Computer Communications 38 (2014) 36–49

Contents lists available at ScienceDirect

Computer Communications

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

Optimal multipath forwarding in planned Wireless Mesh Networks

Germán Capdehourat *, Federico Larroca, Pablo Belzarena

Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Julio Herrera y Reissig 565, ZC 11300 Montevideo, Uruguay

ARTICLE INFO

Article history:
Received 12 November 2012
Received in revised form 19 September 2013
Accepted 5 October 2013
Available online 29 October 2013

Keywords:
Wireless Mesh Networks
Traffic engineering
Load-balancing

ABSTRACT

Wireless Mesh Networks (WMNs) have emerged in the last years as a cost-efficient alternative to traditional wired access networks. In the context of WMNs resources are intrinsically scarce, which has led to the proposal of dynamic routing in order to fully exploit the network capacity. We argue instead in favour of separating routing from forwarding (i.e. á *la* MPLS). Our proposal is a dynamic load-balancing scheme that forwards incoming packets along several pre-established paths in order to minimize a certain congestion function. We consider a particular but very typical scenario: a planned WMN where all links do not interfere with each other. We use a simple and versatile congestion function: the sum of the average queue length over all network nodes interfaces. We present a method to learn this function from measurements and several simulations to illustrate the framework, comparing our proposal with the IEEE 802.11s standard.

1. Introduction

Wireless Mesh Networks (WMNs) [1] are no longer just a promise for the future but a reality today, thanks mainly to the advantage offered in terms of cost compared to traditional wired access networks. In particular, outdoor community mesh networks [2] and rural deployments [3,4] based on IEEE 802.11 have seen tremendous growth in the recent past. An example is Plan Ceibal [5] which provides connectivity to every school in Uruguay, where WMNs are used to reach suburban and rural schools. Lately even service providers are beginning to use this technology, resulting in an increasing presence of carrier-class equipment in the market [6].

Under this scenario, the typical architecture (see Fig. 1) includes one or more internet gateways and several relay routers. Clearly, this intermediate routers increase the coverage of the access network without requiring more, and probably expensive, connections to the internet. However, several problems arise that are specific of this kind of architectures.

The main challenge for this kind of networks, at the wireless mesh backbone level, is routing and forwarding. In the current IEEE 802.11s standard [7] (and in several other proposals [8]) each link has an associated metric value as cost. This cost is expected to change over time, and reflect current conditions (propagation conditions, interference, etc.), so as to maximize a certain criteria (e.g. throughput). To choose a path to its destination, each router executes a shortest path algorithm. This procedure is essentially the same than the one used in wired networks. The main difference is that, just like in the internet until the early eighties, link costs are allowed to change at a time scale of some seconds [9]. The more static configuration that is used nowadays is due to the oscillations caused by these dynamic costs. It seems like history is repeating itself, since early experiments with WMNs have also reported routing oscillations [10,11].

However, a completely static routing approach is not a suitable solution in this context. Static means non-optimized routing. In the wired case this is not such a big issue, since resources, specially in the core, are relatively inexpensive (in fact, most core networks are overprovisioned). On the contrary, in wireless networks resources are intrinsically scarce, and "upgrading" a link's capacity is not always a possibility. Available resources must then be used at its maximum, and for this purpose a certain form of dynamism must be implemented in the network.

We present a novel approach which separates routing from forwarding, just like MPLS does in the wired context. That is to say, each ingress router has several possible paths towards the destinations, and these paths remain unchanged as long as no topological change takes place (e.g. a node failure). Please note that in the context of WMNs we may safely assume that nodes are fixed and do not change status nor position very often. Each new incoming flow will be forwarded along one of these paths, a decision that each ingress router will take depending on the current network condition We shall call this procedure dynamic load-balancing. We propose one such scheme that forwards incoming packets along several pre-established paths in order to minimize a certain objective function. If correctly designed, load-balancing will bring improved performance over static routing, without the difficult to avoid

* Corresponding author.
E-mail addresses: gcapde@fing.edu.uy (G. Capdehourat), flarroca@fing.edu.uy (F. Larroca), belza@fing.edu.uy (P. Belzarena).

0140-3664/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comcom.2013.10.007

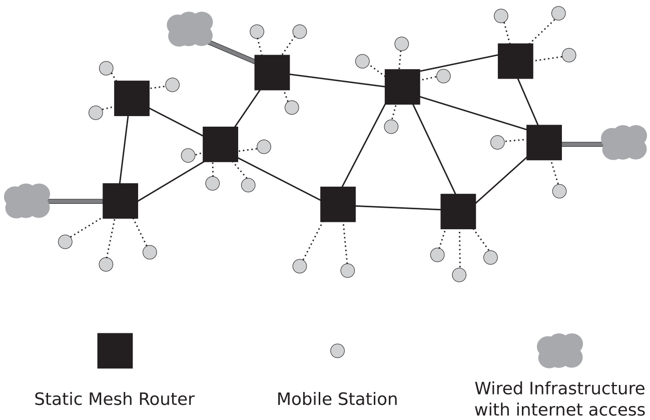

Fig. 1. Wireless Mesh Network (WMN) typical architecture.

oscillations of pure dynamic routing. For more arguments in favour of load-balancing see the discussion presented in [12], where Caesar et al. argue for a separation of timescale between *offline* computation of multiple paths and *online* spreading of load over these paths, or the analysis by Pham et al. [13] where single-path and multi-path routing protocols are compared in a wireless networks scenario, showing that the latter provides better performance.

We consider a particular but very typical scenario: a planned WMN, where all bidirectional point-to-point links do not interfere with each other. This assumption means either that all backhaul links use different channels or that links in the same channel are in different collision domains. There are many scenarios where this assumption holds, for example suburban or rural area networks and even campus networks, deployed with high directional antennas with proper RF design and channel assignment. This assumption also implies that the network topology is already defined, typically at infrastructure deployment phase. This means we cannot decide which backhaul links to establish but only how to use them, i.e. which traffic route through them.

The question that remains is to what purpose should load-balancing serve and be worthwhile. That is to say, what function of the traffic distribution should be optimized (where "traffic distribution" refers to the portion of traffic sent along each path). In this paper we argue that this function should be the sum over all nodes' interfaces of the corresponding average queue length. As shall be discussed in Section 3, this is a very versatile and important performance indicator. The problem we address is then to find the traffic distribution that minimizes the sum over all interfaces of the average queue size. However, instead of relying in analytical expressions based on (arbitrary) models, we will strive at reflecting reality as much as possible, and design a measurement-based scheme. In this framework the relationship between the average queue length and the current traffic distribution will be learned from measurements, and the optimization shall be performed based on this learned function.

This kind of approach, using a network model developed from measurements of queue sizes and traffic loads, has already proved suitable for a wired scenario [14]. In this work, we extend the framework to the previously described wireless scenario. Furthermore, we also consider the dynamic gateway selection problem and we obtain a load balanced solution using the proposed approach. Differently to the wired case, in the considered wireless scenario the average queue size at a given interface now depends not only on the incoming traffic, but also on the activity of the interface at the other end of the link. We model each link with only one average queue (the sum of both interfaces involved) which depends on the traffic in both link directions. A method to learn this bi-variable function is presented, whereas simulations illustrate the framework.

It is important to highlight that we are considering a WMN where links performance is stable and predictable, with a strong correlation between the error rate and the received signal strength. In the context of WMN, as stated in [15], interference (and not multipath fading) is the primary cause of unpredictable performance. In the scenario of interest there is no internal interference, so we expect to have a proper model with the proposed learning technique.

In a nutshell, the contributions of this paper are the following. We propose a load-balancing framework for multipath forwarding in 802.11 WMNs and we show the advantages for this kind of networks. We compare the performance of the proposed method with static routing through shortest path and dynamic routing using 802.11s. Several simulations over canonical topologies show the advantages of the proposed scheme over the alternatives. The proposed framework also copes with the gateway selection problem, typically present in WMNs. The deployment of WMNs in recent years has grown and is expected to continue rising, so it becomes essential to find a proper routing/forwarding to provide adequate service to the also growing traffic demands. The results we present suggest that dynamic load-balancing is an excellent candidate.

The rest of the paper is structured as follows. In the next section we describe some previous work and highlight some recent papers. In Section 3 we introduce the network model and most of the notation used in the paper. The paper continues in Section 4 where we describe the procedure for learning the congestion function model from measurements, while in Section 5 we detail the operation of the proposed method. Finally, in Section 6 we present the simulation experiments and performance comparison, while conclusions and future work are discussed in Section 7.

2. Related work

In the context of WMNs, several previous works presented new metrics for single path routing that take into account information from lower layers [8]. The need to increase the WMNs capacity led to the use of nodes with multiple radio interfaces which was analyzed in [16,17]. In this paper we consider a planned WMN, where all links do not interfere with each other. Even in an unplanned scenario several algorithms have been proposed [18–20] which could be used to schedule the links so that they do not interfere with each other.

There are some recent related works that we would like to highlight. In [21] an optimization framework is presented to reach minimum average delay per packet in a single channel WMN. Starting from a Markov chain model for the medium access of a single node, they derived a closed form representation for the average system delay which is used as the objective function. The model takes into account the neighbours interference but several parameters of the Markov chain need to be calculated or defined which could difficult the implementation.

Another work that uses an analytical model in the context of single channel WMNs is [22]. In particular, the authors developed a queuing-based model which is used to estimate the network capacity and to identify network bottlenecks. Based on a load-aware routing metric they choose the corresponding path for each new incoming flow, and then based on the model a centralized entity performs admission control to guarantee network stability. They focused on per-flow performance and compare the results with shortest-path first routing algorithm.

Concerning dynamic gateway selection, in [23] an heuristic algorithm was proposed to tackle the problem. A single channel