Volume 31, issue 5, 25 March 2008 ISSN 0140-3664

ELSEVIER

computer
communications

Available online at
ScienceDirect
www.sciencedirect.com

Special Issue:
Mobility Management and Wireless Access
Guest Editors: Luciano Bononi
Albert Y. Zomaya

Available online at www.sciencedirect.com

ScienceDirect

Computer Communications 31 (2008) 885–895

computer communications

www.elsevier.com/locate/comcom

Performance analysis of multi-radio AODV in hybrid wireless mesh networks ☆

Asad Amir Pirzada [a,b,*], Marius Portmann [a,b], Jadwiga Indulska [a,b]

[a] *Queensland Research Laboratory, National ICT Australia Limited, Level 20, 300 Adelaide Street, Brisbane, Qld 4000, Australia*
[b] *School of ITEE, The University of Queensland, Brisbane, Qld 4072, Australia*

Abstract

Wireless Mesh Networks (WMNs), based on commodity hardware, present a promising technology for a wide range of applications due to their self-configuring and self-healing capabilities, as well as their low equipment and deployment costs. One of the key challenges that WMN technology faces is the limited capacity and scalability due to co-channel interference, which is typical for multi-hop wireless networks. A simple and relatively low-cost approach to address this problem is the use of multiple wireless network interfaces (radios) per node. Operating the radios on distinct orthogonal channels permits effective use of the frequency spectrum, thereby, reducing interference and contention. In this paper, we evaluate the performance of the multi-radio Ad-hoc On-demand Distance Vector (AODV) routing protocol with a specific focus on hybrid WMNs. Our simulation results show that under high mobility and traffic load conditions, multi-radio AODV offers superior performance as compared to its single-radio counterpart. We believe that multi-radio AODV is a promising candidate for WMNs, which need to service a large number of mobile clients with low latency and high bandwidth requirements.

Keywords: Multi-radio; Routing; Wireless mesh network

1. Introduction

Wireless Mesh Networks (WMNs) have recently gained considerable popularity due to their low cost, self-configuring and rapid deployment capabilities. Application scenarios that have been proposed for WMNs include building automation, intelligent transportation systems, metropolitan area networks and public safety applications. In their most generic form, WMNs are comprised of a combination of relatively static `MESH_ROUTER`s and mobile `MESH_CLIENT`s. `MESH_ROUTER`s form the backbone infrastructure and provide connectivity between `MESH_CLIENT`s. They may also provide access to a wired network. Due to their mobility, `MESH_CLIENT`s are typically more resource constrained than `MESH_ROUTER`s.

Depending upon their architecture and deployment configuration, WMNs can be broadly categorised into three main types [1]: Infrastructure mesh, client mesh and hybrid mesh networks. In the infrastructure mesh scenario, the `MESH_ ROUTER`s collectively provide a wireless backbone infrastructure. This is similar to a traditional WLAN, with the key difference that the wired backbone is replaced with a wireless multi-hop network. `MESH_CLIENT`s simply access the network directly via a single wireless hop to the nearest `MESH_ROUTER`. In this architecture, clients have a passive role and do not contribute to the mesh network infrastructure. In a client mesh architecture, the network is made up of mobile `MESH_CLIENT`s only, and no dedicated network infrastructure is involved. Therefore, `MESH_CLIENT`s need to perform additional network

---

☆ This is an extended version of the paper titled "Evaluation of Multi-Radio Extensions to AODV for Wireless Mesh Networks" published in the proceedings of the 4th ACM International Workshop on Mobility Management and Wireless Access (MobiWac 2006).

[*] Corresponding author. Address: Queensland Research Laboratory, National ICT Australia Limited, Level 20, 300 Adelaide Street, Brisbane, Qld 4000, Australia. Tel.: +61 7 30000503; fax: +61 7 30000480.

E-mail address: asad.pirzada@nicta.com.au (A.A. Pirzada).

0140-3664/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2007.12.012

functions such as routing and packet forwarding. A client mesh WMN is essentially identical to a traditional pure ad-hoc network [2].

A hybrid mesh architecture, such as illustrated in Fig. 1, is the most generic type of WMN, combining the concepts of infrastructure and client mesh networks. A hybrid WMN consists of relatively static `MESH_ROUTER`s which form the backbone of the network. In addition, mobile clients can act as a dynamic extension of the static infrastructure part of the network, by implementing routing and packet forwarding functionality. The hybrid mesh architecture is very flexible and allows combining the benefits of both infrastructure and client meshes. In this paper, we primarily focus our evaluation on hybrid mesh networks. However, in future we intend to extend our evaluations to infrastructure and client mesh networks.

Ad-hoc routing protocols are promising candidates for hybrid and client WMNs, due to their capability to deal with highly dynamic environments. DSR [3] and AODV [4] are two well-known ad-hoc routing protocols, which are currently undergoing extensive research in regard to their suitability for mesh networks [5,6]. So far, these protocols have been studied exclusively in single-radio mode. This paper presents an evaluation of multi-radio AODV in a hybrid WMN scenario. We performed extensive simulations to gather various performance metrics for networks containing `MESH_ROUTER`s with multiple wireless interfaces. We show that in contrast to single-radio AODV, the multi-radio variant of AODV can effectively support high traffic loads and `MESH_CLIENT` mobility.

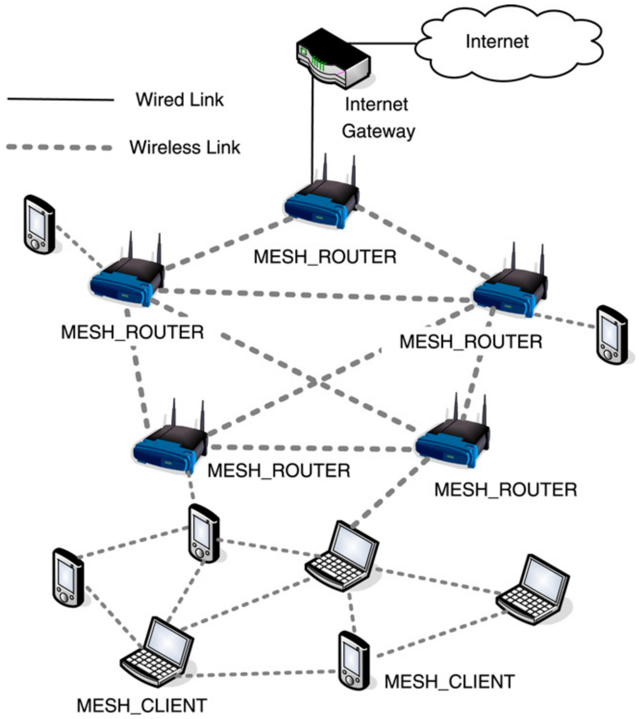

Fig. 1. Hybrid Wireless Mesh Network.

The remainder of the paper is organised as follows: Section 2 provides an overview of the most relevant related work. The multi-radio extensions to AODV are explained in Section 3. Simulation results and their analysis are discussed in Section 4. Section 5 concludes the paper.

2. Related work

2.1. Hyacinth

Hyacinth [7] is a static wireless mesh network that uses multiple radios and channels per node to improve network performance. It supports a fully distributed channel assignment algorithm, which can dynamically adapt to varying traffic loads. It uses a spanning-tree based routing algorithm [8] to load balance the network as well as to rectify route failures. The `MESH_ROUTER`s that have access to the wired network are considered as the root nodes of the spanning tree. The authors prove that even with complete knowledge of the network topology, the channel assignment problem is NP-hard. Hyacinth's channel assignment heuristic breaks a single-channel collision domain into multiple domains, one per distinct channel. The channel assignment algorithm operates in two phases: neighbour-interface binding and interface-channel assignment. In the first phase, each node separates its interfaces into upstream interfaces (UP-NICs) and downstream interfaces (DOWN-NICs). Each node has control over the channel assignment on its DOWN-NICs only. During the second phase, each node exchanges periodic channel usage status messages with its immediate neighbours. Using the per-channel total load information, a node can issue a change channel message to its downstream neighbours in order to switch to a less used channel. The advantage of this channel assignment scheme is that a fat-tree architecture is obtained in which links close to the root of the spanning tree are given higher bandwidth. The channel assignment is further integrated with the routing process. Each node that has a valid path to the root node advertises this information to its one-hop neighbours. This advertisement also contains the cost metric, which is made up of hop-count and uplink capacity. Each node that receives this advertisement makes a decision whether to join the advertising node based on the cost metric. If the node decides to join, it sends an acceptance message to the advertising node and a departing message to the parent node it was previously attached to. New nodes entering the network broadcast `HELLO` packets in order to initiate the joining process.

2.2. Single-radio multi-channel routing

The Multi-Channel Routing Protocol (MCRP) [9] is a routing protocol specifically designed for wireless nodes with a single-radio, but with the ability to switch channels with a delay of less than 80 μs [10]. The protocol assigns channels to data flows rather than to nodes. This implies that all nodes that are involved in a flow need to be on a common channel. The advantage of this mechanism is that once the route is established, nodes are not required to switch channels for the duration of a flow. The protocol considers all nodes in the network to be in one of the four states, i.e. free, locked, switching or hard-locked. The free nodes are nodes that are not supporting any flow at the moment. Locked nodes are nodes that are currently supporting a single flow. A switching node supports two or more flows on different channels. A hard-locked node is one that, due to certain constraints, cannot become a switching node. The MCRP protocol benefits from multiple channels without modifying the MAC layer. The routing scheme is similar to that of AODV. However, each `ROUTE_REQUEST` packet is broadcast on each channel in a round robin manner. Each receiving node rebroadcasts the `ROUTE_REQUEST` packet in a round robin manner. These nodes also create a reverse route to the source and maintain a channel number for the next hop with each reverse route table entry. To convey the channel information of the next hop, each `ROUTE_REQUEST` packet contains the operating channel[1] number of the hop sending the `ROUTE_REQUEST` packet. The `ROUTE_REQUEST` packet also contains a channel table and flow table to be propagated along with each `ROUTE_REQUEST` packet. The channel table contains a list of channels that are used on a single flow path. The flow tables maintain a list of simultaneous flows being carried out on a single channel. These tables are used by the destination to make a decision regarding the selection of the optimal route from multiple received `ROUTE_REQUEST` packets. The `ROUTE_REPLY` packet is unicast from the destination to the source via the optimal path. All nodes forwarding the `ROUTE_REPLY` packet change their operating channels to the channel selected by the destination.

2.3. Multi-radio link quality source routing

The Multi-Radio Link Quality Source Routing (MR-LQSR) protocol [5] has been developed for static community wireless networks. The protocol works in conjunction with the Mesh Connectivity Layer (MCL), which permits higher layer applications to connect to the wireless mesh network using Wi-Fi or WiMAX technology. The MCL implements an intermediary layer between the data link and network layer. It is implemented as a loadable driver, which acts as a virtual network adapter with the ability to multiplex between several physical adapters. The MCL emulates a virtual network adapter and routes packets using MR-LQSR, which is an optimised version of the DSR protocol. The MR-LQSR protocol assumes that the number of wireless interfaces per node is equal to the number of channels being used in the network. The protocol identifies all nodes in the wireless mesh network and assigns weights to all possible links. To do so, link information including channel assignment, bandwidth and loss rates are propagated to all nodes in the network. This propagation is combined with the delivery of DSR control packets. The Expected Transmission Time (ETT) metric on each link is computed using the Expected Transmission Count (ETX), bandwidth and packet size [11]. The ETT metric is then used to compute the Weighted Cumulative Expected Transmission Time (WCETT), which is a path metric that represents a trade off between high throughput and low delay paths. In contrast to native DSR which uses hop-count as the routing metric, MR-LQSR uses WCETT as its path metric.

2.4. AODV-ST

AODV-ST [6] is a hybrid routing protocol developed specifically for infrastructure mesh networks. The protocol uses the ETT routing metric to select routes. AODV-ST has been designed with the aim of providing Internet access to `MESH_CLIENT`s with the help of one or more gateways. AODV-ST uses a proactive strategy to discover routes between the `MESH_ROUTER`s and the gateways, and a reactive strategy to find routes between `MESH_ROUTER`s. In the proactive case, the gateways periodically broadcast special `ROUTE_REQUEST` packets to initiate the creation of spanning trees. All subsequent `ROUTE_REQUEST` packets with a better routing metric are used to update the existing reverse route to the gateway.

2.5. Mult-radio multi-channel routing

The Multi-Channel Routing (MCR) Protocol [12] has been developed for dynamic wireless mesh networks, where nodes have multiple wireless interfaces. The protocol makes use of an interface switching mechanism to assign interfaces to channels. Two types of interfaces are defined: Fixed and switchable. With fixed interfaces, K out of a total of M interfaces are assumed to be operating on K fixed channels. With switchable interfaces, the remaining $M–K$ interfaces are dynamically assigned to any of the remaining $M–K$ channels. The switching is carried out depending on the maximum number of data packets, which are queued for a single channel. Multiple queues are maintained for all switchable interfaces and each node maintains a neighbour table and channel usage list. The neighbour table contains information regarding the fixed channels that are being used by the node's neighbours. The channel usage list contains the nodes and their corresponding fixed channels. Each node periodically transmits a `HELLO` packet on all channels, containing the node's fixed channel number. Each node that receives a `HELLO` packet updates its neighbour table and a channel usage list. The information in this table and list is used to manage the channel and interface switching mechanism. The switching mechanism assists the MCR protocol in establishing routes over multiple channels. MCR uses a new routing metric

---

[1] The operating channel of a node is the default channel on which it is generally listening.

which is computed as a function of the channel diversity cost, interface switching cost and hop-count. Diversity cost is determined by the number of distinct channels used in a route. A route composed of a large number of distinct channels is considered as having a lower diversity cost. The switching cost represents the cost involved in switching interfaces from one channel to another. The route discovery mechanism of MCR is similar to that of the DSR protocol. However, each `ROUTE_REQUEST` packet also contains the channel number and switching cost. When the destination receives the `ROUTE_REQUEST` packet, it computes the channel diversity cost and the total switching cost (sum of all link switching costs). The selection of the path is based on minimising the sum of all these costs.

The related work indicates that two types of switching mechanisms can be used to improve the routing performance of a network, i.e. interface switching and channel switching. In the former mechanism network interfaces are switched in order to connect to other network nodes, while in the latter channels are switched on a single interface to connect to nodes operating on the same channel. Both mechanisms permit improved utilisation of the frequency spectrum. However, the accurate execution of these mechanisms in a mobile network entails the availability of a virtual switching protocol and incurs switching delays [13,5]. The switching protocol may be integrated with the MAC layer or may be implemented in a higher layer protocol. In either case, it requires precise synchronisation between the nodes negotiating an interface or channel switch.

In the following sections we discuss and evaluate the AODV-MR protocol which does not require exact synchronisation between nodes but is still able to make use of multiple radios per node by establishing channel diverse paths. AODV-MR lowers contention for the medium and achieves an overall improvement of network performance.

3. Multi-radio ad-hoc on-demand distance vector routing

The Multi-Radio Ad-hoc On-Demand Distance Vector Routing (AODV-MR) protocol is a multi-homing extension to the AODV protocol. AODV is inherently a distance vector routing protocol that has been optimised for ad-hoc wireless networks. It is an on-demand or reactive protocol since it sets up routes only when required. AODV borrows basic route establishment and maintenance mechanisms from the DSR protocol and hop-to-hop routing vectors from the Destination-Sequenced Distance-Vector (DSDV) routing protocol [14]. Multi-path support has also been added to AODV through a number of extensions [15,16], which permit discovery and establishment of loop-free and disjoint alternate paths. To avoid the problem of routing loops, AODV makes extensive use of sequence numbers in control packets.

When a source node intends to communicate with a destination node whose route is not known, the source broadcasts a `ROUTE_REQUEST` packet. Each `ROUTE_REQUEST` packet contains an ID, source and destination IP addresses, sequence numbers, hop-count and control flags. The ID field uniquely identifies the `ROUTE_REQUEST` packet and the sequence numbers indicate the freshness of control packets. The hop-count represents the path length between the source and destination. Each recipient of a `ROUTE_REQUEST` packet that has not seen a request with the same source IP and ID pair, or does not have a fresher (with larger sequence number) route to the destination, rebroadcasts the same packet after incrementing the hop-count. Intermediate nodes also create and preserve a reverse route to the source node for a certain interval of time.

When the `ROUTE_REQUEST` packet reaches the destination or any node that has a fresher route to the destination, a `ROUTE_REPLY` packet is generated and unicast back to the source of the `ROUTE_REQUEST` packet. Each `ROUTE_REPLY` packet also contains a destination sequence number, source and the destination IP address, a route lifetime parameter, hop-count and control flags. Each intermediate node that receives the `ROUTE_REPLY` packet increments the hop-count, establishes a forward route to the source of the packet, and transmits the packet on the reverse route.

To preserve connectivity information, each node executing AODV can use link-layer feedback or periodic `HELLO` packets to detect link breakages to nodes that it considers as its immediate neighbours. In case a link break is detected for a next hop of an active route, a `ROUTE_ERROR` packet is sent to its active neighbours that were using that particular route.

In our scenario, each node is equipped with multiple radios and each radio can operate on one of the multiple non-overlapping channels. When a route is required in AODV-MR, the `ROUTE_REQUEST` packet is broadcast simultaneously on all interfaces. Neighbour nodes, which share at least one common channel with the sender node, receive the packet. If the `ROUTE_REQUEST` packet is not a duplicate (determined by source IP and destination ID), a reverse route pointing towards the source node is created.

AODV-MR maintains an interface number in the routing table, as shown in Fig. 2. The interface number indicates the network interface via which a next hop node, for a particular path, can be reached. This information permits sending of the `ROUTE_REPLY` packet and future data packets on the correct interface.

The intermediate nodes, after updating their routing tables, broadcast the `ROUTE_REQUEST` packet on all interfaces except the one on which the `ROUTE_REQUEST` packet was initially received. If the number of nodes in the network is n and the number of interfaces per node is i, the maximum number of `ROUTE_REQUEST` packets propagated in the network is $i \cdot (n - 1)$.

The propagation of `ROUTE_REQUEST` packets continues until the IP TTL expires. If the `ROUTE_REQUEST` packet is received by the destination itself or any node with a fresh route to the destination, a `ROUTE_REPLY` packet is

Destination IP	Destination Seq. No.	Destination Valid Flag	Flags	Network Interface	Hop Count	Next Hop	Precursors List	Lifetime

Fig. 2. AODV-MR Routing Table Fields.

sent to the source node using a loop free path. When the `ROUTE_REPLY` packet is received by intermediate nodes, a forward route to the destination is established by creating the corresponding routing table entries, including the interface number.

We can model a multi-radio WMN with two graphs G_1 and G_2. $G_1 = (N, E)$ represents the set of all nodes N in the network and the set of edges E. An edge $e_{i,j}$ exists between nodes n_i and n_j $(e_{i,j} \in E)$ if the two nodes are within transmission range of each other.[2] In single-radio WMNs, where all nodes operate on a single common channel, being within transmission range of another node directly implies network connectivity. This is not necessarily the case for multi-radio WMNs, where it is possible that two nodes that are in close proximity do not have network connectivity due to the fact that none of their interfaces operates on a common channel. To make this distinction, we use L as the set of actual network links between nodes. The graph $G_2 = (N, L)$ represents the network connectivity graph of a WMN. A link $l_{i,j}$ between nodes n_i and n_j exists $(l_{i,j} \in L)$ if the following two conditions are met. First, n_i and n_j need to be within transmission range of each other, i.e. $e_{i,j} \in E$. Secondly, the two nodes need to share at least one common channel.

If C_i and C_j represent the sets of channels on which the interfaces of node n_i and n_j are operating on, this condition can be expressed as $C_i \cap C_j \neq \emptyset$. We make the assumption that no two interfaces on the same node are tuned to the same channel. In single-radio WMNs, where all nodes operate on a common single channel, E and L, and thus also G_1 and G_2, are identical.

4. Simulation results and analysis

4.1. Simulation environment

We have compared the performance of AODV-MR with AODV-ST[3] and the standard single-radio AODV in a hybrid wireless mesh network through extensive simulations in NS-2 [17]. We have used the Enhanced Network Simulator (TeNs) extensions[4] to provide multi-radio support [18].

A dense WMN covering an area of 1 square km is established using 25 static `MESH_ROUTER`s (51–75) arranged in a regular 5 × 5 grid, as shown in Fig. 3. The distance between adjacent `MESH_ROUTER`s is set to 176 m to ensure diagonal connectivity. The network further consists of 50 randomly placed mobile `MESH_CLIENT`s (1–50). Concurrent UDP flows are established between randomly selected source and destination `MESH_CLIENT` pairs.

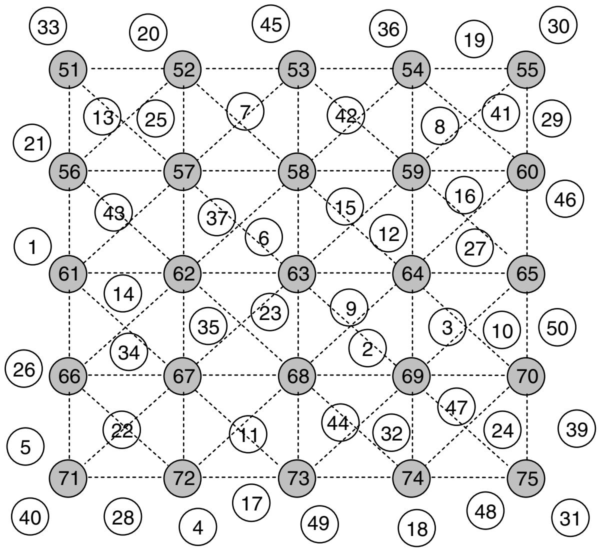

Fig. 3. Mesh Network Topology.

The following four tests were conducted to evaluate the performance of the AODV-MR protocol under varying mobility and traffic load conditions as well as different node configurations:

- Test 1: Varying `MESH_CLIENT` speeds
- Test 2: Varying traffic load
- Test 3: Varying packet size
- Test 4: Varying the number of radios per `MESH_ROUTER`

The performance metrics are obtained by ensemble averaging the results from over 50 individual simulation runs for each test [19]. Confidence intervals have been shown in all test results as vertical bars indicating the minimum, mean and maximum values. All three routing protocols i.e. AODV, AODV-ST and AODV-MR have been evaluated in Test 1, while AODV and AODV-MR have been assessed in Tests 2, 3 and 4 respectively. The maximum speed of the `MESH_CLIENT`s in Test 2, 3 and 4 is set to 1 m/s. A positive minimum speed of 0.1 m/s has been set during the tests as the random way-point mobility model fails to reach a steady state in terms of instantaneous average node speeds when the minimum speed is set to zero [20]. The packet size is set to 512 bytes in Tests 1, 2 and 4. The number of flows in Tests 1, 3 and 4 is set to 30. Except for Test 3, the transmission rate of each flow is

---

[2] We assume bi-directional links and a constant transmission range for all radios.

[3] For the sake of comparison, we have only implemented the reactive portion of AODV-ST in NS-2.

[4] http://www.cse.iitk.ac.in/users/braman/tens/

Table 1
Simulation Parameters

Examined protocols	AODV & AODV-MR
Simulation time	900 s
Simulation area	1000 × 1000 m
Propagation model	Two-ray ground reflection
Mobility model for `MESH_CLIENT`s	Random way-point
Maximum speed of `MESH_CLIENT`s	20 m/s
Transmission range	250 m
Number of flows	30
Traffic type	CBR (UDP)
Packet size	512 bytes
Packet rate	32 pkts/s
Transmission rate	128 kbps/flow
Number of `MESH_ROUTER`s	25
Number of 802.11b radios per `MESH_ROUTER`	6
Number of `MESH_CLIENT`s	50
Number 802.11b radios per `MESH_CLIENT`	1

set to 128 kbps, and the average path lengths between a source and destination node is four hops.

The number of radios per `MESH_ROUTER` in Tests 1, 2 and 3 is set to 6 for AODV-MR and AODV-ST, and is set to 1 for standard AODV. All `MESH_CLIENT`s operate a single radio only. Single-radio nodes have the radio tuned to the 802.11b channel number 1, while the multi-radio nodes operate the radios tuned to 802.11b channel numbers 1, 3, 5, 7, 9 and 11, respectively.[5] All channels are assigned to the radios before the tests are commenced and remain fixed for the duration of the tests. The parameters common to all four tests are listed in Table 1.

4.2. Assumptions

The following assumptions have been made:

- All radios are statically tuned to a channel.
- All `MESH_CLIENT`s and `MESH_ROUTER`s have a radio tuned to a common 802.11b channel.
- The remaining radios on the `MESH_ROUTER`s are tuned to orthogonal 802.11b channels.
- The transmission and reception ranges of the wireless transceivers are equal.
- All antennas are omni-directional.

4.3. Mobility model

We use the random way point mobility model for the `MESH_CLIENT`s in our simulation. `MESH_CLIENT`s first wait for the pause interval of 10 seconds, then move to a randomly chosen position with a velocity chosen randomly between 0.1 m/s and the maximum speed, wait there for 10 s, and then move on to the next random position. A maximum speed of 0 m/s corresponds to a completely static network.

4.4. Communication model

The IEEE 802.11 Distributed Coordination Function (DCF) [21] is used at the MAC layer. All packets are transmitted using the un-slotted Carrier Sense Multiple Access protocol with Collision Avoidance (CSMA/CA) as used by Lucent Technologies WaveLAN-I[6]. In CSMA/CA each broadcasting node waits for a vacant channel by sensing the medium. If the channel is vacant, it makes the transmission. In case of a collision, the colliding stations wait using the Ethernet binary exponential back off algorithm [22]. Virtual Carrier Sensing (RTS/CTS) is disabled during the simulations.

4.5. Performance metrics

In our simulations we consider the following performance metrics:

1. **Packet Delivery Ratio**: The ratio between the number of data packets successfully received by destination nodes and the total number of data packets sent by source nodes.
2. **Routing Overhead**: The ratio of the total number of control packets generated to the total number of data packets that are successfully received.
3. **Average Latency**: The mean time (in seconds) taken by packets to reach their respective destinations.
4. **Path Optimality**: The ratio between the length (number of hops) of the shortest possible path and the actual path taken by the packets.

4.6. Test 1: varying the `MESH_CLIENT` Speeds

In Test 1, we have varied the maximum speed of the `MESH_CLIENT`s from 0 to 20 m/s. The results of Test 1 are shown in Fig. 4. The results indicate that standard AODV loses more than 75% of the transmitted packets under varying node speeds. This loss is primarily attributable to the large traffic volume traversing nodes with single radios. Nodes which have to share a single radio among multiple flows observe congestion at the physical layer and consequently experience frequent packet drops. In contrast with AODV, the `MESH_ROUTER`s used by AODV-ST and AODV-MR have six radios each, which facilitate multiple flows to be distributed over distinct phys-

---

[5] Although, 802.11b can only support three orthogonal channels, we have configured the NS-2 802.11b physical layer to consider all channels to be orthogonal. This allows us to simulate the behaviour of radios that support a high number of orthogonal channels such as 802.11a.

[6] http://ftp4.de.freesbie.org/pub/misc/wavelan/docs/LEGACY/BULLETIN/SALES/cb-prox.pdf

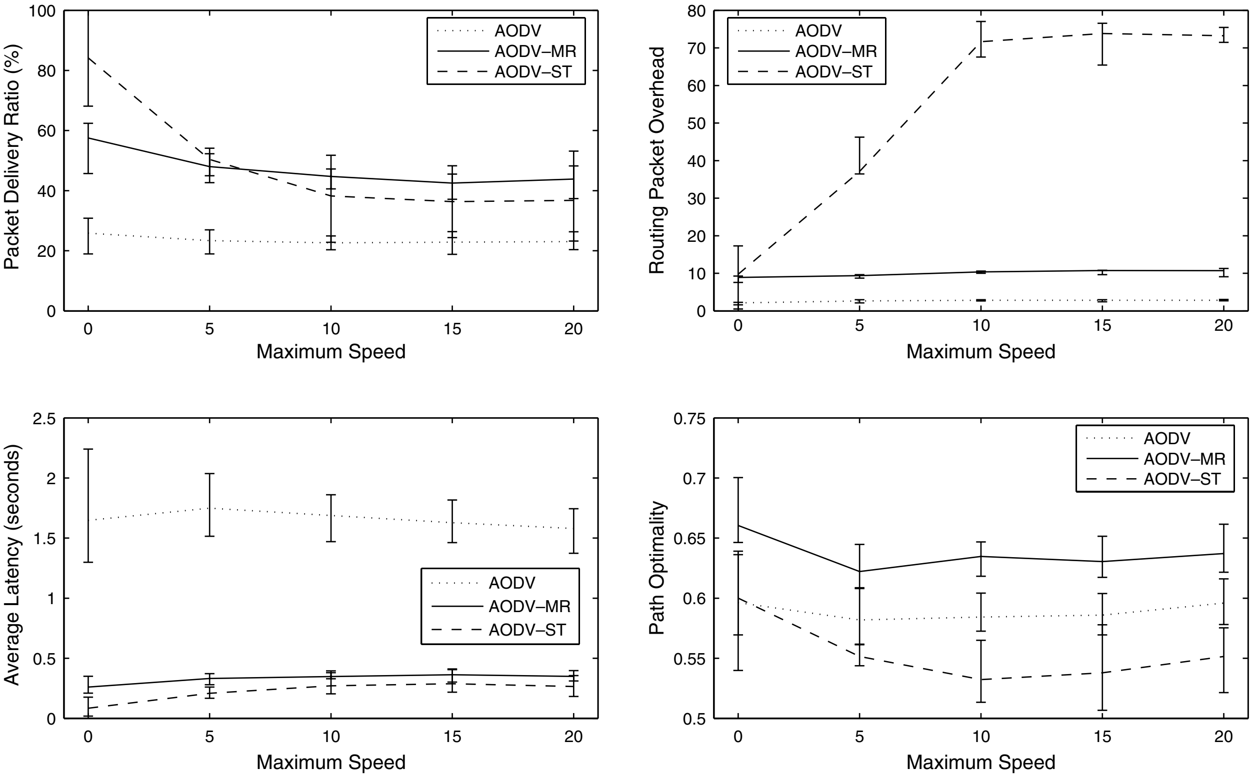

Fig. 4. Test 1 Results: varying the `MESH_CLIENT` speeds.

ical layers. This in turn helps in sustaining multiple ongoing flows, resulting in a lower packet losses.

The packet losses directly affect the packet delivery rate (PDR) of the network. Standard AODV achieves a PDR of around 25% at zero `MESH_CLIENT` mobility. The PDR then gradually decreases to approximately 23% when the `MESH_CLIENT` speeds reach 20 m/s. AODV-MR achieves a PDR of 55% at zero `MESH_CLIENT` mobility, which gradually decreases to 45% when the `MESH_CLIENTs` reach a speed of 20 m/s. At zero `MESH_CLIENT` speeds, AODV-ST achieves a PDR of 85%. However, as soon as nodes are mobile the PDR drops to less than 40%.

The degradation of the PDR in AODV-ST is primarily attributable to the mechanism through which the ETX is computed. When the network topology is fixed, the ETX can be accurately determined using periodic `HELLO` packets. However, at higher node speeds the ETX computation fails to converge due to relatively large time window during which the successful receipt of `HELLO` packets is measured [23]. As a result, stale routes are often selected, causing route breaks. Thus additional route discoveries need to be performed, which increase the routing packet overhead.

The routing packet overhead of AODV-MR is considerably higher than that of the standard AODV. This occurs due to the `ROUTE REQUEST` packets being transmitted on multiple radios during the route discovery process. This essentially increases the routing overhead of the standard AODV by a factor equal to the number of radios in each `MESH_ROUTER`.

The most notable improvement of AODV-MR over standard AODV is in terms of the average packet latency. As mentioned earlier, when multiple flows are directed through a single radio, there is extensive contention for the physical medium. This contention, which occurs at the link layer, causes data packets to be delayed at each hop. In case of AODV-MR the `MESH_ROUTERs` offer multiple links between adjacent `MESH_ROUTERs`. These links allow distributing the load of multiple flows over multiple links.The overall contention on each link between two consecutive `MESH_ROUTERs` is, therefore, considerably reduced. This permits packets to be sent from one `MESH_ROUTER` to another with significantly lower delay. The routes established using AODV-ST have even lower latency than that of AODV-MR. AODV-ST continuously monitors the link qualities, which allows selection of low latency links during route discoveries.

With single radio nodes, there is a high probability that a `ROUTE REQUEST` packet gets lost due to collisions. Therefore, the shortest path may not be discovered in such route discoveries. However, when multiple radios are used by a `MESH_ROUTER`, the probability of all `ROUTE REQUEST` packets getting lost is negligible. The flooding of `ROUTE REQUEST` packets through multiple radios increases the path optimality of the network. AODV-MR shows high path optimality over varying network speeds indicating selection of routes with a smaller number of hops. On the other hand, the routes selected by AODV-ST are not optimal in terms of number of hops and, hence, have lower path optimality.

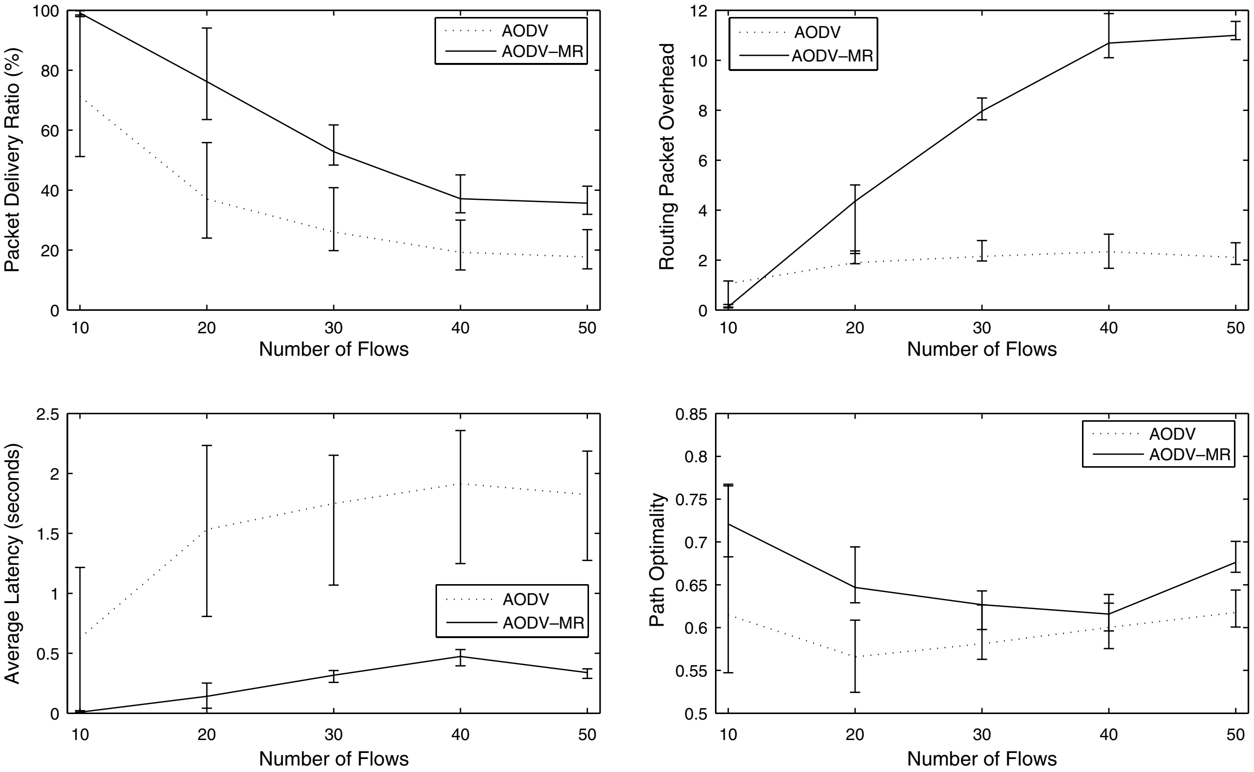

Fig. 5. Test 2 results: varying the traffic load.

4.7. Test 2: varying the traffic load

In Test 2, we varied the number of simultaneous 128 kbps flows between 10 and 50. The results of Test 2, depicted in Fig. 5, indicate that both protocols suffer from packet losses with the increase in traffic load. However, AODV-MR shows a considerably lower packet loss up to 30 simultaneous flows compared to the standard AODV protocol. Beyond this point the aggregate application layer data rate is more than 5 Mbps (40 × 128 kbps). This essentially saturates the physical layer of a single 802.11b channel, which has an effective capacity of no more than 5 Mbps [24].

The packet delivery ratio (PDR) for AODV-MR remains at almost 100% when the number of flows is set to 10. The PDR then drops to almost 35% when the number of flows reaches 50. For a similar increase in traffic load, the PDR of AODV degrades from 70% to less than 20%.

The routing overhead metric is computed as a ratio between the number of transmitted routing packets and the number of successfully received data packets. AODV-MR has a higher PDR and achieves a lower routing overhead when the number of flows is less than 10. However, when the traffic load is increased, the number of `ROUTE REQUEST` packets that are generated during route discoveries increases significantly.

The latency of the packets using AODV-MR remains below 50 ms, even when the traffic load is increased to 50 simultaneous flows. AODV-MR forms channel diverse paths with the help of multiple radios tuned to orthogonal channels. These paths help in minimising the latency, which generally increases when there is higher contention for the physical wireless medium. Standard AODV offers no option for channel diverse paths and, hence, exhibits higher latency with increasing traffic loads.

AODV-MR shows improved path optimality over the standard AODV protocol due to its effective route discovery using multiple radios. However, both protocols depict minimum deviation from the shortest possible paths[7] with increasing traffic loads.

4.8. Test 3: varying the packet size

In Test 3, we varied the packet size from 64 to 1024 bytes, while maintaining the number of simultaneous flows at 30. The results, shown in Fig. 6, indicate that AODV-MR is able to minimise packet losses compared to standard AODV. The packet delivery ratio of AODV-MR shows an improvement of more than 60% over AODV for a packet size of 64 or 128 bytes. This also results in a lower routing packet overhead, which is computed as a ratio of generated control packets per received data packet. The latency with AODV-MR remains considerably lower

---

[7] The shortest possible path is determined by an omniscient entity present in the NS-2 simulator known as the General Operations Director.

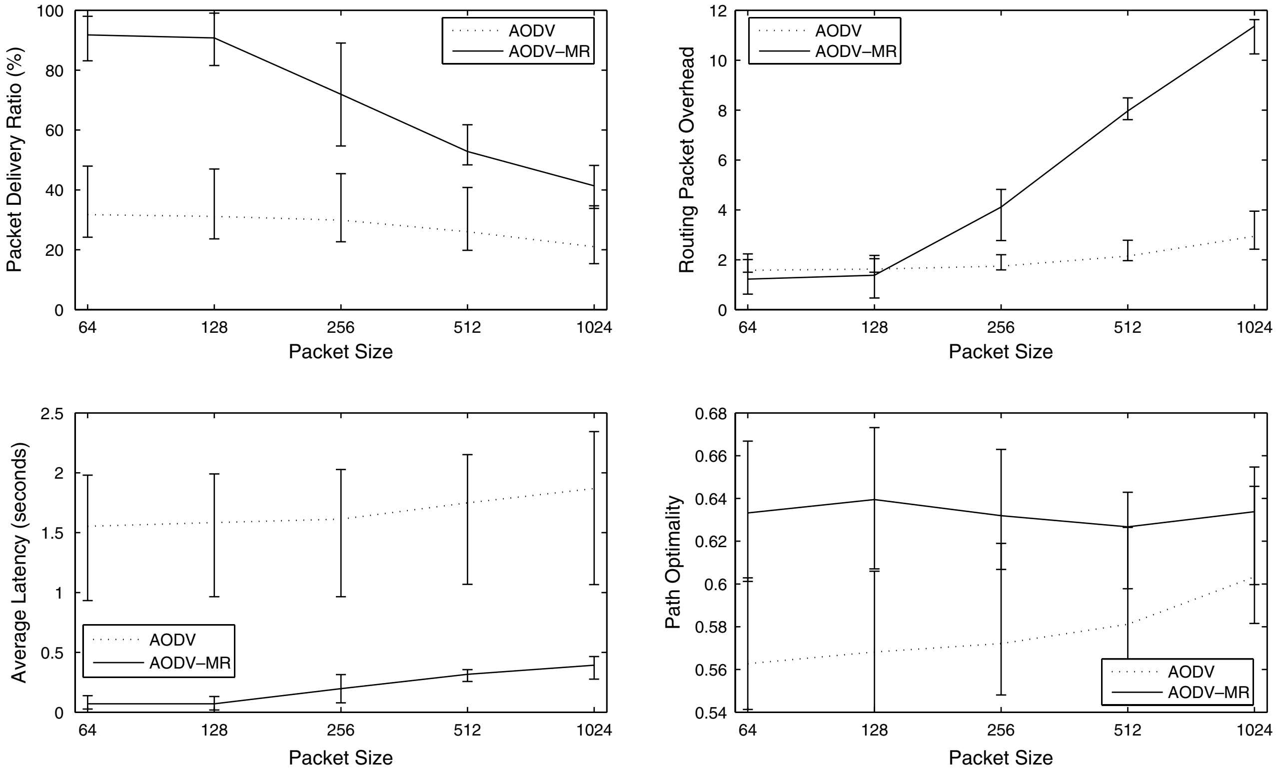

Fig. 6. Test 3 results: varying the packet size.

than that of standard AODV. The higher path optimality of AODV-MR suggests usage of shorter paths compared to those used by AODV. This can occur in scenarios where the routes are being established through a minimum set of `MESH_ROUTER`s rather than a set of `MESH_CLIENT`s.

4.9. Test 4: varying the number of radios on each `MESH_ROUTER`

In Test 4, we have varied the number of interfaces in each `MESH_ROUTER` from 1 to 9, with an increment of 2 interfaces. The results of Test 4, shown in Fig. 7, indicate that the packet losses gradually decrease with the increase in the number of radios. This can be attributed to the increased probability of selecting different interfaces (channels) for the forward and reverse routes.

AODV-MR exhibits improved packet delivery with the increase in the number of radios. The additional radios help AODV-MR to select, with a high probability, interference and contention free optimal links during the creation of routes. However, increasing the number of `MESH_ROUTER` interfaces to more than three does not show any significant further improvements. This is due to the fact that three interfaces operating on orthogonal channels are sufficient to provide the required capacity and channel diversity for the network and traffic pattern considered in our simulation. However, the ideal number of `MESH_ROUTER` interfaces will vary for different types of networks with different size, density and traffic load.

The routing overhead of AODV-MR increases with the increase in number of radios. This essentially occurs due to the increased number of `ROUTE REQUEST` packets being propagated by the intermediate `MESH_ROUTER`s.

The latency in the network is considerably reduced with an increase in the number of radios. Multiple radios on each `MESH_ROUTER` allow the creation of channel diverse routes, which in turn reduce the average packet latency. The path optimality of the network using AODV-MR shows gradual improvement with the increase in number of radios per `MESH_ROUTER` due to reasons discussed earlier.

5. Conclusions

Wireless mesh networks have a great potential for a wide range of applications owing to their self-configuring, self-optimising and self-healing capabilities, coupled with their low cost of equipment and deployment. One of the key shortcomings of wireless mesh networks is their limited capacity and scalability, which is typical of multi-hop wireless networks. It has been shown previously that using multiple radios per node, operating on orthogonal channels, can greatly increase the capacity of wireless mesh networks. In this paper, we have presented a performance evaluation of a multi-radio version of the AODV routing protocol. Multi-radio AODV makes efficient use of the multiple interfaces and increased frequency spectrum to minimise interference and contention in the network. Our simulation

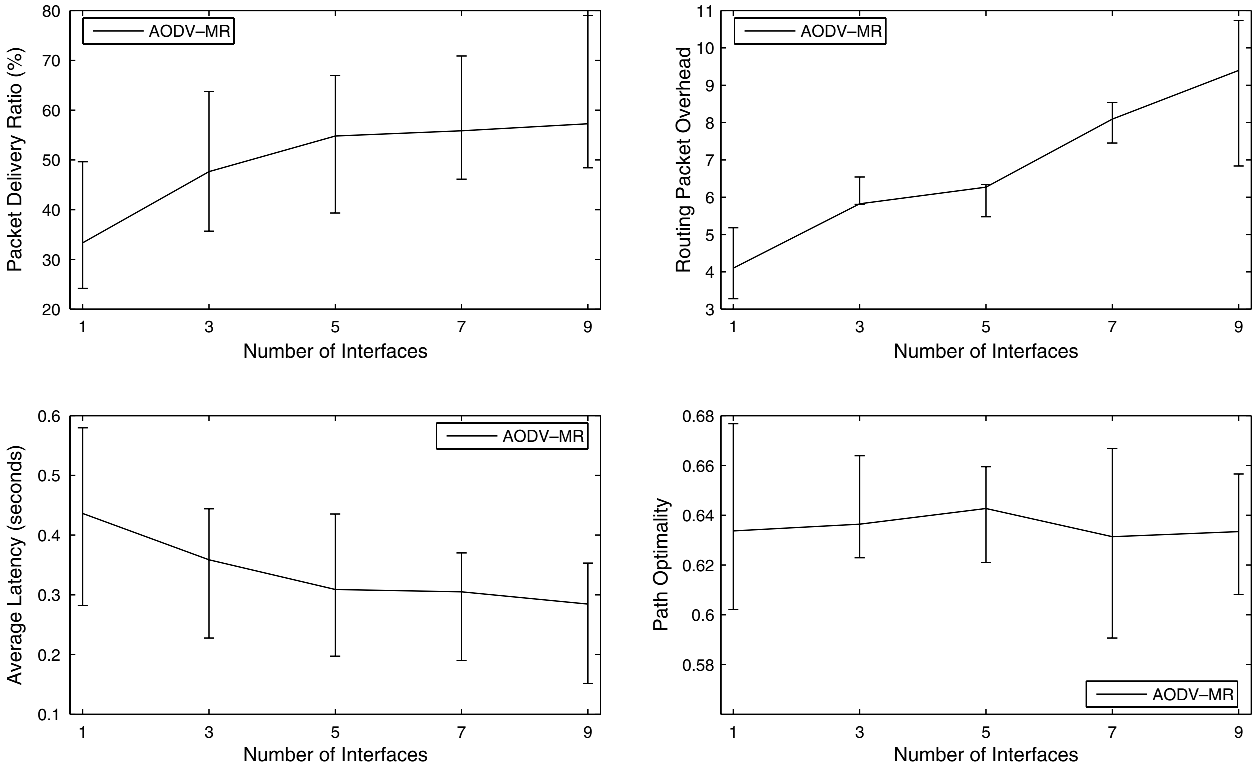

Fig. 7. Test 4 results: varying the number of radios per `MESH_ROUTER`.

results are very encouraging and show a significant improvement in terms of packet delivery ratio and latency over the single-radio AODV, especially under high load conditions. We believe there is a great potential to further improve the performance of multi-radio AODV, for example by using routing metrics that not only consider path length, but also the quality of the individual links [25]. Further improvements can be gained by making more optimal use of channel diversity of the end-to-end path, thereby minimising co-channel interference. These and similar ideas are currently being investigated in our ongoing work.

Acknowledgements

National ICT Australia is funded by the Australian Government's Department of Communications, Information Technology and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Research Centre of Excellence programs and the Queensland Government.

References

[1] I.F. Akyildiz, X. Wang, A survey on wireless mesh networks, IEEE Communications Magazine 43 (9) (2005) S23–S30.

[2] A.A. Pirzada, C. McDonald, Establishing trust in pure ad-hoc networks, Proceedings of the 27th Australasian Computer Science Conference (ACSC), vol. 26, Australian Computer Society, 2004, pp. 47–54.

[3] D.B. Johnson, D.A. Maltz, Y. Hu. The dynamic source routing protocol for mobile ad hoc networks (DSR). IETF MANET, Internet Draft (2003).

[4] C. Perkins, E.M. Royer, S. Das, Ad hoc on-demand distance vector (AODV) routing, IETF RFC 3561 (2003).

[5] R. Draves, J. Padhye, B. Zill, Routing in multi-radio, multi-hop wireless mesh networks, in: Proceedings of the 10th Annual International Conference on Mobile Computing and Networking, ACM Press, 2004, pp. 114–128.

[6] K. Ramachandran, M. Buddhikot, G. Chandranmenon, S. Miller, E. Belding-Royer, K. Almeroth, On the design and implementation of infrastructure mesh networks, in: Proceedings of the IEEE Workshop on Wireless Mesh Networks (WiMesh), IEEE Press, 2005, pp. 4–15.

[7] A. Raniwala, T.C. Chiueh, Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network, Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), vol. 3, IEEE Press, 2005, pp. 2223–2234.

[8] IEEE, Standard for local and metropolitan area networks: media access control (MAC) Bridges IEEE P802.1D/D41 (2003).

[9] J. So, N.H. Vaidya, A routing protocol for utilizing multiple channels in multi-hop wireless networks with a single transceiver, Technical Report, Dept. of Computer Science and Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 2004.

[10] P. Kyasanur, N.H. Vaidya, Routing and interface assignment in multi-channel multi-interface wireless networks, Proceedings of the IEEE Wireless Communications and Networking Conference, vol. 4, IEEE Press, 2005, pp. 2051–2056.

[11] R. Draves, J. Padhye, B. Zill, Comparison of routing metrics for static multi-hop wireless networks, in: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), ACM Press, 2004, pp. 133–144.

[12] P. Kyasanur, N.H. Vaidya, Routing and link-layer protocols for multi-channel multi-interface ad hoc wireless networks, SIGMOBILE Mobile Computing and Communications Review 10 (1) (2006) 31–43.

[13] R. Chandra, P. Bahl, MultiNet: connecting to multiple IEEE 802.11 networks using a single wireless card, Proceedings of the Twenty-third

Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), vol. 2, IEEE Press, pp. 882–893.

[14] C.E. Perkins, P. Bhagwat, Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers, in: Proceedings of the SIGCOMM Conference on Communications, Architectures, Protocols and Applications, ACM Press, 1994, pp. 234–244.

[15] S.J. Lee, M. Gerla, AODV-BR: backup routing in ad hoc networks, in: Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), IEEE Press, 2000, pp. 1311–1316.

[16] M.K. Marina, S.R. Das, On-demand multi path distance vector routing in ad hoc networks, in: Proceedings of the Ninth International Conference on Network Protocols (ICNP), IEEE Press, 2001, pp. 14–23.

[17] NS. The Network Simulator. http://www.isi.edu/nsnam/ns/, 1989.

[18] B. Raman, C. Chebrolu, Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks, in: Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom), ACM Press, 2005, pp. 156–169.

[19] W.H. Yuen, R.D. Yates, Inter-relationships of performance metrics and system parameters in mobile ad hoc networks, Proceedings of the IEEE Military Communications Conference (MILCOM), vol. 1, IEEE Press, 2002, pp. 519–524.

[20] J. Yoon, M. Liu, B. Noble, Random waypoint considered harmful, Proceedings of the Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), vol. 2, IEEE Communications Society, 2003, pp. 1312–1321.

[21] IEEE. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications 802.11. 1997.

[22] A.S. Tanenbaum, Computer Networks, fourth ed., Prentice Hall, 2002.

[23] D.S.J. De Couto, High-throughput Routing for Multi-hop Wireless Networks, Ph.D. thesis, Massachusetts Institute of Technology, 2004.

[24] J.K. Anderson, N. Youell, A closer look at WLAN throughput and performance, Bechtel Telecommunications Technical Journal 1 (1) (2002) 86–94.

[25] Y. Yang, J. Wang, R. Kravets, Designing routing metrics for mesh networks, in: Proceedings of the IEEE Workshop on Wireless Mesh Networks (WiMesh), IEEE Press, 2005.