Ad Hoc Networks 30 (2015) 46–62

Contents lists available at ScienceDirect

Ad Hoc Networks

journal homepage: www.elsevier.com/locate/adhoc

Perfect link routing for energy efficient forwarding in geographic routing

A. Kleerekoper *, N.P. Filer

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

ARTICLE INFO

Article history:
Received 6 October 2014
Received in revised form 8 February 2015
Accepted 3 March 2015
Available online 19 March 2015

Keywords:
Ad hoc networks
Energy efficiency
Geographic routing

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.

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 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

* Corresponding author.
E-mail addresses: kleereka@cs.man.ac.uk (A. Kleerekoper), nick@cs.man.ac.uk (N.P. Filer).

http://dx.doi.org/10.1016/j.adhoc.2015.03.001
1570-8705/

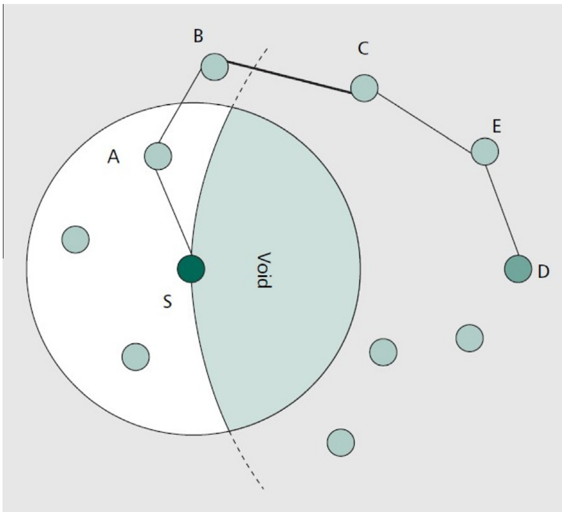

Fig. 1. Greedy forwarding can fail when a void is encountered and techniques are needed to route around such "holes". Figure from [1].

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 [2–4]. 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 PRRxd.

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 $PRR^2 d/(1 + PRR)$ 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^2xd 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, $PRR^2 d/(1 + PRR)$, 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 5)
- 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 5)
- We examine how the efficiency of PLR varies when data packets become larger relative to control packets (Section 5)
- 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