Mobile Netw Appl (2009) 14:649–660
DOI 10.1007/s11036-009-0156-7

Design and Performance Evaluation of IAR: Interference-Aware Routing Metric for Wireless Mesh Networks

Sonia Waharte · Brent Ishibashi · Raouf Boutaba · Djamal-Eddine Meddour

Published online: 5 February 2009

Abstract Multihop wireless mesh networks are an attractive solution for providing last-mile connectivity. However, the shared nature of the transmission medium makes it challenging to fully exploit these networks. In an attempt to improve the radio resource utilization, several routing metrics have been specifically designed for wireless mesh networks. However none of these routing metrics efficiently tackles interference issues. Moreover, although some evaluations have been conducted to assess the performance of these metrics in some contrived scenarios, no overall comparison has been performed. The contributions of this paper are consequently twofold. First, we propose a new routing metric, Interference-Aware Routing metric (IAR), specifically designed for WMNs. IAR uses MAC-level information to measure the share of the channel that each link is able to utilize effectively. As a result, paths that exhibit the least interference will be selected to route the data traffic. Then we evaluate the performance of IAR against some of the most popular routing metrics currently used in wireless mesh networks: Hop Count, Blocking Metric, Expected Transmission Count (ETX), Expected Transmission Time (ETT), Modified ETX (mETX), Network Allocation Vector Count (NAVC) and Metric of Interference and Channel-Switching (MIC). We show under various simulation scenarios that IAR performs the best in terms of end-to-end delay and packet loss, and provides the fairest resource utilization.

Keywords mesh networks · interference · routing metrics · QoS · evaluation

S. Waharte · B. Ishibashi · R. Boutaba
David R. Cheriton School of Computer Science,
University of Waterloo,
Waterloo, Ontario, Canada N2L 3G1

S. Waharte
e-mail: spwaharte@cs.uwaterloo.ca

B. Ishibashi
e-mail: bkishiba@cs.uwaterloo.ca

R. Boutaba
e-mail: rboutaba@cs.uwaterloo.ca

D.-E. Meddour (✉)
Orange Labs, 2 avenue Pierre Marzin,
22300 Lannion, France
e-mail: djamal.meddour@orange-ftgroup.com

1 Introduction

In response to the increasing demand for ubiquitous low latency, high volume communication, the deployment of Wireless Mesh Networks (WMNs) has become an attractive alternative to wired solutions, 3G cellular systems, and WLANs. WMNs can offer high levels of service coverage, while requiring relatively inexpensive deployment costs. Initial deployments [9, 10] have demonstrated WMNs' tremendous potential and market value. WMNs have been utilized to inexpensively share Internet connections in low-income community networks (Meraki, NetEquality), and for deploying coverage across university campuses (e.g. MIT, University of Arkansas). As a result, several companies including Nokia, Microsoft, and Intel are actively promoting full IP-based solutions for WMNs.

Nonetheless, despite the technological progresses and the possibility to concurrently transfer data on multiple channels, transmission rates remain limited compared to the ones offered in wired networks. Consequently, adequate resource management become necessary to meet consumers' increasing demand for quality-of service.

Creating the necessary resource management framework starts with an effective routing protocol, and particularly an adapted routing metric. Recently, the number of proposals of routing metrics tailored for wireless mesh networks has flourished. Through different strategies, the proposals try to evaluate the levels of interference and route the traffic flows around the most congested areas. But so far, none of them has been widely adopted. Several reasons can explain this:

Level of complexity: At the opposite of some topological or traffic-related parameters that can be easily obtained, measuring the level of interference is a challenging task. The channel quality can be hard to assess as it changes in space and time. A sender and receiver can potentially suffer from different levels of interference, that can lead to a poor quality of communication (with a high packet loss) if the transmission rates are not properly adjusted. Also in IEEE 802.11-based networks, the shared nature of the transmission medium makes it difficult to properly evaluate a link utilization. Indeed, it is necessary to account not only for the traffic flows going through a particular link, but also for all the traffic flows going through links at interference range. Unless exact information on the traffic characteristics for all the nodes can be maintained and assuming a perfect data scheduling, only a rough approximation of the actual network status can be obtained. Moreover, the exchange of control messages is usually required to propagate link quality measurements. The cost involved in these operations can overshadow the actual improvement obtained by avoiding lossy or congested links.

Lack of comparisons: To the best of our knowledge, no complete evaluation of the existing contributions in this area has been performed. Each proposed metric has been evaluated in some limited scenarios, with specific parameters, and compared with only a small subset of the existing routing metrics [1].

Lack of insights: The existing evaluations of the different routing metrics for wireless mesh networks have only been conducted for some very contrived scenarios. Insights on the metrics efficiency in different situations have rarely been provided. It is therefore difficult to extrapolate on the performance of a particular metric if different network settings are considered.

Previous experiments conducted in [4] have shown that currently implemented metrics (Hop Count, Expected Transmission Count, Expected Transmission Time) perform similarly. This suggests that the metrics are essentially equivalent. In fact, only Hop Count, the simplest metric, distinguishes itself in mobile networks, as the other metrics do not adapt quickly enough to topology changes [3].

To deal with the above limitations, we propose a routing metric that evaluates each link's effective share of the medium. The Interference-Aware Routing metric (IAR) MAC-level measurements determine the percentage of time each transmission wastes due to interference from other nodes. This wastage occurs in the form of backoff and waiting time, as well as failed transmissions. Routing using IAR selects links that experience the least interference.

We demonstrate the benefit of IAR through simulations, comparing to some of the most commonly used WMN routing metrics—Hop Count, Blocking Metric, ETX, mETX, ETT, NAVC and MIC.

The remainder of this paper is organized as follows. In Section 2, we define a set of criteria against which the chosen metrics will be compared. We then discuss the implementation of existing routing metrics in Section 3. In Section 4, we describe our interference-aware routing metric. The results of the evaluations are presented in Section 5. Finally, we conclude the paper in Section 6.

2 Route selection parameters

Routing in WMNs extends network connectivity to end users through multi-hop relays. Packets can be routed via one or multiple paths, possibly using several different channels. Depending on the application requirements, a routing protocol can focus on optimizing one or more routing metrics. Path length, end-to-end delay, and packet loss represent some parameters whose importance varies depending on the level of quality requested by an application. Interference should also be accounted for, as if can result in severe performance degradation during concurrent data transmissions. The shared transmission medium constrains all nodes in the interference range of a sender or receiver to inactivity until completion of the ongoing communication. In their seminal work [5], Gupta and Kumar have shown that in a wireless network with n identical nodes, the achievable per node throughput is $\Theta(1/\sqrt{n \log n})$ with random node placement and communication pattern. Under the assumption of an optimal node placement and communication pattern, this throughput becomes $\Theta(1/\sqrt{n})$.

Interference can occur:

- Within a single flow (intra-flow interference): a communication between two nodes on one path can block the upstream and downstream nodes that are within interference range.
- Between multiple flows from one or multiple sources (inter-flow interference): this phenomenon is referred to as the route coupling effect.

Therefore, when a new flow is to be sent across the network, it is important to realize that the actual expected performance can not simply be estimated without considering the flows already established and without considering the impact of adding this new flow on top of the existing ones. In order to give a clear overview of the focus of the routing metrics considered and before delving into the details of their design, we first define a set of criteria against which we may compare these routing metrics. This list, although not exhaustive, encompasses a set of factors that we consider to have the greatest impact on the performance of a wireless mesh network.

Different parameters can enter into the computation of a routing metric. Among them, the ones that can be considered as the most characteristic of wireless networks are the following:

- *Path Length*: The number of hops between a source router and a destination router is an important (and the most commonly used) comparison criterion as longer paths mean more self-interference (interference among links along the same path) and consequently potentially greater end-to-end delay. Flows transmitted via a long path also interfere with a greater number of links located geographically close to this path.
- *Bandwidth*: Network links can support different data rates as a result of technical limitations or in the case of wireless networks, environmental noise and signal strength. This difference in capacity affects not only the link considered but also the residual capacity of geographically close links. Indeed, the use of a lower-capacity link not only increases the transmission delay of the flow crossing the link considered, but reduces the achievable rate of neighboring transmissions by increasing their interference level. As current hardware allows rate adaptation depending on the quality of the transmission medium, obtaining and maintaining this information can help improve the network performance.
- *End-to-end Delay*: Delay-sensitive applications require bounded end-to-end delay in order to function properly. Therefore it is important to evaluate the time it takes for a packet to reach its destination, as well as to estimate the variability (jitter) over all data transmissions.
- *Interference*: Owing to the shared nature of the transmission medium, nodes transmitting on the same channel can interfere with each other if they are located in the same geographical area. Integrating interference into the design of the routing metric can therefore help to combat network congestion and increase overall network performance.
- *Packet Loss*: Channel quality can be assessed by estimating the number of retransmissions necessary for a transmission to be successfully performed.
- *Effective Link Share*: As access to the transmission medium is shared among nodes located in the same area, a communication on a particular link is affected by the transmissions on neighboring links. It follows that a node may have to wait for concurrent communications to complete before it is able to send its own data. Obtaining an estimate of the channel occupation (and therefore the congestion level) is therefore a desirable task.

We will also distinguish the level of complexity of the routing metrics based on some implementation parameters such as:

- *Per-node/Per-link metric*: A per-link metric can potentially allow fine-grained information of each link to be maintained, whereas a per-node metric assumes by default that all the links attached to a node have the same cost. On the downside, a per-link metric might be costly to maintain (e.g. by incurring extra control messages).
- *Knowledge*: A metric can be computed based on different information: packet loss, number of nodes, number of neighbors, traffic characteristics, etc.
- *Interference*: Different strategies with different levels of complexity might be implemented to account for the interference.

3 Existing routing metrics description

In this section, we present some routing metrics that are currently used in WMNs. They were either specifically tailored for WMNs or previously developed for other types of networks (e.g. ad hoc networks) but adopted for use in WMNs due to the underlying similarities with WMNs. We consider the following metrics: Hop Count, Blocking Metric, Expected Transmission Count (ETX), Expected Transmission Time (ETT), Modified

Expected Number of Transmissions (mETX), Network Allocation Vector Count (NAVC) and Metric of Interference and Channel-Switching (MIC).

3.1 Hop count

Hop count is the most commonly used metric in wireless multihop networks. The path selected is the one minimizing the number of links between a given source and destination node. It became very popular in ad hoc networks due to its ease of computation as it only considers the route length as the differentiating criterion. However, on the downside, this routing metric fails to account for the specifics of wireless environments (links may have different transmission rates, loss ratios, etc.) and it does not consider the congestion level resulting from the shared use of the transmission medium.

3.2 Blocking metric

A simple improvement over hop count has been presented in [11] in order to account for the interference along a certain path. In this work, the interference level referred to as the Blocking Value, is defined as the number of neighbors a node is interfering with. Each node is therefore weighted according to this Blocking Value. The Blocking Metric of a path is then defined as the sum of all the blocking values along the path. Paths with minimum cost will consequently be used to carry the traffic flow.

This technique presents the advantage of being simple, without any additional overhead other than to maintain some information on the number of neighbors. However, this metric still does not incorporate any characteristics concerning the traffic flow or link capacity and only superficially addresses the issue of interference. Little improvement over hop count is therefore to be expected.

3.3 Expected transmission count (ETX)

Expected Transmission Count is defined as the number of transmissions required to successfully deliver a packet over a wireless link [2]. The ETX of a path is then defined as the sum of the ETX of each link along the path. Let p_f and p_r be the packet loss probability in the forward and reverse directions. The probability p of an unsuccessful transmission is:

$$p = 1 - (1 - p_f)(1 - p_r) \tag{1}$$

Therefore, the expected number of transmissions to successfully deliver a packet in 1 hop can then be expressed as:

$$ETX = \sum_{k=1}^{\infty} kp^k (1-p)^{k-1} = \frac{1}{1-p} \tag{2}$$

The delivery ratios are measured using 134-byte probe packets. One probe packet is sent every τ second (set to 1 s in the experiments that follow later in the chapter). The packet loss ratio is computed by counting the number of probe packets received over a predetermined period of time (10 s in the experiments).

ETX favors paths with higher throughput and lower number of hops as longer paths have lower throughput due to increased self-interference. However, this metric does not consider differences in transmission rates. It does not completely account for the interference on the transmission medium as the sender of a probe packet can defer its transmission if it senses the channel is busy. As the transmission rate of the probe packets is typically low, it does not give a good indication of how busy a link really is. It also does not give any indication of the effective link share.

3.4 Expected transmission time (ETT)

ETT is an improvement over ETX as it includes the bandwidth in its computation [3]. Let S be the packet size and B the bandwidth of the link considered, then ETT is computed as follows:

$$ETT = ETX \frac{S}{B} \tag{3}$$

In a similar fashion to ETX, the expected transmission time of a path is computed according to the sum of the links' ETT along the path.

The authors later improved over ETT by proposing a Weighted Cumulative ETT (WCETT) [3]. This metric was designed to favor channel-diverse paths. For a path p, WCETT is defined as follows:

$$WCETT(p) = (1-\beta) \sum_{\text{link } l \in p} ETT_l + \beta \max_{1 \le j \le k} X_j \tag{4}$$

where β is a tunable parameter less than 1 and X_j represents the number of times channel j is used along path p.

Nevertheless, this metric still suffers from the same limitations as ETX/ETT by not estimating the effective link share and does not completely capture the inter-flow interference.

3.5 Modified expected number of transmissions (mETX)

An enhancement over ETX has been proposed by [7] based on the observation that ETX does not account for the channel variability and only considers the average channel behaviour. The authors therefore defined mETX as follows:

$$mETX = exp\left(\mu_\Sigma + \frac{1}{2}\sigma_\Sigma^2\right) \quad (5)$$

where μ_Σ and σ_Σ^2 represent the mean and variability of the error probability.

The main challenge in the implementation of this metric is to properly model and quantify the variability of the transmission channel.

3.6 Network allocation vector count (NAVC)

NAVC [8] essentially accounts for the interflow interference by averaging the values of the Network Allocation Vector experienced by a node along a link for a given observation period. According to the value obtained, a level of congestion is attributed to the node. During the route discovery process, two parameters, *heavy_node_number* and *nav_sum*, are maintained. Upon reception of a ROUTE REQUEST packet, a node has therefore three options depending on the value of the measured NAVC.

1. If $NAVC > 0.65$: increase *heavy_node_number* by 1 and add the square of NAVC to *nav_sum*;
2. If $0.25 \leq NAVC \leq 0.65$: increase *nav_sum* by the square of NAVC;
3. If $NAVC < 0.25$: do nothing.

The cost of a path comprises the sum of the *heavy_node_number* of each node along the path and the sum of the *nav_sum*. Paths are therefore given priority first depending on the *heavy_node_number* and then on the *nav_sum*.

3.7 Metric of interference and channel-switching (MIC)

MIC has been designed to improve over WCETT by capturing more information on the effective link share [12]. For a network composed of N nodes and a path p, MIC averages the time to transmit on a particular link over the minimum time to transmit over all the existing links. Similarly to WCETT, MIC adds a term to account for channel diversity called Channel Switching Cost (CSC).

$$MIC(p) = \frac{1}{N \times \min(ETT)} \sum_{\text{link } l \in p} IRU_l + \sum_{\text{node } i \in p} CSC_i \quad (6)$$

$\min(ETT)$ represents the smallest ETT in the network and IRU_l represents the interference-aware resource usage defined as:

$$IRU_l = N_l \times ETT_l$$

$$CSC_i = \begin{cases} w_1 \text{ if } CH(prev(i)) \neq CH(i) \\ w_2 \text{ if } CH(prev(i)) = CH(i) \end{cases}$$

$$0 \leq w_1 < w_2$$

N_l is the number of nodes link l is interfering with, ETT_l is the expected transmission time on link l, $CH(i)$ is the channel assignment of node i and $prev(i)$ represents the node before node i along path p. IRU_l can therefore be interpreted as the total channel time consumed by link l. CSC is a weight allocated to a link as a function of the channel used by the link preceding the link considered on a particular path. If both links use the same channel, a greater weight is assigned to the link.

This metric presents some major drawbacks in terms of implementation. First, the overhead required to maintain up-to-date information of the ETT for each link can significantly affect the network performance depending on the traffic load. Second, this metric assumes that all the links located in the collision domain of a particular link contribute to the same level of interference, which does not take into account the difference in traffic load at each node.

3.8 Routing metrics summary

Table 1 summarizes the characteristics of each of the metrics just discussed. In particular, we highlight if the metric is computed on a per-node or per-link basis, what information is required in the computation (number of nodes, neighbors, packet loss, etc.) and how it handles interference, if applicable.

ETT (and its extension WCETT) satisfies most of the criteria that we identified as important for WMNs but still fails to provide any information on the effective link share. MIC takes into account the number of neighbors for each node but its computation is expensive and only provides an estimation of the actual link utilization. In this paper, we address this issue by proposing a novel routing metric based on the evaluation of

Table 1 Comparison of routing metrics

	Per node	Per link	Knowledge	Interf. awareness
Hop count	X		None	N/A
Blocking Metric	X		Nb. of neighbors	Nb. of neighbors
ETX		X	Link pkt loss	Per-link pkt loss ratio Averaged over time
ETT		X	Link pkt loss Bandwidth Packet size	Per-link pkt loss ratio Averaged over time
mETX		X	Link pkt loss Channel variability	Per-link pkt loss ratio
NAVC	X		NAV	Node waiting time
MIC		X	Link pkt loss	Per-link pkt loss ratio Avg over time and over neighbors

the effective link share. We discuss its implementation when a single channel is used and describe how to extend this metric to integrate multiple channels. As our work is solely focused on wireless mesh networks that are characterized by a fixed backbone, we are not concerned with node mobility.

4 Interference-aware metric

4.1 Motivations and design choices

The impact of interference on the network performance is a difficult parameter to estimate. However, integrating interference into the design of a WMN routing protocol is of paramount importance. Interference can be considered as a measure of the quality of the transmission channel. If the channel quality is poor, a packet has a high probability of requiring several retransmissions before successfully reaching its destination. Measuring interference also gives an estimation of the network utilization level. If several concurrent transmissions occur in the neighborhood of a source-destination pair, the nodes within transmission distance have to wait for the medium to be cleared before they have access to it. The higher the number of nodes, the greater the probability of collision due to simultaneous transmissions. Consequently, deriving a metric that is able to account for these different states can increase the network performance by avoiding lossy links and congested zones.

Therefore, we believe that a metric best suited for WMNs should incorporate the following characteristics.

- *Low overhead*. Exchange of control messages on the link status can be costly in terms of resource usage. It is therefore preferable to favor a non resource consuming solution based on local monitoring.
- *Interference-Awareness*. Both intra- and inter-flow interference have to be accounted for. This means that it is necessary to account for the waiting time as well as the number of retransmissions due to packet loss.
- *Differentiation on link capacities*. Not all the links have the same transmission rates due to environmental noise or technological limitations. Higher capacity links should be favored when they are not congested.
- *Channel diversity*. If the network nodes are embedded with multiple interfaces, this should be exploited to favor the use of high-quality links (higher transmission rate, less packet loss) and by reducing the interference by spreading the traffic over multiple channels.

4.2 IAR: description

Before we describe the actual computation of our proposed metric, it is important to have a clear view of the different states in which a node can be. There are five states:

- *Idle*: The node does not have any packets of its own to transmit neither does it have any packets to relay. It therefore does not contribute to increasing the interference in the network and should consequently be ignored.
- *Success*: The state refers to the case where a node has successfully received the acknowledgment of the packet it has sent.
- *Collision*: In this state, a node sent a DATA packet but never received an acknowledgement for the packet. Either the receiver node was in the range of another transmission and therefore received

several packets at the same time. Or the receiving node was itself initiating a communication.

- *Wait*: As only one communication can occur at the same time in the same geographical area, if a node senses the medium is busy, it has to wait until the ongoing communication is completed before it starts its own.
- *Backoff*: Even though a node has some data to transmit and the medium is free, IEEE802.11 Standard enforces a random waiting period (during which the medium has to remain idle) before it starts sending its data.

The period of time between the moment when a node generates a packet (or receives a packet it then has to relay) and the moment it successfully transmits the packet to the next hop node (possibly the destination of the packet) is a succession of *Success*, *Collision*, *Wait* and *Backoff* states (Fig. 1).

We therefore designed a routing metric, *Interference-Aware Routing metric* (IAR), that could address the shortcomings of the existing metrics we previously highlighted.

In particular, it more realistically reflects the link usage, and includes all possible states a node is in - particularly the waiting periods caused by neighboring nodes' transmissions.

Let $T_{Success}$, T_{Wait}, $T_{Collision}$ and $T_{Backoff}$ be the time spent respectively in the *Success*, *Wait*, *Collision* and *Backoff* states. The communication cycle is defined as the period between the generation of a packet up to its successful transmission. The duration of each state is summarized in Table 2. We have ignored in the computation the propagation delay although it is accounted for in the actual implementation. The Wait state has a variable duration as it depends on other nodes' transmissions.

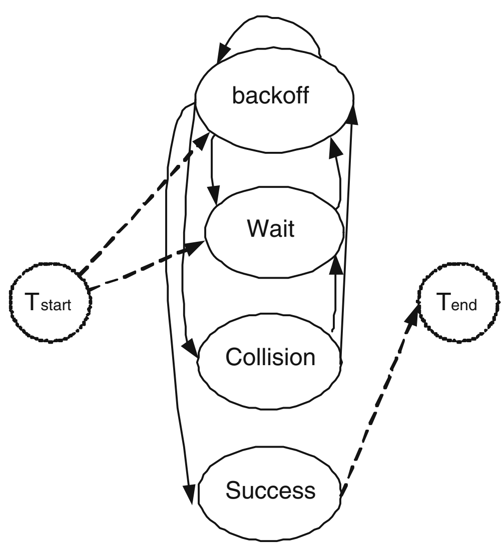

Fig. 1 Four of the communication states of a node (IDLE state is not represented)

Table 2 Duration for the 4 channel states

State	Duration
Backoff	Time slot
Wait	*Variable*
Collision	DATA+SIFS+ACK+DIFS
Success	DATA+SIFS+ACK+DIFS

For each link, we calculate the *unproductive busyness* α_{ub}, that is to say the percentage of time spent in states in which communication on this link is not possible.

$$\alpha_{ub} = \frac{T_{Wait} + T_{Collision} + T_{Backoff}}{T_{Wait} + T_{Collision} + T_{Backoff} + T_{Success}}$$

Therefore, for a link l, IAR is defined as:

$$IAR(l) = \frac{1}{1-\alpha_{ub}} * \frac{S}{B} \quad (7)$$

IAR can be interpreted as the time to transmit a packet of size S over a medium of actual bandwidth $(1-\alpha_{ub}) * B$.

The cost of a path p is consequently defined as the sum of the cost of each link along the path.

$$IAR(p) = \sum_{l \in p} IAR(l) \quad (8)$$

The amount of time spent in each of these states can be determined by passive measurements using the actual traffic in transmission or by active probing. Similar to ETT, IAR can be modified to handle the multi-channel scenario with the addition of a switching channel cost factor (cf. the computation of WCETT).

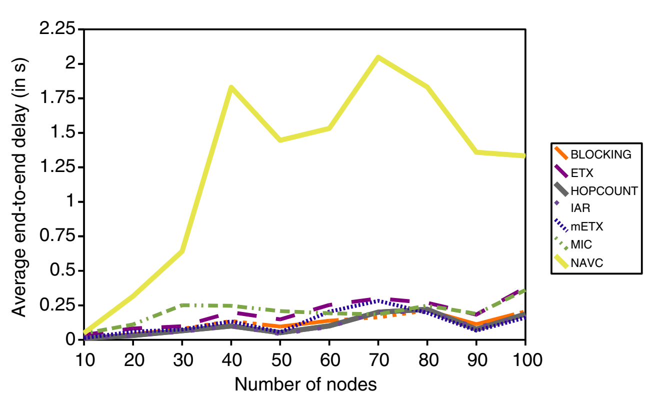

Fig. 2 End-to-end delay with increasing number of nodes

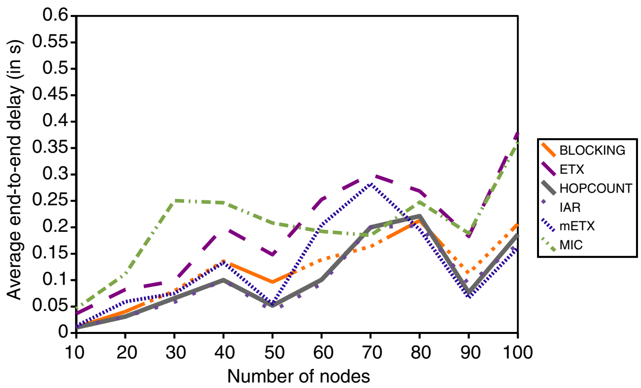

Fig. 3 Closer look at the end-to-end delay with increasing number of nodes (excluding NAVC)

5 Performance evaluations

5.1 Implementation details

We compared the performance of the routing metrics through simulations implemented in NS2 [6]. We used the default settings that the network simulator provides to model the wireless transmissions: two ray ground propagation mode, 250 m transmission range and 550 m interference range. The network topologies have been randomly generated in a 2000 × 2000 m^2 area. UDP is used at the transport layer and all flows are sent at a constant bit rate, with a packet size of 512 bytes or 1512 bytes. The source and destination of each flow are randomly chosen in order to avoid the appearance of a single bottleneck. We only performed simulations in a single-channel environment. This decision was motivated by the fact that we wanted to conduct a fair comparison of the performance of the metrics, which is difficult to achieve between single and multi-channel metrics. Besides, it is worth noting that even though some metrics have not been initially designed to handle channel diversity, the addition of a cost factor similarly to what has been done for ETT or MIC can resolve this issue.

For each configuration, we evaluated the end-to-end delay, the path length and the packet loss. We assumed that all links have the same nominal capacity and that the packet size is fixed. In this context, as ETX and ETT necessarily lead to the same results, we only refer at ETX in the remainder of the experimental analysis (although the same results apply for ETT).

The packet loss ratio is determined via periodic transmissions of probing packets (sent every second in the simulations). The routing tables are recomputed periodically. To make the implementation oblivious to the specifics of a particular routing protocol, we assumed the existence of a central entity responsible for computing and keeping the routers updated with the optimal routing tables at any given time.

5.2 Simulation results

5.2.1 Impact of the network size

First, we evaluated the impact of the network size on the performance of each routing metric. We increased the size of the network from 10 to 100 nodes with 5 traffic flows of 20 pkt/s. The results obtained consist of an average of 50 simulations over all the flows. Figures 2, 3, 4, 5 and 6 show the average end-to-end delay, the average number of hops and the average loss probability. We observe that overall NAVC performs poorly in terms of delay and packet loss compared to the other

Fig. 4 Packet loss for all routing metrics with increasing number of nodes

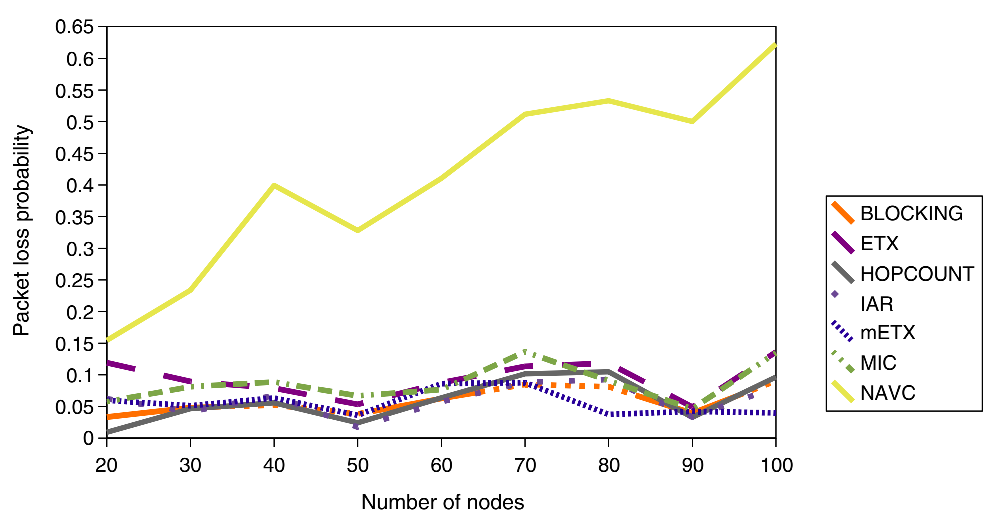

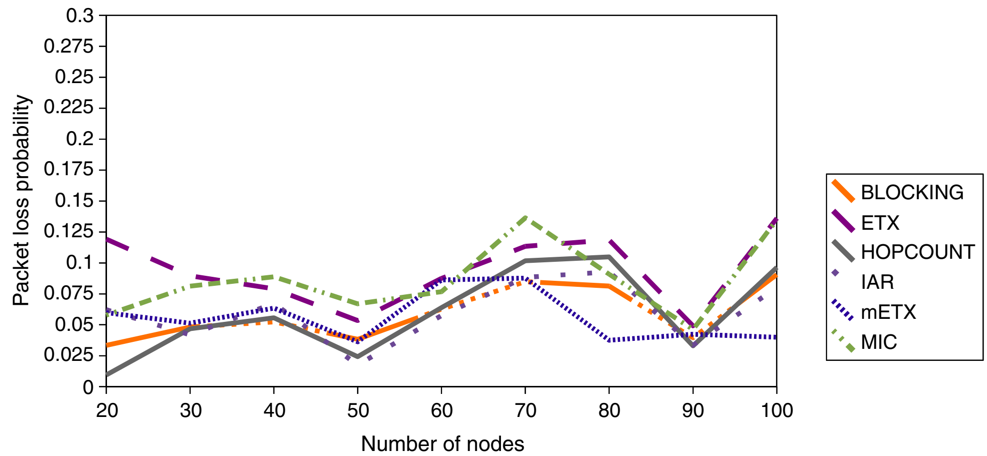

Fig. 5 Closer look to the packet loss probability for the routing metrics with increasing number of nodes (excluding NAVC)

metrics implemented. Incorporating the value of the network allocation vector in the metric computation could theoretically provide some useful information on the effective link share at each node since this parameter indicates the duration of the data transmission to be initiated. We believe that the bad performance of NAVC results from the way the threshold values for *heavy_node_number* and *nav_sum* are computed. These thresholds are solely predetermined based on simulations without any strong justification supported by analysis. Moreover the difference in link capacities is not accounted for, nor are the traffic characteristics. As the network size increases, the performance degrades significantly, eventually leading to a situation in which only flows for which the source and destination are within direct reach of each other can successfully be transmitted. This explains why the average path length is significantly better for NAVC than for the other metrics.

NAVC aside, we can observe that IAR performs the best in terms of end-to-end delay (Fig. 3) and packet loss probability (Fig. 5), followed closely by Hop Count. Hop Count favors shortest paths but at the expense of greater end-to-end delay and packet loss probability, whereas IAR avoids highly congested areas, which results in longer routing paths. In general, routing implemented with Hop Count, Blocking Metric or IAR results in path lengths on average 10 to 15% shorter than with ETX, mETX or MIC.

We also looked at the per-flow performance and computed Jain's fairness index in a 50-node network with 10 traffic flows (Fig. 7). We observe a fairer traffic load distribution in the case of IAR, ETX, mETX and MIC than with Blocking and Hop Count. This results from the fact that Blocking and Hop Count can lead to the starvation of some flows to the benefit of others. This result is not surprising as IAR, ETX, mETX and MIC favor less congested paths whereas Hop Count and Blocking Metric favor shortest but potentially more congested paths.

This first scenario demonstrates that IAR stands out as the best solution as: 1/ it offers a better or similar level of performance in terms of end-to-end delay and packet loss as Hop Count and Blocking; 2/ it offers a fairer load distribution than Hop Count and is easier to implement than ETX, mETX or MIC.

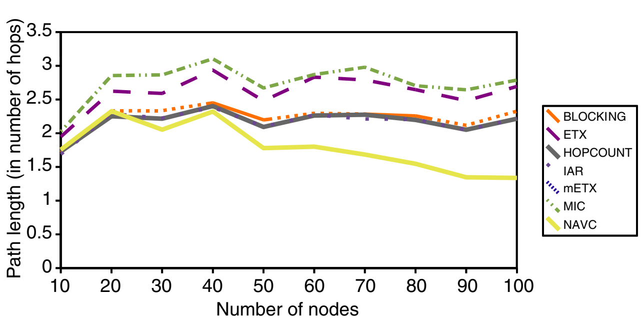

Fig. 6 Path length with increasing number of nodes

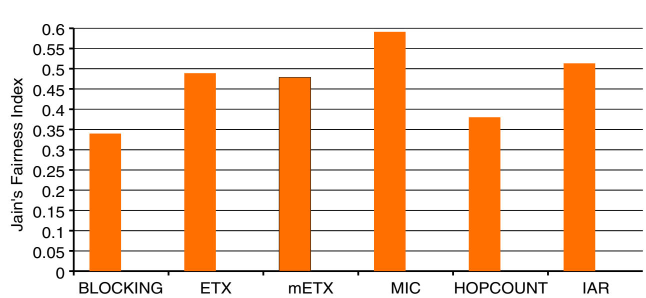

Fig. 7 Fairness analysis: we compute Jain's fairness index for a 50-node network with 10 traffic flows (1 is the best value)

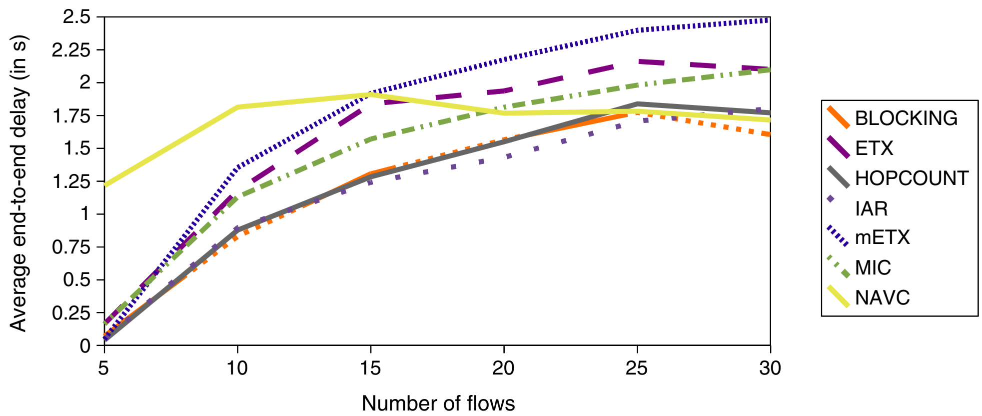

Fig. 8 End-to-end delay with increasing number of flows for a 50-node network

5.2.2 Impact of the traffic load

In the second set of simulations, we studied how the traffic load can impact the network performance by progressively increasing the number of flows from 5 to 30 for a network of 50 nodes uniformly distributed over a 2000×2000 m^2 area (Figs. 8, 9 and 10). As in the previous case, NAVC performs very poorly compared to the other routing metrics. In terms of end-to-end delay and packet loss, IAR still performs the best followed by Hop Count and Blocking. Similarly to the previous case, Hop Count leads to shorter paths than the other routing metrics but IAR, ETX and MIC lead to a fairer load distribution.

We ran similar experiments while increasing the size of the topologies. We considered networks with 100 and 150 nodes and analyzed the resulting network performance. As the path length increases, with a similar number of flows, the probability of collision increases. Therefore, flows on shorter paths have a greater chance of being successfully transmitted. Hop Count, IAR and Blocking Metric still perform the best in terms of packet loss whereas ETX, mETX and MIC, although trying to avoid congested areas, lead to a poor network utilization by electing longer paths and therefore contributing even more to the interference level.

We also analyzed the impact of the packet size on the network performance. We ran the same sets of simulations with packets of 1512 bytes. With only 5 flows, given the network characteristics, the network gets immediately congested. The packet loss probability is in the order of 70% for 5 traffic flows and goes over 80% with 30 flows. Moreover, if a packet has to be retransmitted due to a collision, a greater packet size incurs some extra delay for transmission and consequently an increased end-to-end delay. When the number of flows increases, similarly to the previous observation, the flows between the closest source-destination pairs (1 hop away) are favored, starving the other traffic flows. This is a direct consequence of the way the MAC protocol has been designed. As the number of collisions increases, the backoff time (mandatory waiting time before attempting another transmission) exponentially increases. In addition, given the packet size, retransmitting a packet due to a collision takes 3 times longer compared to the previous experiments.

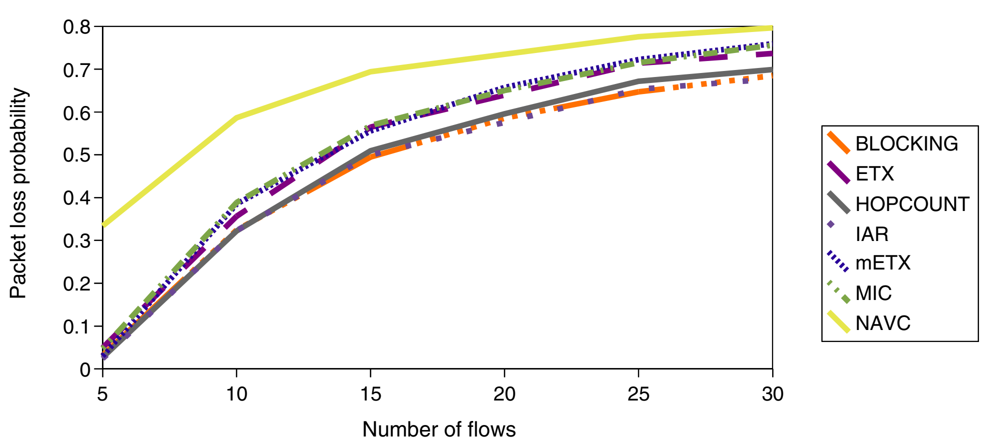

Fig. 9 Packet loss with increasing number of flows for a 50-node network

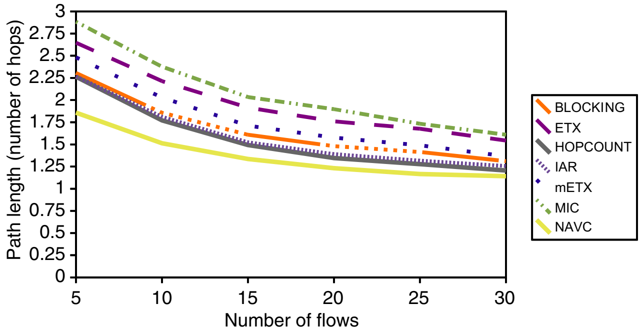

Fig. 10 Number of hops with increasing number of flows for a 50-node network

6 Conclusion

As user expectations for ubiquitous connectivity and quality service increase, wireless mesh networks represent a promising solution. By extending network coverage through the use of multi-hop wireless communication, WMNs offer versatility, along with easy and inexpensive deployment. However, routing in such networks is a challenging research issue with tremendous impact on network performance, particularly when interference is considered. With both self-interference between hops along the same path, and interference between different paths, it is important that the routing protocol integrate these effects into the routing decision.

In this paper, we have considered the abilities of various routing metrics to address interference issues in a WMN. The different metrics utilize different types and degrees of network state information. Some are simple (e.g. Hop count, Blocking), others are more sophisticated (e.g mETX or MIC). While the more advanced approaches directly consider interference issues, they also require more complex network state information. This can be costly to obtain and maintain, with overhead of control messages competing with data transmissions.

The performance of six popular routing metrics has been studied using simulation. The impact of network size and traffic loads were evaluated in terms of end-to-end delay, packet loss, and path length. These studies have demonstrated that despite the consideration of interference, the sophisticated metrics fail to consistently outperform the simple approaches. In fact, in many scenarios, particularly as traffic increases, the performance of the advanced metrics suffers. However, it appears that the simple approaches are inherently unfair.

Based on these observations, we have proposed a novel Interference-Aware Routing metric (IAR). IAR allows a node to estimate its effective share of the link capacity using local measurements. This approach accounts for intra- and inter-flow interference, as well as packet loss resulting from poor channel quality. The simulation results demonstrate that IAR outperforms hop count in many scenarios, particularly in terms of end-to-end delay. It does so, while maintaining a much fairer delivery of packets.

We believe that this work motivates the need for further work into developing an appropriate routing metric for WMNs. IAR demonstrates that such a metric can outperform the simple approaches, despite the additional overhead costs involved in collecting the required information. However, we intend to continue to investigate how channel quality can be accurately evaluated and incorporated into the metric. Additionally, the use of multiple channels and support for channel diversity must be considered in continuing to develop an interference-aware metric.

References

1. Campista MEM, Esposito PM, Moraes IM, Costa LHM, Duarte OCM, Passos DG, de Albuquerque CVN, Saade DCM, Rubinstein MG (2008) Routing metrics and protocols for wireless mesh networks. IEEE Netw 22(1):6–12, January–February
2. De Couto D, Aguayo D, Bicket J, Morris R (2003) A high-throughput path metric for multi-hop wireless routing. In: Proceedings of the 9th annual international conference on mobile computing and networking (MobiCom), San Diego, 14–19 September 2003, pp 134–146
3. Draves R, Padhye J, Zill B (2004) Comparison of routing metrics for static multi-hop wireless networks. In: Proceedings of the 2004 conference on applications, technologies, architectures, and protocols for computer communications (SIGCOMM), Portland, August 2004
4. Eriksson J, Agarwal S, Bahl P, Padhye J (2006) Feasibility study of mesh networks for all-wireless offices. In: Proceedings of the international conference on mobile systems, applications, and services (MobiSys), Uppsala, 19–22 June 2006
5. Gupta P, Kumar PR (2000) The capacity of wireless networks. IEEE Trans Inf Theory 46(2):388–404, March
6. The network simulator—ns2 (2009) The network simulator—ns2 homepage. http://www.isi.edu/nsnam/ns/
7. Koskal CE, Balakrishnan H (2006) Quality-aware routing metrics for time-varying wireless mesh networks. IEEE J Sel Areas Commun 24(11):1984–1994, November
8. Ma L, Zhang Q, Xiong Y, Zhu W (2005) Interference aware metric for dense multi-hop wireless networks. In: Proceedings of the IEEE international conference on communications (ICC), Seoul, May 2005
9. Rayner K (2003) Mesh wireless networking. Commun Eng 1(5):44–47, October–November
10. Schrick B, Riezenman MJ (2002) Wireless broadband in a box. IEEE Spectrum 39(6):38–43, June
11. Wei H-Y, Ganguly S, Izmailov R, Haas Z (2004) Interference-aware IEEE 802.16 WiMax mesh networks. In: Proceedings of the IEEE vehicular technology conference (VTC), Los Angeles, 26–29 September 2004

12. Yang Y, Wang J, Kravets R (2005) Designing routing metrics for mesh networks. In: Proceedings of the 1st IEEE workshop on wireless mesh networks (WiMesh), Santa Clara, September 2005

Sonia Waharte is a Ph.D. Candidate in the School of Computer Science at the University of Waterloo. She received her M.Sc. in Computer Engineering from the National Institute of Telecommunications, Paris, France in 2002 and her M.Sc. in Computer Science, with high honors, from the University Paris VI in 2002. Her research interests include the design and analysis of network protocols for sensor networks and resource management in wireless mesh networks.

Brent Ishibashi received the B.Sc. degree in Computer and Information Sciences from the University of Guelph, Canada, in 2000, and the M.Math degree in Computer Science from the University of Waterloo, Canada in 2004. He is presently working towards his Ph.D. degree at the University of Waterloo under the supervision of Prof. Boutaba. His research interests include wireless resource management, wireless ad hoc and mesh networks, and cognitive radio.

Raouf Boutaba is a Professor in the School of Computer Science of the University of Waterloo. Before that he was with the Department of Electrical and Computer Engineering of the University of Toronto. Before joining academia, he founded and was the director of the telecommunications and distributed systems division of the Computer Science Research Institute of Montreal (CRIM). Dr. Boutaba conducts research in the areas of network and distributed systems management and resource management in multimedia wired and wireless networks. He has published more than 150 papers in refereed journals and conference proceedings. He is the recipient of the Premier's Research Excellence Award, the NORTEL Networks research excellence Award and several Best Paper awards. He is a fellow of the faculty of mathematics of the University of Waterloo and a distinguished lecturer of the IEEE Computer Society. Dr. Boutaba is the Chairman of the Working Group on Networks and Distributed Systems of the International Federation for Information Processing (IFIP), the Vice Chair of the IEEE Communications Society Technical Committee on Information Infrastructure, and the Director of standards board of the IEEE Communications Society. He is the founder and acting editor in Chief of the IEEE eTransactions on Network and Service Management, on the advisory editorial board of the Journal of Network and Systems Management, on the editorial board of the KIKS/IEEE Journal of Communications and Networks, the editorial board of the Journal of Computer Networks and the Journal of Computer Communications. He has also served as a guest editor of several special issues of IEEE Journal of Selected Areas in Communications (JSAC), the Journal of Computer Networks, the Journal of Computer Communications and the Journal of Network and System Management. He acted as the program chair for several conferences including IFIP Networking, IEEE CCNC, IEEE/IFIP NOMS, IFIP/IEEE MMNS, IEEE FIW, IEEE ACC and symposia in IEEE ICC.

Djamal-Eddine Meddour He received in 2000 the B.S. degree in computer science engineering from INI (Institut National de formation en Informatique), Algiers, Algeria, his M.S. degree in advanced network architecture from Versailles University, France, in 2001, and his Phd in 2004 from Paris 6 University. Dr. Meddour has a number of refereed publications and standards contributions and holds several patents. He is currently with Orange Labs (France Telecom Group's R&D), his research interests include Home networking, QoS routing in wireless mesh networks, and resource management in multimedia wired and wireless networks.