only local information. These types of networks typically have lossy links where the instantaneous quality of a wireless link can vary significantly presenting a trade-off between hop length and link quality. In this paper we revisit the question of energy efficient geographic routing for such networks and argue in favour of Perfect Link Routing, an extreme form of blacklisting with a fall-back option.

Existing research has favoured cost-based methods where all links are considered for routing. We argue, however, that a discontinuity exists between the cost of perfect links (those with virtually guaranteed delivery) and other links. This is based on a more careful use of acknowledgements which we suggest ought to be considered a function of individual links. Revisiting the original analysis we find that for energy efficiency, perfect links should be favoured except in low-density networks where such a scheme leads to very poor delivery rates. A hybrid approach is proposed which we call Perfect Link Routing and this method is shown to outperform alternatives for a number of ARQ schemes.

Keywords: Ad Hoc Networks; Energy Efficiency; Geographic Routing; Wireless Sensor Networks

1. Introduction

Multihop wireless networks, such as ad hoc and sensor networks, have received significant research attention in the past decade. Since these networks are typically constructed from low-power, resource-constrained devices, the research effort has been largely directed at designing efficient protocols. Because the devices are resource constrained, they cannot usually be expected to store large routing tables and in large networks the energy cost of calculating and

*Corresponding author

Email addresses: `kleereka@cs.man.ac.uk` (A. Kleerekoper), `nick@cs.man.ac.uk` (N. P. Filer)




Figure 1: Greedy forwarding can fail when a void is encountered and techniques are needed to route around such “holes”. Figure from [1]

maintaining long routes is prohibitive. Therefore, localised routing protocols are preferred especially for networks with significant point-to-point traffic.

Prime among these localised protocols is the family of geographic routing (also referred to as position-based routing). They rest on the assumption that nodes are aware of their own location either through GPS or some other form of localisation and can share this information with their one hop neighbours. Routing is then performed primarily through greedy forwarding with each node forwarding the packet to a node closer to the destination with the expectation that the packet will eventually arrive.

There are two significant problems with geographic routing based on greedy forwarding. The first is that packets can become stuck if a node cannot find a neighbour closer than itself to the packet destination as illustrated in Fig. 1. Significant research effort has been applied to devise methods for finding routes around such “voids” as surveyed in [1].

Perhaps a bigger problem, and the one that this paper focuses on, is that the majority of geographic routing protocols adopt the unit disk graph model of wireless links. This model states that, so long as two nodes are within a fixed distance of each other, every transmission from one will be received perfectly by the other unless there is interference. In effect, the model posits two regions around a transmitting node: a connected region in which the packet reception rate (PRR) is close to 100% and a disconnected region in which the PRR is close to 0%. The dividing line between these two regions is a fixed distance.

The reality, however, is that low-power links do not behave this way. It has been known for some time that there are in fact three distinct regions with a large transitional region between the connected and disconnected regions [2, 3, 4]. In the transitional region there is no definite relationship between distance and PRR and a node close to the transmitter may have a low rate while a node far away may have a high reception rate.

If simple greedy forwarding is used which only takes the geographic progress of the link into account then nodes far from the transmitting node are preferred and these may well have poor performance requiring numerous retransmissions and resulting in low energy efficiency. As a recent survey notes, “this observation by itself annihilates all geographic routing protocol solutions and has been largely overseen” [5].

In fact there has been some research directed at adapting the geographic routing to the reality of the wireless channel. One of the main contributions to this was made by Seada *et al.* in two papers [6, 7]. They recognised that when selecting a neighbour to forward packets to it is essential to consider the PRR of the link to that neighbour as well as how much closer it is to the final destination. They analysed and compared a number of different methods for incorporating PRR into neighbour selection and concluded that the optimal method is to find the neighbour with the largest product of PRR and distance, a metric they called PRR_{xd}.

In their original work they considered two scenarios: networks using Automatic Repeat reQuest (ARQ) to guarantee end-to-end delivery and networks without ARQ. In this paper we extend their analysis in three important ways. Firstly, we consider a third scenario which uses negative acknowledgements to improve end-to-end delivery rate. Secondly, we incorporate into the analysis the probability that control packets may fail. Thirdly, we argue that the use of acknowledgements should be a function of the link quality leading to an updated analysis with variable energy costs.

Our analysis shows that high quality links should be preferred and, in the case of ARQ, the routing metric ought to be $PRR^2d/(1 + PRR)$ which favours high quality links over long ones. Moreover, we argue that in all cases an extreme form of reception-based blacklisting, which we call Perfect Link Routing (PLR), is optimal in terms of energy efficiency. In PLR only links with 100% PRR are considered for routing. This obviously leads to lower end-to-end delivery rates at low densities and to compensate we suggest a simple hybrid method whereby PLR uses the PRR_{xd} metric where no perfect links are found.

This work is a major extension of our early discussion [8] on this topic. Our main contributions are:

- We consider the possibility that control packets may fail and extend the original analysis, showing that for the ARQ scenario a new metric, $PRR^2d/(1 + PRR)$, is optimal (Section 4.1)
- We analyse a new, third, scenario which uses negative acknowledgements and show that in this scenario high quality links must be preferred (Section 4.2)

- We argue that the use of both positive and negative acknowledgements should be considered a function of the link quality and not as a network-wide decision. This leads to a variable link energy cost and we analyse the energy efficiency of greedy forwarding using this new cost function (Section 5).
- We show that a Perfect Link Routing (PLR) is more energy efficient than PRRxd in all three scenarios (Section 6)
- We consider the range of network densities for which PLR is appropriate and show that modifying PLR into a simple hybrid protocol captures the benefits of both methods (Section 7)
- We examine how the efficiency of PLR varies when data packets become larger relative to control packets (Section 8)
- We further the case for PLR based on the temporal nature of wireless links (Section 9)

2. Related Work

Stojmenovic gives an important survey of geographic routing, listing 20 geographic routing methods [9]. He identifies the main advantage of this type of routing as increased scalability compared to path-based methods. The underlying greedy forwarding element means that routing decisions are localised and therefore largely independent of the total network size. Moreover, changes in one part of the network due to node failure or mobility only affect a small part of the network and do not require the recalculation of entire paths.

Stojmenovic further identified that there are two elements to geographic routing: greedy mode and recovery mode. In the greedy mode the packet is forwarded towards the destination using local information and decisions. However, greedy methods can get stuck at local maxima which can happen when a node that receives the packet cannot find a forwarder. When this happens recovery mode techniques are needed to find a route around the empty space (called a void) in the network. We focus on the greedy mode which is the main operating mode.

Although the basic principle of greedy forwarding is to ensure that each packet makes progress towards the destination, there are different methods in the literature for determining which forwarder represents the optimal choice [10, 11, 13, 14, 12]. In this paper we follow [6] and adopt the method proposed by Finn called “Cartesian Routing” which takes progress to be the difference between the source-destination distance and the relay-destination distance.

The traditional approach to geographic routing is to utilise the unit disk graph (UDG) model, also known as “perfect-within- R ”, in which every node has a circular transmission range of radius R . In UDG the source has a perfect link to any node inside the circle and no link to any node outside it. This simplified model is very attractive because it enables rigorous mathematical




Figure 2: The Packet Reception Rate (PRR) of low-power links does not depend entirely on distance and in the transitional region two nodes at the same distance from the transmitter may have very different rates. Figure from [4]

analysis. Studies have shown, however, that its oversimplification results in the poor performance of geographic routing protocols that attempt to maximise the distance of transmissions [2, 3, 15, 16].

Geographic routing research originally made heavy use of the simplified UDG model but more recently this model has been shown to be far from accurate and new models have been proposed. Woo *et al.* conducted extensive experiments and identified three regions around a transmitting node, rather than the two used in the UDG model [2]. Close to the node the packet reception rate (PRR) was almost always 100% and far away from the node it was almost always 0%, corresponding to inside and outside the transmission range in UDG. However, between these two regions exists a transitional region in which the PRR can vary significantly between links of the same length. In their important experiments, Zhao and Govindan found the same three regions noting that the transitional region can be a large proportion of the total communication range [3]. They also investigated the spatial correlation of link quality, concluding that it is reasonable to assume that the PRR of each link is independent of the PRR of nearby links. The three areas surrounding a transmitting node, and the PRR of links at different distances, are illustrated in Figure 2.

Zhou *et al.* were among the first to investigate the impact of the realistic wireless channel behaviour on routing performance [15, 17]. Based on empirical data they developed a “Radio Irregularity Model” (RIM) which includes a parameter, DOI which stands for *degree of irregularity*. The path loss in a given direction is calculated by modifying the path loss by a randomly selected proportion of the DOI. The larger the DOI the more variation there is in PRR values

at different locations around the transmitter. Simulation results using the RIM model show that geographic routing performs poorly as the channel becomes more irregular and less idealised. Similar results were reported by Couto *et al.* [16].

Zuniga and Krishnamachari made use of the extensive literature regarding wireless channel behaviour for cellular networks to develop an accurate model of the behaviour of wireless channels for low powered networks [4]. Their model is also based on the log-normal shadowing model and is similar, in many ways, to RIM and is described in more detail in Section 6.

Once the problems with traditional geographic routing strategies were identified and became more accepted, research efforts were directed towards developing techniques that were able to cope with the lossy nature of wireless links.

Both Woo *et al.* [15, 17] and Couto *et al.* [16] proposed replacing geographic routing metrics with a new metric based on link quality. They both proposed the same metric with they called “Minimum Transmission” (MT) and “Expected Transmission” (ETX) respectively. The ETX of a link is the expected number of transmissions including retransmissions (but not acknowledgements) that would be needed to successfully transmit a packet using that link. Sanchez and Ruiz proposed the “Expected Transmission Energy” (ETE) metric which considers the energy consumption of both the packets and the acknowledgements to produce a scheme which minimises the expected energy consumption of using a given link [18].

Seada *et al.* were the first to consider how best to modify geographic routing to account for lossy links [6]. Utilising the model developed by Zuniga and Krishnamachari they analysed the energy efficiency of a link both with and without Automatic Repeated Request (ARQ). In both cases they found that in order to maximise energy efficiency it is necessary to find the link with the highest product of PRR and length. They therefore proposed a new metric for geographic routing which they called $PRR \times d$.

Stojmenovic *et al.* proposed a general approach for designing geographic routing protocols, namely that the chosen link should “minimise the ratio of cost over progress” [19]. The definition of cost depends on the assumptions that are made and how success is measured and the authors go on to investigate a number of options [20, 19, 21]. Lee *et al.* independently proposed the same framework [22, 23]. Both sets of researchers found that for energy efficiency the $PRR \times d$ method is optimal.

Although the initial work of Seada *et al.* has been followed up by some (for example [24]), it has unfortunately been largely overlooked in the field [5]. In fact, a recent, otherwise comprehensive, survey on geographic routing discusses the weakness of the UDG model without mentioning the transitional region or the $PRR \times d$ metric [25].

In this paper we extend the work of Seada *et al.* by considering the impact of the control packets. In the original analysis, which we will explain in more detail in the next section, the possibility of acknowledgements failing was not considered. In this paper we focus on the implications of the possible failure of control packets and the variable link cost when control packets are used most

efficiently. Our results initially provide a new routing metric for the ARQ case and later show that extreme reception-based blacklisting, which we call Perfect Link Routing (PLR), is optimal in terms of energy efficiency in all the scenarios we consider. PLR, however, results in low delivery rates in some cases and we therefore experiment with two simple hybrid methods and show that by modifying PLR to become a hybrid method the problem of low delivery rates is avoided.

3. Energy Efficiency Model

In this section we present the analysis of Seada *et al.* [6, 7] in more detail and point out the two simplifying assumptions they made that we relax in our analysis. Consider a chain of nodes where greedy forwarding is to be used to route packets between a source node at one end of the chain and a destination node at the other. Let n be the total number of packets generated by the source and r be the proportion of those packets that arrive successfully at the destination. Further let t be the total number of packets transmitted in order to achieve the delivery rate, r , and e be the total energy consumed performing each transmission. The energy efficiency of the network, ξ , is:

$$\xi = \frac{nr}{et} \quad (1)$$

At each hop, the optimal length of the hop is that distance, d_{opt} that maximises the efficiency of the network given a packet reception rate (PRR) model which maps transmission distances to PRR. The particular model used in this paper was derived by Zuniga and Krishnamachari [4], but any suitable model would suffice and the results would not alter significantly.

Let p_d be the expected PRR at a distance d . If no acknowledgements are used in the routing then the expected delivery rate is given by the product of the PRR of each hop. If there are h hops then the end-to-end delivery rate, r , is:

$$r = (p_d)^h \quad (2)$$

Since the delivery rate at each hop is not 100% some packets will traverse more hops towards the destination than others. Therefore, the number of packets that each node on the route must transmit varies and so the total number of packets transmitted is:

$$t = n \sum_{i=1}^h (p_d)^{i-1} = n \frac{(p_d)^h - 1}{(p_d - 1)} \quad (3)$$

Substituting equations (2) and (3) into equation (1), gives the efficiency for the no ARQ case:

$$\xi_{NoARQ} = \frac{(p_d)^h(1 - p_d)}{e(1 - (p_d)^h)} \quad (4)$$

This result shows that high quality (ie high PRR) links are required for high efficiency with no ARQ. As the quality of the links in the route fall the $(p_d)^h$ term in the numerator decreases rapidly whilst the $(1 - (p_d)^h)$ term in the denominator increases leading to an overall reduction in energy efficiency. This is intuitively correct since without ARQ, using links with low PRR leads to a large proportion of packets failing to reach the destination in which case any energy spent transmitting them along the route has been entirely wasted.

If ARQ is used, such that any failed transmissions are repeated, the end-to-end delivery rate is guaranteed to be 100%. The number of transmissions required to achieve this rate is inversely proportional to the PRR of the link. Let D be the distance between the source and destination and d is the distance traversed at each hop. The total number of transmissions is:

$$t = \frac{D}{d} \frac{n}{p_d} \quad (5)$$

Substituting equation (5) into (1) gives the efficiency in the ARQ case:

$$\xi_{ARQ} = \frac{dp_d}{eD} \quad (6)$$

Since the energy per hop, e , is considered to be constant and D is certainly constant, maximising the energy efficiency is equivalent to maximising the product of the link length, d , and the PRR of the link giving the PRRxd metric.

Whilst the analysis is complete in the no-ARQ case, in the case of ARQ, the impact of the acknowledgement packets have not been considered. The energy cost does not include them (a point we will discuss fully in Section 5), the number of packets transmitted per attempt does not include them either and the possibility that they may fail to be received correctly is not considered. Hence, the expected number of transmissions per hop is assumed to be $1/p_d$ which accounts for only one packet per attempt and does not allow for the possibility of a false negative whereby the data packet is successful but the acknowledgement fails. In the next section we consider both the acknowledgement packet and the possibility of its failure as well as considering the case where retransmissions are controlled by the use of negative acknowledgements.

4. Possibility of Acknowledgement Failure

In this section we extend the work of Seada *et al.* by considering the acknowledgement packet in the number of transmissions as well as the possibility that it may fail. In Section 5 we consider its energy cost as well. First we examine the case when ARQ is used and then subsequently consider the case of negative acknowledgements. In both cases we find that, just as in the no-ARQ case, high quality links should be preferred and we derive a new metric for the ARQ case.

State	Data	ACK	No. Packets	Probability
Successful	✓	✓	2	p^2
Correct Failure	✗	N/A	1	$1 - p$
Incorrect Failure	✓	✗	2	$p(1 - p)$

Table 1: When ARQ is used there are three possible outcomes from each transmission attempt.

4.1. ARQ with Failed Acknowledgements

When ARQ is used, a node assumes that all its transmissions fail unless it is explicitly informed otherwise by an acknowledgement packet (ACK) being sent back by the receiver. In lossy networks, just as the data packet may fail to be received correctly the ACK packet might also fail. In this case a false negative can occur whereby the data packet was received successfully but the transmitting node retransmits it believing it to have failed because the ACK failed. This gives rise to three states: (i) successful, (ii) correct failure and (iii) incorrect failure. Table 1 summarises the states and their properties. In the successful state both the data packet and the ACK are correctly received and a total of two packets are required. In the correct failure state the data packet fails in which case no ACK is sent and only one packet is required but another attempt at transmission is needed. In the incorrect failure state the data packet was received correctly but the ACK packet failed leading to the consumption of two packets and a retransmission.

Let p be the probability that the data packet is received successfully. The standard assumption is that for low-power networks the ACK packet will be approximately the same size as the data packet (though later in Section 8 we relax this assumption). The probability of a successful state is the probability that both the data and ACK packets are received successfully which is p^2 . The correct failure state has probability $(1 - p)$ and the incorrect failure state occurs with probability $p(1 - p)$.

For guaranteed end-to-end delivery the success state has to be encountered and the expected number of transmissions required to achieve that is $1/p^2$. The number of packets per attempt is not constant and the expected number of transmissions per hop, t_h , is:

$$t_h = \frac{1}{p^2} (2p^2 + (1 - p) + 2p(1 - p)) = \frac{1}{p^2}(1 + p) \quad (7)$$

The total number of transmissions depends on the number of hops, given by D/d and substituting into equation (1) gives:

$$\xi_{ARQ} = \frac{n}{eD} \frac{dp^2}{(1 + p)} \quad (8)$$

This gives rise to a new metric for maximising the energy efficiency of systems using ARQ. Unlike the PRRxd metric originally found, this new metric favours high quality links because of the p^2 term which makes it harder for long links

State	Data	NACK	No. Packets	Probability
Successful	✓	N/A	1	p
Correct Failure	✗	✓	2	$p(1 - p)$
Incorrect Failure	✗	✗	2	$(1 - p)^2$

Table 2: When negative acknowledgements are used there are three possible outcomes from each transmission attempt.

to compensate for a reduction in quality. In the PRRxd metric a link that was 10% longer but 10% lower quality would have the same score whereas when the possibility of false negatives are taken into account we find that a 10% reduction in link quality from a perfect link would need an increase of 17% in link length to give the same score.

4.2. Retransmissions with Negative Acknowledgements

In this subsection we consider a retransmission method which uses negative acknowledgements. In this method, which we refer to as the NACK method for convenience, transmitting nodes assume that their transmissions were successful unless they are explicitly told otherwise by the receiving node using a negative acknowledgement (NACK) packet. The NACK method may be more efficient in networks which usually utilise higher quality links since most transmissions are successful and therefore fewer NACK packets are required than ACK packets would be if ARQ is used. This method was recently suggested for use in WiMAX networks [26]. The NACK method requires the receiver to be expecting a data packet which is appropriate in many ad hoc networks where the MAC layer uses long preambles, explicit scheduling or RTS/CTS for collision avoidance.

As with the ARQ scheme, the NACK method has three states for each transmission, summarised in Table 2. If the data packet is received correctly then no NACK is needed and only a single packet is used. A correct failure is encountered when the data packet fails but the NACK succeeds. An incorrect failure occurs when both the data and NACK packets fail. Crucially, however, unlike the ARQ case when an incorrect failure is a false negative, in the NACK scheme an incorrect failure is a false positive. The transmitting node believes that the data packet was received successfully as it did not receive a NACK packet informing it otherwise. It therefore does not retransmit the data packet and the packet is lost. Even allowing for infinite transmissions the NACK scheme cannot guarantee perfect end-to-end delivery because of the false positives. We do not consider a scheme in which the receiving node is able to retransmit NACK packets and leave this for future work.

The probability of the three possible states are p , $p(1 - p)$ and $(1 - p)^2$ respectively. In order for a data packet to successfully reach the next hop the success state must occur before the incorrect failure state. The success state will be reached first if it is the result of the first attempt, or if the first attempt results in the correct failure state and the second attempt results in the success state, or if the first two attempts result in the correct failure state and the

second attempt results in the success state etc ad infinitum. The probability, P , of the packet reaching the next hop successfully is:

$$\begin{aligned} P &= p [1 + (p(1-p)) + (p(1-p))^2 + (p(1-p))^3 \dots] \\ &= \frac{p}{p^2 - p + 1} \end{aligned} \quad (9)$$

Following the logic applied to the no-ARQ case, the expected delivery rate is P^h and the total number of transmissions is:

$$t = n \sum_{i=1}^h P^{i-1} = n \frac{P^h - 1}{P - 1} \quad (10)$$

Substituting equation (10) into (1) gives the efficiency for the NACK case:

$$\xi_{NACK} = \frac{P^h(1-P)}{e(1-P^h)} \quad (11)$$

As with the no-ARQ case, to maximise energy efficiency P must be maximised because as P falls, P^h in the numerator tends to zero whilst $(1 - P^h)$ in the denominator increases. In order to maximise P , the PRR of the chosen link, p , must also be maximised. Therefore, as with the no-ARQ case, in the NACK case the most energy efficient method is to select only links with high PRR.

5. Retransmissions as a Function of Link Quality

In the previous section we analysed the energy efficiency of greedy forwarding when considering the possibility of the failure of acknowledgements. Our results suggested that because acknowledgements might fail and thereby either require extra, unnecessary packets in the ARQ case or result in dropped packets in the NACK case, it is preferable to select links with high PRR. In fact, for the no-ARQ and NACK cases, the only way to maximise energy efficiency is to select links with high PRR. For the ARQ case, however, our analysis showed that whilst PRR was a more important factor in maximising energy efficiency, there are situations in which selecting a longer link with a lower PRR would be more efficient than selecting a shorter link with a higher PRR.

In this section we make the case that the use of ARQ should be considered a function not only of the network but of the individual link quality as well. Up to now, and including in the original analysis from Seada *et al.*, the use of acknowledgements has been considered as a network wide decision so that if ARQ is used, it is used on every link. We argue that if a threshold of link quality is specified then only links with below-threshold qualities should use acknowledgements and all other links should not. Based on this we show that even in the ARQ case the energy efficiency of the network is maximised by selecting only high quality links.

Suppose that for a given link there is a minimum quality requirement, q , stating that at least $q\%$ of packets must be received successfully. This could be a constant for all links or a variable value depending on the expected number of hops between the source and the destination combined with a desired end-to-end delivery rate. If the PRR of the link is less than q then acknowledgements must be used to raise the effective PRR to at least q . Even if the desired link quality, q , is less than 100% meaning that not all unsuccessfully transmitted packets need to be retransmitted, acknowledgements must still be sent after all transmissions (in the ARQ case) so that the source node can monitor the effective PRR of the link. The source node may decide after an unsuccessful transmission that the effective PRR is high enough to not require a retransmission and accept that some packets are lost. Without acknowledgements after every successful transmission, however, it cannot know what the current effective PRR of the link is and so cannot make informed decisions about when to retransmit.

In contrast, if the PRR of the link is already known to be above q , then there is no need for acknowledgements. The source node is aware that some packets will not be received successfully but the proportion being lost is acceptable. The only purpose of sending the acknowledgements would be to allow the source to retransmit a failed packet. Since no retransmissions are ever needed on this link, transmitting acknowledgements is entirely wasteful. Therefore, the decision on whether to send an acknowledgement should be a function of the link quality.

The implication of not using acknowledgements for high quality links is twofold. Firstly, there is no possibility of the acknowledgement failing and secondly, the energy cost of a high quality link is significantly different to that of a low quality link.

In Section 4.1 we showed that the possibility of an ACK failing leads to false negatives and an increase in the expected number of transmissions needed for successful traversal of a hop. If above-threshold links do not use ACKs then this precludes the possibility of false negatives and leads to two qualitatively different efficiency measures.

For above-threshold links, the probability that the data packet is successfully received is p and so the expected number of transmissions needed to reach the threshold quality is q/p . But since, by definition, $p \geq q$ and there cannot be fewer than one transmission per attempt the total number of transmissions needed for an above-threshold link is just 1.

This compares to q/p^2 for below-threshold links where a packet is only considered to have been successfully received if both the data and the ACK packet are correctly received. In fact, comparing the energy efficiency for the two types of links, equation (12), shows that for the above-threshold links the actual PRR of the link is irrelevant and only the distance traversed by the hop is important.

$$\xi_{ARQ} = \begin{cases} \frac{n}{eD} \frac{dp^2}{q(1+p)} & p < q \\ \frac{nd}{eD} & p \geq q \end{cases} \quad (12)$$

From equation (12) a case can be made for prioritising above-threshold links because as p increases towards q and q tends to 1 ($\lim p \rightarrow q \rightarrow 1$) the efficiency becomes:

$$\lim_{p \rightarrow q \rightarrow 1} \xi_{ARQ} = \frac{n}{eD} \frac{d}{2} \quad (13)$$

Equation (13) suggests that the maximum efficiency when using below-threshold links is half that when using above-threshold links. The only way to compensate for the loss of link quality would be to at least double the distance covered by the hop.

However, when the variable energy cost is considered the difference between above-threshold and below-threshold links increases further. For a below-threshold link, the energy cost to the transmitting node, E_{TX} , of an attempted transmission includes the cost of transmitting the data packet and the cost of waiting for the acknowledgement. Even if the data packet fails and so no acknowledgement is sent, the transmitting node must still listen to the channel because it has no way of knowing that no ACK is being sent. Since in low-power devices the energy consumed by the radio in idle listening is orders of magnitude greater than the energy consumed by processing a received packet, the energy cost of idle listening can be considered the same as the energy cost of successfully receiving a packet. Thus the energy cost to the transmitter for a low-quality link is:

$$E_{TX} = e_{tx} + e_{rx} \quad (14)$$

where e_{tx} and e_{rx} are the energy cost of transmitting and receiving a packet respectively. Note that at this stage we adopt the standard assumption that the data and ACK packets are of similar size. Later, in Section 8, we relax this assumption.

While the cost to the transmitter is independent of the success or failure of the data packet, the cost to the receiver, E_{RX} , is, since an ACK is only transmitted if the data packet succeeded. Thus the energy consumed by the receiver per attempted transmission is:

$$E_{RX} = \begin{cases} e_{rx} + e_{tx} & \text{received successfully} \\ e_{rx} & \text{otherwise} \end{cases} \quad (15)$$

If the PRR of the link is p then the expected number of attempts required for successful transmission is q/p . One of those transmissions will be successful and all the others will not be. Therefore, the total energy cost of the link, E_{link} , is:

$$\begin{aligned} E_{link} &= \frac{q}{p}(e_{tx} + e_{rx}) + \left(\frac{q}{p} - 1\right) e_{rx} + (e_{tx} + e_{rx}) \\ &= \frac{q}{p}(e_{tx} + 2e_{rx}) + e_{tx} \end{aligned} \quad (16)$$

Equation (16) applies only to the links with below-threshold quality since it assumes that every successful transmission is followed by an ACK. For high-quality links ACKs are never needed and so the transmitter does not need to listen for them and the receiver does not need to send them. Moreover, since the link quality is above the threshold the success of the data packet can be assumed and the data packet only ever needs to be sent once. Thus the energy cost of a high-quality link is simply $e_{tx} + e_{rx}$. The energy cost of a link can therefore be summarised as:

$$E_{link} = \begin{cases} \frac{q}{p}(e_{tx} + 2e_{rx}) + e_{tx} & p < q \\ e_{tx} + e_{rx} & p \geq q \end{cases} \quad (17)$$

Substituting the link energy cost into the equation for energy efficiency in the ARQ case (equation (12)) gives:

$$\xi_{ARQ} = \begin{cases} \frac{n}{D} \frac{dp^2}{q(1+p)} \frac{1}{\frac{q}{p}(e_{tx} + 2e_{rx}) + e_{tx}} & p < q \\ \frac{n}{D} \frac{d}{e_{tx} + e_{rx}} & p \geq q \end{cases} \quad (18)$$

The difference in efficiency between a high quality and any other link is stark. Now in the limit ($\lim p \rightarrow q \rightarrow 1$) the efficiency of the below-threshold link tends to:

$$\lim_{p \rightarrow q \rightarrow 1} \xi_{ARQ} = \frac{n}{D} \frac{d}{4(e_{tx} + e_{rx})} \quad (19)$$

Equation (19) shows that if a below-threshold quality link is chosen then in order for the energy efficiency of the network to be the same as a network using only above-threshold links, the distance traversed by the below-threshold link must be at least four times that of the above-threshold link. Since this is unlikely for most networks it is therefore clear that to maximise energy efficiency only above-threshold links should be used for routing. If the the desired end-to-end delivery rate is 100% then then chosen routing method should be to only use links with PRR of 100%. This is an extreme form of reception-based blacklisting that we refer to as Perfect Link Routing (PLR).

6. Experimental Evaluation of Perfect Link Routing

In this section we use simulations to validate our conclusion that Perfect Link Routing (PLR) is more energy efficient than using the PRRxd metric for all three cases of no-ARQ, ARQ and NACK. Throughout the rest of this paper we follow Seada *et al.* [6, 7] by adopting the accurate PRR model developed by Zuniga and Krishnamachari [4]. Their model is based on the well-known log-normal shadowing model, given in equation (20) where d is the length of the

Variable	Value
η	4.15
σ	4.15
$PL(d_0)_{dB}$	55dB
P_t	0dBm
P_n	-115dBm
b	50 bytes
e_{tx}	49.5mW
e_{rx}	28.8mW

Table 3: Summary of the model variable values used in the simulations

link, d_0 is the reference distance, η is the path loss exponent and $\mathcal{N}(0, \sigma)$ is a zero-mean random Gaussian variable.

$$PL(d)_{dB} = PL(d_0)_{dB} + 10\eta \frac{d}{d_0} + \mathcal{N}(0, \sigma) \quad (20)$$

The path loss returned by the log-normal model is used to calculate the signal to noise ratio, $\gamma(d)$, which also relies on the output transmission power used by the transceiver, P_t $_{dB}$, and the noise floor, P_n $_{dB}$:

$$\gamma(d)_{dB} = P_t \text{ }_{dB} - PL(d)_{dB} - P_n \text{ }_{dB} \quad (21)$$

Combining these elements with transceiver specific constants and the number of bits per packets, b , the final model for the PRR of a link is:

$$PRR(d) = \left(1 - \frac{1}{2} \exp^{-\frac{\gamma(d)}{2} \frac{1}{0.64}} \right)^b \quad (22)$$

Throughout, we use the same values for the variables given by Zuniga and Krishnamachari and used by Seada *et al.* which are summarised in Table 3.

Additionally, we follow the general experimental methodology of Seada *et al.* with two minor changes. In their experiments, they simulated networks of static nodes in which (for each run) a randomly selected node transmits 100 packets to a randomly selected destination. The results presented are the average of 100 runs.

It is not clear from the original work whether the simulated network consisted of a chain of nodes or a 2-dimensional network. They define the network size in terms of the total number of nodes and the number of nodes per nominal transmission range. This definition suggests a chain network because knowing these two pieces of information is only enough to calculate the length of a chain and not a 2-dimensional area, unless some unstated assumptions are made. In our experiments we use a square network of length 250m, unless stated otherwise. We vary only the network density.

The network density is given in terms of the expected number of neighbours per node. We differ slightly from Seada *et al.* who defined the neighbourhood range of a node as being 40m which, for the model and variables being used, corresponds to roughly twice the size of the connected and transitional regions. The decision was justified on the basis that any nodes further than this distance from the transmitting node are extremely unlikely to have a valid link (greater than 1% PRR). However, the consequence is that the number of neighbours is unusually high, ranging from a “low” value of 25 neighbours per node up to 200 neighbours.

Since the results are fundamentally unaffected by the choice of neighbourhood range so long as the number of neighbours is adjusted, we take 10m as the neighbourhood range. This corresponds roughly with the outer edge of the connected region. We believe that this is more similar to the majority of studies which adopt the UDG model and therefore only consider the connected region. It allows us to give densities starting from just 2 neighbours per node. The community is used to thinking in terms of 6-10 neighbours per node as a connected network and our results more easily fit that mindset. ¹.

6.1. No-ARQ

In this section we consider the no-ARQ case where our conception of acknowledgements as a function of link quality is not relevant. In the original analysis by Seada *et al.* they concluded that for the no-ARQ case links with high PRR should be chosen to ensure high delivery rates and therefore higher efficiency. They simulated the no-ARQ case only to show that the delivery rate in such a scheme using original greedy or PRRxd is significantly lower than when using ARQ and therefore that the energy efficiency ought to be lower. Our simulation results reveal that the reality is not so clear cut.

Figure 3 shows the delivery rate for different network densities. As expected, the PRRxd method struggles to achieve high delivery rates because it selects nodes with less than perfect PRR and therefore packets occasionally fail. On the other hand, PLR achieves high delivery rates as soon as the density is high enough to avoid the problem of blacklisting the only available neighbours.

Our results show that using high-PRR links is not sufficient to achieve high delivery rates without retransmissions. Figure 4 shows the average PRR of the links chosen by the PRRxd method and the proportion of the links that have perfect PRR. The average PRR of the chosen links varies between 91.18% and 96.67% with between 16.37% and 41.23% of the links having perfect PRR. Despite this the delivery rate levels off at around 75%. It is therefore clear that using only high-quality links is not sufficient to achieve high delivery rates. Instead, it is necessary to use only perfect links.

¹See for example the famous work of Kleinrock and Silvester who give six as the “magic number” for connectivity [27] although this number has been revised up slightly to 8-10 in [10, 12]




Figure 3: With no retransmissions, end-to-end delivery can only be guaranteed if only perfect links are used. The PRRxd method cannot achieve much greater than a 75% delivery rate though this depends on the size of the network.




Figure 4: With the PRRxd method, the chosen links have very high PRR and a significant proportion of the chosen links are perfect ones with PRR of 100%. This explains why the network efficiency is high.




Figure 5: With no retransmissions the PRRxd method is more efficient than PLR despite having a lower delivery rate.




Figure 6: The path lengths without retransmissions are very short at low densities because packets are only successfully delivered if, by chance, some long but high quality links can be found in a chain. As the density increases paths become initially longer and then start to fall slowly with density.




Figure 7: As the network becomes larger PLR retains perfect delivery but the delivery rate for PRRxd falls as more hops are required and therefore there are more opportunities for packet delivery to fail.

On the other hand, Figure 5 shows that despite a lower delivery rate the PRRxd method is actually more energy efficient. The reason for this is that because PLR restricts itself to only perfect links its hop length is usually shorter and its path length is therefore significantly longer than the PRRxd method. Figure 6 shows that paths are on average 38.4% longer with PLR than PRRxd, once the lengths have settled. This means that even though PRRxd wastes a lot of energy sending packets along a path that fails before reaching the destination, this is outweighed by the extra hops needed by PLR.

These results would appear to contradict the claim that a high delivery rate is needed for high efficiency. However, further experimentation showed that the claim is, in fact, correct but the definition of “high“ is lower than one might initially think.

We kept the network density constant at 20 neighbours per node and varied the network length. Figure 7 shows the delivery rate for different network sizes. PLR retains 100% delivery despite the changing network size because the density is high enough to ensure that a valid path of perfect links exists between any source and destination. For the PRRxd method, however, the delivery rate falls off as the network becomes larger. This is because the paths becomes longer and there are more opportunities for packet delivery to fail.

The result of a falling delivery rate is that, eventually, PLR becomes more efficient than PRRxd. Figure 8 shows how the efficiency varies with increases network size. It is not at all surprising that as the network becomes bigger the efficiency falls since all paths are longer and it requires more energy to transmit the same sized packet from source to destination. The key result is that once




Figure 8: As the network becomes larger and the path lengths increase the efficiency falls since more energy is needed for end to end delivery. Crucially, though, once the delivery rate drops low enough PLR becomes more efficient.

the network size increases beyond some value (800m in our experiments) PLR becomes more efficient. This corresponds to the point at which the delivery rate for PRRxd falls below 35%. We can therefore confirm that a “high” delivery rate is needed for high efficiency but the delivery rate which is considered high may be as low as 35%.

It is important to note that the explanation for the cross-over of efficiency must lie in the delivery rate and not in the effects on path length. Figure 9 shows the impact network size has on path length. In all cases the path length will increase but crucially it rises faster with PLR than PRRxd. This means that at the point when PLR becomes more efficient its path length is approximately 43.3% longer than PRRxd whereas it was only 40.4% longer in the smaller network. Thus PLR becomes more efficient relative to PRRxd because the delivery rate falls under PRRxd and not because the path lengths become more similar.

6.2. ARQ

The results for the no-ARQ case show that PLR is necessary for high end-to-end delivery rates, especially with longer networks and that in longer networks PLR is also more efficient. In the ARQ case, however, the impact of acknowledgements becomes a factor as well as path length. When ARQ is used, as we have argued, not using a perfect link requires the use of acknowledgements which raise the energy cost of the link significantly. We would therefore expect PLR to be more efficient even with the relatively small network size of 250m.

Figure 10 shows the energy efficiency of the two methods when ARQ is used. At low densities the PRRxd method is more efficient. This is because




Figure 9: As the network size increases the paths under PLR become, relatively, even longer than PRRxd.




Figure 10: At low densities the PRRxd method has greater efficiency because PLR wastes energy on packets that are dropped further down the route. At higher densities, when PLR can always find a route, its energy efficiency becomes larger and stays higher.




Figure 11: With ARQ, the delivery rates reach 100% even for the PRRxd method since only voids in the network can prevent greedy forwarding eventually working.

the blacklisting imposed by PLR often prevents it from finding a route to the destination. When that happens, any energy spent on the packet is all wasted when the packet is dropped. After a density of about ten neighbours, however, the efficiency of PLR becomes higher than PRRxd and becomes increasingly higher. With ten neighbours per node, PLR is 8.86% more efficient than PRRxd but this rises to 32.77% with 30 neighbours. This confirms that PLR is more efficient for most densities with ARQ, even with a smaller network size of 250m.

As mentioned, PLR has reduced delivery rates at low densities. Figure 11 shows that the use of ARQ raises the delivery rate of the PRRxd method to 100% even at low densities whereas PLR requires a moderately dense network before delivery is guaranteed.

The blacklisting of links not only reduces the delivery rate at low densities but also increases the path length when delivery is successful. Most of the longer links which sit inside the transitional region have less than perfect PRR even if their PRR is still very high. PLR blacklists these links and thereby restricts itself to usually using shorter links resulting in longer paths. PRRxd, on the other hand, can take advantage of high quality (but not perfect) links when their length is high as well and therefore has shorter paths.

Figure 12 shows the path lengths of the two methods. As with the no-ARQ case at low densities both show short path lengths because nodes are usually distant from each other. The path length rises sharply before starting a slow decline as density increases. Once the decline starts the difference in path lengths between PLR and PRRxd remains approximately constant with PLR having an average of 38.68% more hops than PRRxd.




Figure 12: PLR blacklists all links with less than perfect PRR which includes most of the longer links in the transitional region. PRRxd can still use some of these and therefore has much shorter paths.

6.3. NACK

When no acknowledgements are used (the No-ARQ case) the PRRxd method can be more efficient than PLR even though it has a lower delivery rate. As the network size increases, however, the path lengths increase and the delivery rate using PRRxd falls until PLR becomes more energy efficient. In contrast, when acknowledgements are always used (the ARQ case) except for perfect links, PLR is more efficient even with small networks because it completely avoids the added cost of these acknowledgements.

The NACK case, in terms of delivery rate, lies between the no-ARQ and ARQ cases. On the one hand it uses retransmissions to raise the effective PRR of links and therefore can achieve higher delivery rates than without acknowledgements. On the other hand, it cannot achieve perfect delivery rates like ARQ because the NACK packets can fail leading to false positives where the sender believes the data packet to have been received successfully when it was not.

In terms of efficiency the NACK case is also somewhere between the previous two cases. Unlike the no-ARQ case, it does make use of retransmissions which are avoided on perfect links which would tend to favour PLR. However, since NACK packets are only needed after failures and PRRxd favours high quality packets, acknowledgements are far less frequent even on non-perfect links than in the ARQ case. This makes it less important to utilise only perfect links.

Figure 13 shows the delivery rate for the NACK case. Using the PRRxd method results in high delivery rates at low densities but it never reaches 100%, levelling off at about 94.42%. Since PLR only uses perfect links there can be no false positives and so the delivery rate does reach 100% but only slowly because




Figure 13: Because NACKs may fail and create a false positive delivery is not guaranteed with PRRxd. Using only perfect links avoids this problem.

of the blacklisting.

As found in the no-ARQ case, the PRRxd method favours high-quality links and so in the NACK case the number of acknowledgement packets needed is significantly smaller than in the ARQ case. Therefore, the extra cost of these acknowledgements appears far less often. In fact, as Figure 14 shows, the PRRxd scheme is more efficient in the NACK case than PLR. The difference becomes slightly smaller as the density increases but settles down with the PRRxd method being about 11.5% more efficient than PLR.

The PLR method still retains its cost of additional path length, as illustrated in Figure 15. Once the density settles the path length remains around 38.96% longer under PLR than PRRxd.

As with the no-ARQ case, although we have found that PRRxd is more efficient than PLR this is only for relatively small networks. As the network size increases, the opportunities for false positives increases and the delivery rate under PRRxd falls. At some point PLR becomes more efficient because it retains its perfect end-to-end delivery rate.

Figure 16 shows the fall in delivery rate as the network size increases. Since PLR avoids false positives by only using perfect links it retains 100% delivery rate at all network sizes. PRRxd, however, has decreasing delivery rates because larger networks lead to longer paths and therefore more opportunities for false positives.

Figure 17 confirms that the drop in delivery rate results in a drop in relative efficiency between PRRxd and PLR. For networks greater than about 1,700m PLR becomes more efficient, reaching up to 26.15% more efficient at 4,000m. The crossover point corresponds to a delivery rate of 70% for PRRxd. Compared




Figure 14: Since the NACK scheme requires far fewer acknowledgements than the ARQ scheme the PRRxd method is more efficient.




Figure 15: The PLR scheme results in increased path length once the density is high enough to allow reasonable delivery rates.




Figure 16: As the network becomes larger and average path lengths increase there is more opportunity for false positives using the NACK scheme leading to a falling delivery rate for PRRxd.




Figure 17: Because the delivery rate falls with PRRxd as the network becomes larger, eventually the amount of wasted energy from failed packets makes it less efficient than PLR.

to the no-ARQ case, PLR becomes more efficient only with larger networks but at almost double the delivery rate.

The results in this section have shown that PLR is more energy efficient in all cases under certain conditions. In the no-ARQ case PLR is more efficient when the network is large enough that the delivery rate for PRRxd becomes very low (less than 35% in our experiments). In the ARQ case PLR is more energy efficient even in small networks so long as the density is high enough to allow for perfect delivery rates. In the NACK case PLR is again more efficient only with larger networks, when the impact of false positives causes the delivery rate to drop for the PRRxd method (less than 70% in our experiments).

7. Modifying PLR

In the previous section we showed that in all cases PLR is more efficient for large, dense networks because it guarantees end-to-end delivery without the need for acknowledgements and retransmissions. For low-density networks, however, there is no such guarantee and in fact the delivery rate is worse than with PRRxd because PLR blacklists, and therefore ignores, numerous usable links. This problem has made blacklisting schemes unpopular. In this section we modify PLR to become a hybrid scheme and also consider an alternative hybrid. Both methods prioritise perfect links but can use other links if necessary.

The modified version of PLR, which we refer to as Hybrid-PLR in this section to distinguish it from the original version, first tries to find a perfect link and, if this fails, switches to using PRRxd to find the next hop. This is an attempt to achieve the best of both worlds; to utilise PLR at moderate to high densities and PRRxd at lower densities. The second scheme that we investigate is Best Reception (BR) which was already considered by Seada *et al.* This method always uses the link with the highest PRR using distance only as a tie-breaker. Although not strictly a hybrid scheme it can be thought of as using PLR when it can and using BR otherwise. Seada *et al.* concluded that, although BR is not optimal, it is nevertheless a good choice for systems with ARQ.

In this section we consider both hybrid schemes (along with the original PLR and PRRxd) for the ARQ scheme. In the No-ARQ and NACK cases we have already shown that only PLR can guarantee end-to-end delivery because it avoids failed transmissions and false positives. Therefore, it is only in the case of ARQ at low densities where the problems of low delivery rate are apparent for PLR.

Figure 18 shows the delivery rate for all four methods. Hybrid-PLR, as one might expect, shows almost identical performance as the PRRxd scheme since, if it cannot make progress using PLR, it becomes PRRxd and therefore is very unlikely to be unable to find a route when PRRxd can. Note that there is no guarantee that Hybrid-PLR will always find a route when PRRxd has one available because in theory a link chosen by PLR might lead to a void whereas one chosen by PRRxd might avoid it. However, there is no method for predicting this and the opposite is also possible where PRRxd leads to a void and PLR does not.




Figure 18: The Hybrid-PLR scheme has almost identical delivery rates as PRRxd while Best Reception shows a small improvement over PLR.

BR shows a much smaller improvement over PLR. This is because it only considers link reception and (almost) never takes into account the link progress which leads it to select usually shorter links. The result is that the path length under BR (see below) is similar to that of PLR. This means that it has more chance to run into a void than Hybrid-PLR.

The results for efficiency, in Figure 19, tell an interesting story. At high densities BR and Hybrid-PLR are indistinguishable from PLR and all three are significantly more efficient than PRRxd. At lower densities, however, the Hybrid-PLR scheme actually outperforms both PLR and PRRxd. Because it can take advantage of PRRxd when PLR fails, it has a higher delivery rate than PLR and therefore reduces the number of times energy is wasted by sending a packet part of the way before it is dropped. On the other hand, by using more perfect links than PRRxd it avoids the need for as many acknowledgements and is more efficient than that method alone.

Best Reception only has one of these improvements in that it has a better delivery rate than PLR and therefore has a higher efficiency than PLR when delivery rates are low. However, when perfect links are not available, and acknowledgements are going to be used regardless, BR does not have the flexibility to select longer hops to reduce overall energy usage and therefore is not as efficient as PRRxd. Only when the delivery rates reach 100% and there are always perfect links to use does the efficiency of BR match that of Hybrid-PLR and beat PRRxd.

The cost of increased efficiency and delivery rate is increased path length. Figure 20 shows that at high density, when both hybrid schemes are essentially the same as PLR, the path length is indistinguishable from that of PLR. At lower




Figure 19: Both hybrid schemes are more efficient than PLR at low densities and than PRRxd at high densities with Hybrid-PLR even managing to be more efficient than PRRxd at low densities.




Figure 20: The two hybrid schemes have virtually the same path lengths as PLR at high densities because they are essentially identical to it in that range. At lower densities their lengths are a reflection of their improved delivery rates.

densities the path length is affected by the prevalence of voids. PLR and BR which always prefer very short hops run into voids frequently and therefore have a low delivery rate. They can only successfully route a packet if there happens to be a chain of nodes that do not hit voids. This is most likely to be the case when only a few hops are used because then there are fewer opportunities to find a void.

Hybrid-PLR, however, is able to recover from voids that PLR falls into and therefore can have successful packets even when using a larger number of short hops. That Hybrid-PLR uses many short hops is evident in its path length which is significantly longer than PRRxd despite them both having the same delivery rate.

The results in this section show that Hybrid-PLR, which uses PLR whenever possible and PRRxd at other times, has the delivery rate of PRRxd and efficiency that is at least as high as either of the base methods and is higher than PRRxd at low densities.

8. When Acknowledgements Become Less Significant

We have shown that PLR is more efficient than PRRxd for all cases and for most network sizes and densities. In the cases of ARQ and NACK this is, in part at least, because of the cost of the acknowledgements that are avoided in PLR. Up to now we have followed the standard assumption of considering acknowledgements to be approximately the same size as the data packets in low-power networks. In this section we relax this assumption and investigate the impact that a less significant acknowledgement packet has on the relative performance of PRRxd, PLR and Hybrid-PLR for both the ARQ and NACK cases.

In the experiments in this section we introduce a new variable, α , which is the relative size of data packets to acknowledgements. That is, if $\alpha = 2$ then data packets are twice the size of acknowledgements. For completeness' sake we consider also one case when acknowledgements are actually larger than the data packets and therefore vary α between 0.5 and 5.

8.1. ARQ

For the case of ARQ we consider the impact of changing α for three different densities of 6, 12 and 24 neighbours per node which are roughly low, medium and high densities. Figure 21 shows the effects of α at low density. Because it is low density, PLR is less efficient than PRRxd but, as seen earlier, Hybrid-PLR is more efficient. In all cases as α increases the efficiency falls since data packets require more energy but the efficiency of PLR and Hybrid-PLR fall faster than for PRRxd. There are two ways to think of this. One approach is to consider that since PLR and Hybrid-PLR avoid acknowledgements where possible then the data packets make up a larger proportion of their energy consumption and therefore when the data packets become larger their energy usage increases by a larger proportion than PRRxd. A second viewpoint is




Figure 21: At low densities, as seen earlier, PLR always has lower efficiency than PRRxd. Hybrid-PLR is more efficient than PRRxd when acknowledgements are expensive (alpha is small) but as the acknowledgements become less significant then PRRxd becomes slightly more efficient.

that as data packets increase relative to acknowledgements then the relative “penalty” of using acknowledgements falls and the costs of shorter hops begins to outweigh the benefits of avoiding acknowledgements.

Figures 22 and 23 show the results for the medium and high density networks respectively. A similar result to the low density network is seen but in these cases PLR has almost identical efficiency to Hybrid-PLR.

There are two important factors that change as the density increases. Firstly, the value of α at which PRRxd becomes more efficient is higher at higher densities. The values of α at the crossover point are approximately 1.75, 1.87 and 2.0 for the low, medium and high densities respectively.

The second important factor is the difference in efficiency. When the density is low then when $\alpha = 5$ PRRxd is 26.32% more efficient than Hybrid-PLR. This value falls to 25.51% and 21.92% at the medium and high densities.

8.2. NACK

For the NACK case we have seen that PLR becomes more efficient than PRRxd when the network is relatively large and the delivery rate for PRRxd falls. Since acknowledgements are used less often in the NACK case than the ARQ case the advantage that PLR has over PRRxd from the avoidance of all acknowledgements is far less pronounced. This means that not only is PRRxd the more efficient for longer than in the ARQ case but also that changes in the relative size of acknowledgements are likely to have less of an impact. Nevertheless, we might reasonably expect that if acknowledgements are smaller than




Figure 22: At medium densities Hybrid-PLR and PLR perform very similarly and the relative size of data packets to ACKs must be higher before PRRxd becomes more efficient.




Figure 23: At high densities the results are similar to medium densities but again the value of alpha needed before PRRxd becomes more efficient is higher.




Figure 24: As acknowledgements become less significant the point at which PLR (and Hybrid-PLR) become more efficient than PRRxd increases quickly before settling.

data packets then the network size required for PLR to become more efficient than PRRxd would increase.

To test this we kept the network density constant and high and varied the value of α . At each value we found the approximate network size at which PLR becomes more efficient than PRRxd. Figure 24 shows the results. Initially the crossover point increases rapidly as the advantage of avoiding acknowledgements quickly shrinks. At some point, however, the efficiency gain of avoiding acknowledgements becomes very small and PLR outperforms PRRxd mostly because of its better delivery rate. When this happens the increase in the crossover point slows dramatically.

These results show that even when acknowledgements are of little consequence, PLR is still more efficient than PRRxd for large networks. This is what we observed with the no-ARQ case where acknowledgements play no part. Avoiding false positives is the main advantage for PLR in the NACK case as it leads to a high delivery rate which causes high efficiency, just as in the no-ARQ case.

The results in this section show that modifying the original version of PLR to utilise PRRxd when no perfect links are available results in the best performance. Therefore, for the remainder of this paper when we refer to PLR we intend the modified version that we called Hybrid-PLR in this section.

9. Additional Factors Favouring Perfect Link Routing

In this paper we have shown that Perfect Link Routing (PLR) is optimal in terms of energy efficiency for networks using ARQ with at least moderate

node density. For networks using NACK or no acknowledgements at all then only PLR is able to guarantee end-to-end delivery of any given packet. Our arguments for the optimality of PLR in the ARQ case have been based on the additional energy cost of needing to send acknowledgements which is avoided when only perfect links are used. In this section we consider additional factors that favour the use of perfect (or at least high quality) links. These factors are based on the spatial and temporal characteristics of low-power links.

A number of studies have been conducted into the nature of low-power links including [3, 28, 29] and the results have been recently summarised by Baccour *et al.* [30]. Based on the nature of wireless links we argue that, aside from the savings in terms of acknowledgements, PLR also reduces energy cost for other reasons and will likely result in reduced interference and latency.

9.1. Temporal Characteristics

The quality of low power links can vary significantly over relatively short amounts of time in hard to predict ways. This temporal variation makes it difficult to calculate the packet reception rate (PRR) of a given link at a given moment. However, it has been found that links with very high or very low PRR are much more stable than other links. This means that if a link is found to have a very high PRR on numerous occasions then it is very likely that it will always have a high PRR. Intermediate quality links, by contrast, can sometimes have very high PRR and sometimes very low.

The fact that high quality links are more stable has an important implication for routing protocols that need to know link quality. The quality of links with high PRR can be established with less frequent beaconing because changes are known to occur relatively infrequently. By contrast, intermediate quality links require more frequent beaconing to ensure that estimates are not out of date. It may even be necessary to establish the PRR of links immediately before every transmission. This favours PLR since it only considers high quality links for routing and can therefore have less frequent beaconing than PRRxd.

It has also been shown that in low-power links there is a strong correlation between the link quality over very short time periods. That is, if two packets are sent with little or no break then the success or failure of the packets are highly correlated. In practical terms this means that if a retransmission is needed and it is sent immediately there is a high likelihood that the second transmission will fail just as the first did. There are two possibilities, therefore. Either protocols retransmit immediately following a failure in which case the retransmissions are likely to fail as well which wastes energy; or protocols introduce a delay between a failed transmission and a retransmission which introduces latency. PLR does not suffer from this problem since the probability of the initial packet failing is insignificant. If extra latency is used to maintain efficiency then it may well be the case that while PLR has a longer path to the destination it nevertheless has a smaller latency.

9.2. *Link Asymmetry*

In this paper, and in the literature, link symmetry is assumed. That is, the PRR of a link between two nodes is assumed to be the same in both directions. In reality however, links can exhibit asymmetry in which a link has a high quality in one direction and a low quality in the reverse direction. If an asymmetric link is used for routing then this hampers the usefulness of acknowledgements. It may be that a data packet is received without trouble but the acknowledgement informing the transmitter of this fact fails frequently thus wasting energy on unnecessary retransmissions. In the NACK case the delivery rate will suffer as false positives increase when necessary NACK packets do not arrive informing the transmitter that the data packet failed to arrive safely.

Although it is not well established whether link asymmetry is transient or not, it is established that intermediate links are usually asymmetric whilst links with either very high or very low PRR or usually symmetric. Again this observation favours PLR since it only utilises high quality links and therefore the symmetric assumption holds. Routing protocols such as PRRxd may find their performance degrade in real networks because of link asymmetry.

9.3. *Link Quality Estimation*

Both PLR and PRRxd rely on accurate estimates of the link quality. Much research has been undertaken in this area and a number of estimation techniques have been proposed. The cheapest methods are hardware based such as RSSI or LQI. These methods come at no extra cost and are quick to compute. However, it is known that these methods are not accurate at estimating the value of the quality and are only usable for establishing whether or not a link is high quality. That is, RSSI can be used to say that a link has a very high PRR or not but if it the link is not high quality then RSSI cannot be relied upon to estimate its actual PRR.

The result is that the energy cost of finding high quality links is very small since only one RSSI sample is needed for this. The energy cost of finding the PRR of intermediate quality links, however, is much higher. Results have shown that it may require anywhere between 40 and 120 packet samples to accurately estimate the quality of such links. Therefore PLR is not only cheaper when high quality links are used it is also significantly cheaper when finding possible links to use.

10. **Conclusion and Future Work**

Geographic routing is widely considered a good choice for ad hoc networks because of its scalability and low resource requirements. In the vast majority of research, the Unit Disk Graph model is adopted which makes unrealistic assumptions about the quality of wireless links. In reality, a large proportion of wireless links have link qualities that can vary widely. For geographic routing protocols to be effective they must consider the Packet Reception Rate (PRR) of the links as well as the progress they make towards the destination.

In an early but often overlooked work, a metric was proposed that combined PRR and distance into a single metric of PRRxd. This was argued to be optimal in both no-ARQ and ARQ cases. In this paper we extended the original analysis to consider the possibility of control packet failures and considered also the case of negative acknowledgements. Our analysis and results show that, for energy efficiency, geographic routing protocols should prioritise the use of perfect links which leads us to advocate the use of an extreme form of reception-based blacklisting which we call Perfect Link Routing (PLR). We found that, for large, dense networks PLR was optimal regardless which retransmission method is used (or if none is used).

At low densities, however, PLR can result in a loss of delivery rate. We solve this problem by modifying PLR slightly into a hybrid protocol which uses PLR wherever possible and, if PLR does not find a next hop, uses PRRxd to find a forwarding node. Hybrid-PLR performs at least as well as the two underlying protocols and, in some cases, can outperform both.

We have strengthened the case of PLR with observations of the temporal character of wireless links, the problem of link asymmetry and the energy cost of establishing link quality. We believe that further experimentation and discussion is warranted on these factors to establish for certain the nature of low power links.

Importantly, we note that the vast majority of perfect links are inside the connected region and few, if any, links inside the transitional region are perfect. This leads to the tantalising prospect that PLR may be, in effect, very similar or identical to geographic routing using the UDG model. If this equivalence holds then the past decade or more of research that relied on the UDG model would be validated for real applications using PLR.

References

- [1] Dazhi Chen and Pramod K Varshney. A survey of void handling techniques for geographic routing in wireless networks. *Communications Surveys & Tutorials, IEEE*, 9(1):50–67, 2007.
- [2] 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*, pages 14–27. ACM, 2003.
- [3] Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems*, pages 1–13. ACM, 2003.
- [4] M. Zuniga and B. Krishnamachari. Analyzing the transitional region in low power wireless links. In *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*, pages 517–526. IEEE, 2004.

- [5] Thomas Watteyne, Antonella Molinaro, Maria Grazia Richichi, and Mischa Dohler. From manet to ietf roll standardization: A paradigm shift in wsn routing protocols. *Communications Surveys & Tutorials, IEEE*, 13(4):688–707, 2011.
- [6] K. Seada, M. Zuniga, A. Helmy, and B. Krishnamachari. Energy-efficient forwarding strategies for geographic routing in lossy wireless sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 108–121. ACM, 2004.
- [7] Marco Zúñiga Zamalloa, Karim Seada, Bhaskar Krishnamachari, and Ahmed Helmy. Efficient geographic routing over lossy links in wireless sensor networks. *ACM Trans. Sen. Netw.*, 4(3):12:1–12:33, June 2008.
- [8] A. Kleerekoper and N. Filer. Revisiting blacklisting and justifying the unit disk graph model for energy-efficient position-based routing in wireless sensor networks. *Wireless Days*, 2012.
- [9] I. Stojmenovic. Position-based routing in ad hoc networks. *Communications Magazine, IEEE*, 40(7):128–134, 2002.
- [10] Hideaki Takagi and Leonard Kleinrock. Optimal transmission ranges for randomly distributed packet radio terminals. *Communications, IEEE Transactions on*, 32(3):246–257, 1984.
- [11] Ting-Chao Hou and Victor Li. Transmission range control in multihop packet radio networks. *Communications, IEEE Transactions on*, 34(1):38–44, 1986.
- [12] Ivan Stojmenovic and Xu Lin. Power-aware localized routing in wireless networks. *Parallel and Distributed Systems, IEEE Transactions on*, 12(11):1122–1133, 2001.
- [13] Gregory G Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical report, DTIC Document, 1987.
- [14] Evangelos Kranakis, Harvinder Singh, and Jorge Urrutia. Compass routing on geometric networks. In *in Proc. 11 th Canadian Conference on Computational Geometry*. Citeseer, 1999.
- [15] Gang Zhou, Tian He, Sudha Krishnamurthy, and John A Stankovic. Impact of radio irregularity on wireless sensor networks. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 125–138. ACM, 2004.
- [16] D.S.J.D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. *Wireless Networks*, 11(4):419–434, 2005.

- [17] Gang Zhou, Tian He, Sudha Krishnamurthy, and John A. Stankovic. Models and solutions for radio irregularity in wireless sensor networks. *ACM Trans. Sen. Netw.*, 2(2):221–262, May 2006.
- [18] J.A. Sanchez and P.M. Ruiz. Improving delivery ratio and power efficiency in unicast geographic routing with a realistic physical layer for wireless sensor networks. In *Digital System Design: Architectures, Methods and Tools, 2006. DSD 2006. 9th EUROMICRO Conference on*, pages 591–597. IEEE, 2006.
- [19] I. Stojmenovic, A. Nayak, and J. Kuruvila. Design guidelines for routing protocols in ad hoc and sensor networks with a realistic physical layer. *Communications Magazine, IEEE*, 43(3):101–106, 2005.
- [20] J. Kuruvila, A. Nayak, and I. Stojmenovic. Hop count optimal position-based packet routing algorithms for ad hoc wireless networks with a realistic physical layer. *Selected Areas in Communications, IEEE Journal on*, 23(6):1267–1275, 2005.
- [21] I. Stojmenovic, A. Nayak, J. Kuruvila, F. Ovalle-Martinez, and E. Villanueva-Pena. Physical layer impact on the design and performance of routing and broadcasting protocols in ad hoc and sensor networks. *Computer Communications*, 28(10):1138–1151, 2005.
- [22] S. Lee, B. Bhattacharjee, and S. Banerjee. Efficient geographic routing in multihop wireless networks. In *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, pages 230–241. ACM, 2005.
- [23] S. Lee, B. Bhattacharjee, S. Banerjee, and B. Han. A general framework for efficient geographic routing in wireless networks. *Computer Networks*, 54(5):844–861, 2010.
- [24] Flavio Fabbri, Marco Zuniga, Daniele Puccinelli, and Pedro Marrón. On the optimal blacklisting threshold for link selection in wireless sensor networks. In *Wireless Sensor Networks*, pages 147–162. Springer, 2012.
- [25] Fraser Cadger, Kevin Curran, Jose Santos, and Sandra Moffett. A survey of geographical routing in wireless ad-hoc networks. 2012.
- [26] Zdenek Becvar and Robert Bestak. Overhead of arq mechanism in ieee 802.16 networks. *Telecommunication Systems*, 46(4):353–367, 2011.
- [27] Leonard Kleinrock and John Silvester. Optimum transmission radii for packet radio networks or why six is a magic number. In *Proceedings of the IEEE National Telecommunications Conference*, volume 4, pages 1–4. Birmingham, Alabama, 1978.

- [28] Alberto Cerpa, Jennifer L Wong, Louane Kuang, Miodrag Potkonjak, and Deborah Estrin. Statistical model of lossy links in wireless sensor networks. In *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*, pages 81–88. IEEE, 2005.
- [29] Kannan Srinivasan, Prabal Dutta, Arsalan Tavakoli, and Philip Levis. An empirical study of low-power wireless. *ACM Transactions on Sensor Networks (TOSN)*, 6(2):16, 2010.
- [30] Nouha Baccour, Anis Koubâa, Luca Mottola, Marco Antonio Zúñiga, Habib Youssef, Carlo Alberto Boano, and Mário Alves. Radio link quality estimation in wireless sensor networks: A survey. *ACM Trans. Sen. Netw.*, 8(4):34:1–34:33, September 2012.