Ad Hoc Networks 36 (2016) 35–48

Contents lists available at ScienceDirect

Ad Hoc Networks

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

On-demand beaconing: Periodic and adaptive policies for effective routing in diverse mobile topologies

Ioannis Manolopoulos[a,b], Kimon Kontovasilis[a,*], Ioannis Stavrakakis[b], Stelios C.A. Thomopoulos[a]

[a] Institute of Informatics & Telecommunications, National Center for Scientific Research "Demokritos", GR-15310 Agia Paraskevi, Greece
[b] Department of Informatics & Telecommunications, National & Kapodistrian University of Athens, GR-15784 Ilissia, Greece

ARTICLE INFO

Article history:
Received 7 November 2014
Accepted 27 May 2015
Available online 19 June 2015

Keywords:
On-demand beaconing
Periodic/adaptive beaconing
Diverse network conditions
Routing protocols
Delay tolerant networks

ABSTRACT

Locally aware routing protocols base their next-hop selection on information about their immediate neighborhood, gathered by means of a beaconing mechanism. In general, beacons may be proactively broadcasted from nodes to their neighbors ('receiver-initiated' beaconing) or may be solicited by the node carrying the routed message ('on-demand' beaconing). On-demand beaconing is of growing importance, mainly in more dynamic and sparse environments (e.g., delay tolerant networks), and is addressed in this paper.

A generic analysis is provided for the case of periodically issued beacons, linking the beacon period to the trade-off between the quality of neighborhood perception (determining the routing effectiveness) and the required amount of signaling (related to energy expenditure at the nodes). The analysis leads to upper and lower bounds for the length of the beacon period, expressed in terms of mobility characteristics.

The paper also investigates policies where the inter-beacon intervals vary adapting to the environment, an approach most beneficial when routing is based on metrics bearing some relevance to time. This is the case with the MAD routing protocol, which incorporates the notion of 'retaining time', an estimate of the time that the carrying node will retain the message. It is shown that linking the beacon intervals to the each time applicable retaining time leads to an effective and efficient beacon policy.

The paper provides simulation-based evaluation results, validating the beacon period bounds and demonstrating that, for the case of MAD, adaptive beaconing is capable of providing even better performance.

1. Introduction

In recent years there has been a growing interest in locally aware routing protocols, for use in mobile ad hoc networks and networks with even sparser topologies, e.g., delay tolerant networks (DTNs). Such protocols can cope with the variability encountered in mobile environments by focusing only on information about the local area around the routed message, rather than trying to find or maintain an end-to-end path. Specifically, the node that carries a message collects relevant status from its neighbor nodes and then selects the most suitable neighbor for forwarding/routing the message. Typical suitability criteria include, among others, distance from the destination (to select the closest neighbor, which is suitable for routing in relatively dense topologies),

* Corresponding author. Tel.: +30 2106503130.
E-mail addresses: imanolo@iit.demokritos.gr (I. Manolopoulos), kkont@iit.demokritos.gr (K. Kontovasilis), ioannis@di.uoa.gr (I. Stavrakakis), scat@iit.demokritos.gr (S.C.A. Thomopoulos).

http://dx.doi.org/10.1016/j.adhoc.2015.05.015
1570-8705/

or velocity (to select the best directed and/or fastest-moving node, which is suitable for routing in sparser topologies), or appropriate mixtures of both.

Regardless of the details associated with the particular suitability metric employed, there are two mechanisms that underlie the routing operation. The first is a mechanism to discover the position of the message destination, typically implemented through querying a location service [1]. The second is a signaling mechanism, referred to as beaconing, that enables the collection of neighbors-related status (position and/or velocity) by the node carrying a message.

There are two main approaches to beaconing. The first is 'receiver-initiated': the node possessing the message (also referred to as the "current node" in the following) waits for beacons issued proactively from its neighbors [2,3]. The second approach is 'on-demand': the current node triggers itself the neighbors to send information, when it sees fit [4,5]. Finally, there exist some schemes that attempt to eliminate beacons altogether, aiming at a reduced signaling overhead [6–8]. However, beacon-less schemes are associated with significantly higher end-to-end delays [9] and need to resort to local packet broadcasting, rather than next hop forwarding, something that leads to a higher collision probability and to a lower packet delivery ratio [10]. For a comprehensive taxonomy of the scheme categories just outlined see [11,12].

In this work we focus on the 'on-demand' approach because it is better for routing in highly diverse topologies, which is our main target. Indeed, in dynamic and probably sparse topologies the routing information becomes frequently outdated and the current node deciding about the next hop should make sure that it has up-to-date information, in order to decide optimally and avoid forwarding failures. The 'on-demand' approach guarantees that an accurate depiction of the local topology is available, contrary to the 'receiver-initiated' approach, where it is necessary to incorporate additional mechanisms for validating the cached topology information [13,14]. Another important characteristic of the 'on-demand' approach is the potential for energy savings at the nodes, something that contributes to a prolonged network lifetime. This is because in 'on-demand' schemes only nodes in the local area around the message need to be awake when triggered by a beacon from the current node. However, exploitation of this potential typically requires the usage of an additional interface, through which nodes are triggered to wake up.

An important aspect of 'on-demand' beaconing is the pattern according to which the current node issues beacons to trigger its neighbors for information. An obvious possibility is to use a scheme of regularly issued beacons (referred to as 'periodic on-demand beaconing' in the sequel). Additionally, in many cases it is also possible to vary the inter-beacon intervals, in response to changes in the values of appropriate status parameters. This is called 'adaptive on-demand beaconing' in the following.

Periodic 'on-demand' beaconing has the merit of being generic, thus suitable for use with any existing routing protocol. Given the periodic nature of the beacons, the issue is to select an appropriate beacon period, as the value of this parameter has clear implications on the trade-off between performance and signaling overhead (and associated energy expenditure). Indeed, if the beacon period is very short, the local area around the routed message is burdened with unnecessarily heavy signaling (and the associated energy depletion and bandwidth consumption side-effects), without significant gains in status updates, as it is likely that almost nothing will have changed from the previous check of the neighborhood. On the other hand, if the beacon period is very long the signaling becomes negligible, but the current node now has only a poor perception of its neighborhood and may miss forwarding opportunities, so the routing becomes less effective.

Despite these important implications, to the best of the authors' knowledge the issue of determining appropriate values for the beacon period has not been studied yet. In an attempt to fill this gap, the paper contributes a generic analysis linking the beacon period to the quality of neighborhood perception. The analysis leads to upper and lower bounds for the length of the beacon period, expressed in terms of mobility characteristics. By employing a beacon period between the two boundary values, the current node can successfully balance the trade-off between routing effectiveness and signaling overhead. In accordance with the generic overall character of the periodic 'on-demand' beaconing, the analysis is applicable to any routing protocol that may make use of the beaconing scheme.

No matter how attractively simple and ubiquitously applicable the periodic 'on-demand' beaconing may be, there are cases where adaptive 'on-demand' beaconing schemes offer a greater potential, due to their intelligent adaptation of beacon intervals to status changes. Ideally, an adaptive 'on-demand' scheme should issue a beacon to trigger a neighborhood exploration only when the current node "senses" that some neighbor is likely to be a more suitable carrier than itself. Obviously, such functionality can be accomplished only if the routing protocol provides appropriate support, through relevant parameters and/or metrics.

The recent routing protocol MAD, developed in [15] and refined in [16], is particularly suitable for use with such adaptive 'on-demand' beaconing, as it incorporates a notion of time that can be naturally exploited for adaptivity. Specifically, MAD employs the, so called, 'retaining time', an estimate of the time that a node will keep the message once selected for carrying it. The retaining time encapsulates information relevant to the local environment around the routed message (nodal density and mobility) and to the current node (location and motion attributes).

In view of these remarks, the second contribution of the paper is an adaptive 'on-demand' beaconing scheme based on the current node's retaining time. As with periodic beaconing, to the best of the authors' knowledge this is the first attempt on beacon adaptation for the domain of 'on-demand' schemes. However, for completeness we now compare our approach with the recent works [14,17], addressing the 'receiver-initiated' beaconing context.

[17] proposes two different adaptive schemes. In the first, called 'distance-based' beaconing, a node broadcasts a beacon whenever it has moved a given distance *d*. Also, a node removes from its neighborhood list a neighbor after moving for more than *k* times the distance *d* without hearing a beacon from this neighbor, or after a maximum time-out elapses. The main drawbacks of this scheme are that a slow node has