

MRP: Wireless Mesh Networks Routing Protocol

Jangeun Jun, Mihail L. Sichitiu

Department of Electrical and Computer Engineering

North Carolina State University

Raleigh, NC 27695-7911

Email: {jjun, mlsichit}@ncsu.edu

Abstract

Wireless Mesh Networks (WMNs) are one of the few commonly implemented types of mobile ad-hoc networks (MANETs); several companies offer WMNs for broadband Internet access and for extending the coverage of wireless local area networks. Several particularities differentiate WMNs from MANETs. First, in WMNs, most of the traffic originates or terminates at the gateways (nodes connected to the wired infrastructure/Internet). Second, in most applications, WMN nodes tend to be neatly differentiated as either stationary nodes (providing connectivity and coverage) or mobile nodes (utilizing the coverage afforded by the stationary nodes). While general MANET routing protocols can be used in WMNs, it is expected that a protocol that takes the particularities of WMNs into account will outperform the general protocol. In this paper, we propose such a routing protocol and evaluate its performance via simulations. We show that, for WMNs, the proposed routing protocol outperforms general purpose MANET protocols in terms of routing overhead, packet delivery ratio, throughput and delay.

I. INTRODUCTION

Wireless Mesh Networks (WMNs) [1]–[3] are a relatively new wireless multihop technology that has much in common with the mobile ad hoc networks (MANETs). In its general form (Fig. 1), a WMN is a set of wireless nodes that can communicate with each other, forwarding each other's packets. Like in MANETs, each node is both a host and a wireless router. Clients can connect to the WMN routers using common networking interfaces (e.g., Ethernet, 802.11, Bluetooth) or, in some cases, a PCI or a PCMCIA bus (i.e., the WMN router is embedded in a network card). In most proposed applications, the WMN provides connectivity to an infrastructure network, typically connected to the Internet. We will call the nodes providing Internet connectivity *gateways*.

There are four types of links presented in Fig. 1. Except for the intra-mesh links (that have to be wireless), all other links can be either wireless or wired. The same or different technologies can be used




Fig. 1. A wireless mesh network connecting several stationary and mobile clients to the Internet.

for the four link types. The link choice usually represents a tradeoff between the cost/complexity and the performance of the WMN. Standardization efforts are initiated by IEEE (802.11s, 802.15.5, 802.16a).

Several companies [4]–[18] developed WMN products for different applications. One of the most popular WMN applications is providing broadband Internet access [9], [10], [12], [14], [15], [17], [18]. In this scenario, WMN routers are installed on the roofs of the clients and/or light poles in the coverage area of the WMN. Mobile clients may roam while being handed over from one wireless router to another (some products may be able to handle highway speeds [10]). The main advantage of WMNs in comparison to traditional broadband Internet access technologies (cable-modem and xDSL) is the dramatically reduced initial investment and deployment time. The main advantage in comparison to fixed wireless metropolitan area networks (WMANs) (e.g., IEEE 802.16) is the market coverage [19] (especially in areas with significant obstructions - trees, high-rise buildings) and reliability (multiple available routes can avoid failed nodes and poor links). Furthermore, some implementations allow for mobile user access, while the current IEEE WMAN standard only allows stationary users (although work for mobility extensions is underway).

A second major class of WMN products targets the booming wireless local area network (WLAN) market (mainly IEEE 802.11) [4], [6]–[8], [12], [16]. An important drawback of WLAN technology in multi-access point (AP) deployments is the requirement to separately provide wired network connectivity to each AP, offsetting some of the advantages of a wireless network. WMNs in this category solve the problem by placing the APs in range of each other and allowing them to forward each other's packets to and from a common gateway (another AP connected to the distribution system). The main drawback of these deployments is the reduced bandwidth available to the users (this can become a major problem in scenarios with many active users [20]). Some companies are counteracting the reduction in bandwidth by

using multiple radios on different channels and directional antennas [4], [8].

Other companies found niche applications for WMNs. Firetide [5] uses WMNs to provide network connectivity. Basically, the WMN cloud behaves like one big Ethernet switch that can be used to connect all devices plugged in any of the wireless routers (including, for example, 802.11 APs).

Routing is a fundamental characteristic of WMNs. The routing protocol's strengths and weaknesses are reflected directly in the WMN's characteristics. Several advantages of WMNs over competing technologies are directly enabled by the routing protocol:

- **Reliability:** The routing protocol should be able to reroute fast around failed nodes and broken links; upon the failure of a gateway, it should be able to redistribute the orphaned clients among neighboring gateways. For this property, fast reconfiguration and support of multiple gateways is essential.
- **Mobile user connectivity:** To ensure seamless mobile user connectivity, the routing protocol should enable fast hand-offs.
- **Scalability/Efficiency:** If the routing protocol has a high overhead, it will be impossible to scale the WMN to a large number of nodes.
- **QoS:** In addition to support from the medium access control (MAC) layer and the forwarding engine, selecting the "best" routes for different traffic classes is an essential ingredient for QoS support.

Taxonomically, WMNs are a particular type of mobile ad hoc network (MANET) [1], [21]. WMNs share the same multihop characteristics and mobility-related issues as MANETs. However, there are also significant differences between WMNs and general MANETs:

- **Gateways:** Most WMNs are designed to provide connectivity to a distribution system (usually connected to the Internet). Therefore, they have specialized nodes (the gateways) that provide connectivity to the distribution system.
- **Traffic pattern:** In WMNs, most of the traffic is expected to flow between the clients and the Internet (via the gateways). In general MANETs, the common assumption is that any node is equally likely to be the source or the destination of a traffic flow.
- **Mobility:** In most WMNs, nodes belong to two distinct categories: either stationary (e.g., on lamp poles, rooftops, etc.) or mobile, capable of roaming in the coverage area provided by the stationary nodes. In MANETs, it is often assumed that all nodes have homogeneous mobility characteristics.

In this paper, we propose a new mesh routing protocol (MRP), specifically geared toward WMNs. There have been many routing algorithms proposed for MANETs [21]–[35]. However, the characteristics and requirements of WMNs are considerably different than those of general MANETs such that a new routing protocol can significantly outperform the general MANET routing protocols. The situation is similar to the case of wireless sensor networks, where differences from general MANETs prompted the development

of specialized MAC and routing protocols [36].

MANET protocols rely on variations of a flooding mechanism for route discovery and recovery. Some of them attempt to reduce the flooding overhead by utilizing location information [32]–[35], limiting the flooding diameter [27], [37]–[39], or electing dedicated relay nodes [40]. However, none of them completely eliminates flooding. In contrast, our route discovery and recovery schemes do not require flooding. The main difference between the proposed and existing routing protocols is the absence of flooding.

The proposed routing protocol only maintains routing trees between clients and the gateways of the WMN. This mirrors the flow of data in the network and eliminates the overhead associated with maintaining direct routes between the clients. Furthermore, we show that, by preferring stable routes, the backbone of stationary nodes present in WMNs can be effectively used to improve the performance of the routing protocol. Finally, we introduce a novel link failure detection scheme that takes into account the inherently unreliable nature of wireless links.

The remainder of the paper is organized as follows. Section II provides a brief overview of the related work. In Section III, two versions of the proposed protocol are described, compared, contrasted, and merged into a third hybrid version. The performance of the three versions of the proposed protocol is evaluated via simulations in Section IV. Section V concludes the paper.

II. RELATED WORK

There are hundreds of proposed routing protocols. Many of them have been standardized by IETF and have been in use for many years. Some of those protocols have proven themselves in the Internet and are expected to continue to perform well for many years to come [41]–[45].

In the ad-hoc networking arena, several classes of routing protocols have been proposed and carefully analyzed [22], [23], [37]. The first class of proposed protocols was derived from existing “table-driven” protocols in the Internet (e.g., [25]) optimizing many of the aspects that reduced the efficiency of the existing Internet routing protocols in MANET environments.

In contrast to table-driven routing protocols, on-demand routing protocols (e.g., AODV [28], [46] and DSR [29]) were designed specifically for ad hoc networks with frequent disconnections (due to topology changes). They often outperform the table-driven routing protocols in scenarios with large networks with relatively few active connections.

In a class of their own, the geographical routing algorithms (e.g., LAR [32], DREAM [33] and ZRP [47]) take advantage of node location information (that can be relatively inexpensive to distribute [34]) to reduce routing overhead and improve the performance of the protocol. Geocasting [48], [49] takes a

rather unique view to addressing, as it considers a valid destination any node within a given geographical region. Several specialized applications can greatly benefit from geocasting.

It is well known [37] that current routing protocols scale poorly with the number of nodes, number of flows and increase in mobility. Several protocols were designed for scalability [33], [38], [40], [50]–[52] and show significantly better performance in large ad hoc networks.

The WMN companies are using a variety of routing protocols to satisfy their needs. Some are proprietary and held secret (e.g., [10]), while others use well-known ad hoc routing protocols (e.g., Firetide [5] uses Topology Broadcast based on Reverse-Path Forwarding (TBRPF) [50]). Other companies rely on the IEEE 802.1 spanning tree protocol for routing at layer 2 (e.g., MeshDynamics [8]).

It has long been recognized that many ad hoc networks would benefit from a connection to a fixed infrastructure, and several solutions have been developed to support this type of connectivity. In particular, the SURAN [53] and WINGS [54] projects solved many of the problems associated with these networks (MAC layer, channel access, integration with existing Internet protocols, mobility management, transport layer efficiency, etc.).

In [55], the authors propose extensions to existing routing protocols (AODV [28], DSR [29] and SOAR [56]) such that they optimize access to a set of nodes called netmarks (similar to what we call gateways). They evaluate the performance of the extended SOAR routing protocol (a link state routing protocol) and show via simulations that it outperforms both DSR and AODV. We do not propose to extend a link state protocol, but rather a new design specifically optimized for WMNs. Furthermore, our proposed protocol takes advantage of the fixed nodes in WMNs (by selecting stable routes).

In [57], the authors address routing issues in a hierarchical, two-level MANET. They propose a hybrid scheme using existing protocols in both hierarchies (DSDV [25] at the lower level, i.e., in the clusters connected to the gateways, and AODV in the backbone). They also propose an extension of AODV: Hybrid-AODV (H-AODV). Both approaches are shown to outperform classical AODV in terms of packet delivery ratio and delay for large networks. Our work is different in that it does not consider mobile backbones, and it does not use or extend existing protocols.

The LUNAR [58] system provides a simple routing protocol for small MANETs (limited to three hops). The protocol is predictable as well as easy to implement and analyze. It features auto-configuration and gateway routing, thus, making it a complete set of protocols. Its goals and functionality are, however, very different from those of the protocol proposed in this paper.

The K-hop routing protocol (KRP) [39], [59] is designed to provide service in an ad hoc network with nodes at most K hops away from a gateway. An extension of AODV (the flooding mechanism in AODV is limited to K -hops) is used to discover routes to the gateways. The gateways themselves publish

and update a bulletin board of reachable nodes ensuring network-wide dissemination of information. The approach allows for multiple routes to multiple gateways. Our approach is significantly different from KRP in that we use neither flooding nor a bulletin for route discovery and maintenance.

Landmark Routing (LANMAR) [51], [52], is an extension of the Fisheye State Routing (FSR) [38] that takes advantage of logical groups that tend to move together (e.g., soldiers in a platoon). Our approach is significantly different, as it is not a link state approach, and it does not take advantage of group mobility (unrealistic for general WMNs).

In [60], a gateway discovery scheme was proposed as a part of an architecture connecting generic MANETs and IPv6 networks. The scheme in [60] relies on flooding for route discovery, while the routing protocol proposed in this paper does not. Furthermore, the return paths (i.e., from the gateway to nodes in the WMN) are not established, a MANET routing protocol being assumed for this purpose.

ROMER [61] is a credit-based, opportunistic forwarding mechanism capable of using multiple routes at short time-scales and building long-term reliable routes at a cost of some redundancy. It is shown that for WMNs with significant variability in link quality, ROMER can significantly improve the packet delivery ratio of the network. Similarly, results presented in [62] demonstrate that multi-path routing can improve the reliability of WMNs even when the paths used are not disjoint.

AODV-ST [63] is a hybrid routing protocol designed for WMNs. It uses proactive tree formations by using a beaconing protocol initiated at the gateways to form spanning trees rooted at the gateway. For intra-mesh routing, it uses AODV to discover routes that bypass the gateways. Similar to our work, the main assumption is that the common-case traffic is to and from the gateway, and, thus, uses the proactively constructed trees. In contrast to our work, the common-case protocol (that forms the routing trees) is proactive, periodically flooding the network (AODV, used for intra-mesh routing, also floods RREQ messages). Our protocol does not use flooding.

Several articles considered routing-related issues for WMNs; in [64], [65], routing is considered in conjunction with channel assignment for multichannel WMNs based on commodity IEEE 802.11 compatible hardware. Several articles [65]–[70] show that non-conventional metrics may result in an (significantly) increased capacity in WMNs. In this paper we do not propose any new *routing metric*. Instead, we provide the hooks to allow the use of any of these (and future) routing metrics.

In MANETs, existing routing protocols employ flooding to discover and maintain routes between arbitrary pairs of nodes. In WMNs, most of the traffic is assumed to flow to and from the gateways connected to the Internet. With this assumption, we designed a protocol that does not use flooding either for route discovery or for recovery. Furthermore, we introduce a novel link failure detection procedure that is shown to effectively resolve the problem of misidentifying temporary link errors as permanent link

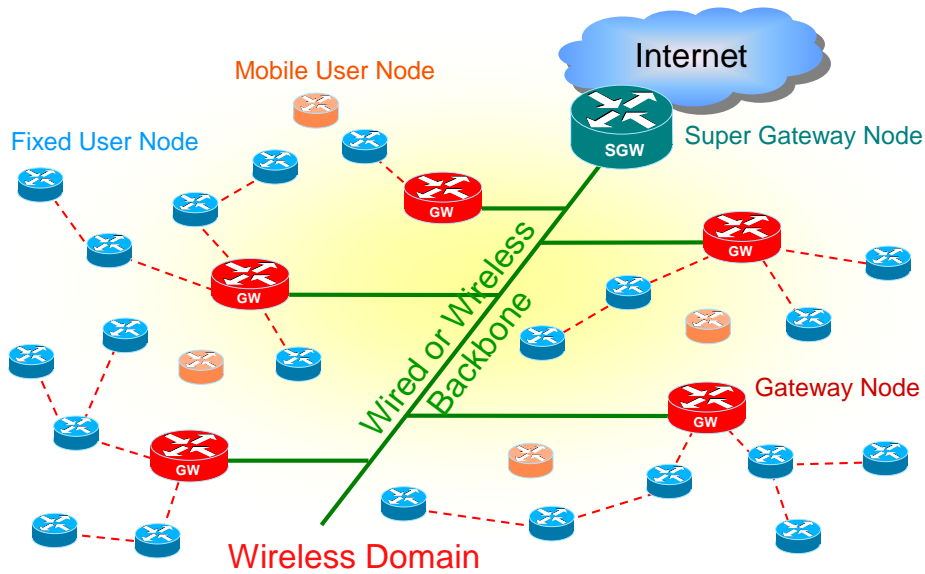


Fig. 2. General WMN network topology.

breaks.

III. PROTOCOL DESCRIPTION

In the design of the protocol, we decided to optimize the common case. In WMNs, the most common traffic flows are to and from the Internet (*downloads from the Internet* are by far the most common case, but TCP acknowledgments form streams in the opposite direction). Thus, we decided that any node in a WMN will only know how to reach one gateway and is, in general, reachable only from a gateway. Any small amount of client-to-client traffic can be routed through the common parent of the clients (potentially the gateway). In essence, the routes to and from the gateway form a tree rooted at the gateway.

In this version of the protocol, every client chooses a single gateway to connect to the Internet. If the node moves or if this gateway fails, the node may choose a different gateway. We, thus, assumed a network topology similar to the one shown in Fig. 2. The gateways are connected using wired or wireless links to a single super-gateway that is further connected to the Internet. All traffic is funneled through the gateways to the super-gateway and, eventually, to the Internet. Mobile users are free to roam as long as they are in the coverage area of the network. Small deployments may use a single gateway and forego the use of a super-gateway.

One of the design requirements we observed was the possibility to implement the protocol in the user space. In other words, the protocol does not have access to the forwarding engine other than through setting the routing table. We chose to transport the MRP messages using UDP packets for reasons of performance (TCP's overhead and delay would be too high). The choice requires a loss-tolerant protocol design. All three versions of the protocol presented in this paper interface with the kernel only through

the routing table changing calls and by intercepting ICMP packet delivery failure messages.

A. MRP On-Demand

When designing the protocol, we faced the well-known proactive vs. on-demand question commonly encountered in MANETs [22]. The first version of the protocol is purely on-demand: when a node is joining the network, it will ask the closest gateway or neighboring user nodes for a route (we will shortly elaborate on the criteria used to choose a gateway). Thus, the first step for a node that wishes to join the network is to broadcast *locally* a route discovery (RDIS) message. In contrast to existing MANET protocols (e.g., AODV and DSR), the RDIS packets in MRP are *not flooded* through the network, and are only received by the one-hop neighbors of the source. Initially, the joining node is in *disconnected* state.

We take advantage of the fact that all connected neighbors already know a route to the gateway and the significant routing metrics of those routes. Thus, all of the nodes receiving the RDIS message (one-hop neighbors of the joining node) will reply with a route advertisement (RADV) packet with the metrics of their current routes. The very first user node to join the network will receive this RADV packet only from the neighboring gateway(s).

The RADV packets are unicasted after a small random delay to avoid collisions. Without this delay, several RADV packets (triggered by the same RDIS) from (potentially hidden from each other) neighbors of the joining node, will collide at the joining node. If there are no connected neighboring nodes, or all of the advertisements are lost, the joining node will periodically broadcast the RDIS message until it receives an advertisement. The new node, joining the network will store all RADV packets (wait for a time chosen slightly longer than the longest random delay used to generate RADVs). Once the joining node receives all the RADVs, it will select one or more upstream routes as a function of the node's requirements and the offered routes.

The current design of the protocol is open to a variety of metrics that can be used for differentiating routes:

- **Hop-count:** the basic and common metric in routing protocols.
- **Route stability:** this metric allows a joining node (given the choice) to select routes that include stationary nodes (those routes are presumably more stable than those including mobile nodes).
- **Minimum delay:** important for delay sensitive applications (e.g., VoIP, telnet, etc).
- **Maximum bandwidth:** important for bandwidth intensive applications (e.g., ftp, peer-to-peer).
- **Minimum packet loss:** important for loss sensitive applications.

In contrast to other papers, we do not propose any new metrics or methods for computing the metrics above. Computing many of them is an area of active research [68], [70]; we simply include the possibility that they can be used in making routing decisions.




Fig. 3. MRP route establishment message sequence.

In the simulations, routes with high stability (i.e., employing stationary nodes) are preferred to routes with minimum hop-counts. In Section IV, preferring stability to hop-counts, is shown to reduce routing overhead, increase the data throughput, and reduce the energy consumption of the mobile nodes at the expense of longer routes.

Routing overhead and packet loss is reduced as the number of route breakage is minimized. In many applications of wireless mesh networks, the mobile nodes have limited energy resources (unless the nodes are mounted on vehicles); thus, power consumption becomes an important issue. If routes through stationary (and powered) nodes are preferred, packet forwarding of the mobile nodes is minimized. Other energy efficiency aspects of MRP are discussed in Section III-C.

Once the joining node selects a route (based on the RADVs it received), the node practically has a route *toward* the Internet (as all of the neighbor nodes have routes toward the Internet). The node can start to send data immediately toward the Internet. We call this the *half-connected* state (as the node has a route to the Internet, but none of the nodes in the Internet can reach it). In the second phase, the joining node registers with the gateway. The registration has the main function of providing a reverse path from the Internet to the joining node. We could have simply used the reverse path of the packets going toward the Internet to provide the return path, but that would imply that the routing protocol can monitor the forwarded data packets (a requirement we eliminated for portability, as explained above).

The registration process proceeds in two steps. In the first step, a registration request (RREG) is unicasted to the gateway (it goes through the MRP layer at every intermediate hop). As the RREG travels toward the gateway, intermediate nodes set up an entry in their routing tables enabling the return path (*from* the Internet to the joining node).

The existence of the return path assumes bidirectional links, but as pointed out in the literature [71], unidirectional links automatically exclude many of the common MAC protocols (e.g., MACA, 802.11, 802.16, and certainly all of those that require acknowledgments). Upon receipt of the RREG packet, the gateway sends a registration acknowledgment (RACK) directly to the joining node, and it also forwards the RREG message to the super-gateway (such that the super-gateway will know to which gateway to forward the packets for the joining node). If any of the RREG or RACK messages are lost, the joining node reverts to disconnected state and re-initiates the route acquisition process (as very likely the route it initially chose just became disconnected).

Upon receipt of a registration acknowledgment, the joining node enters the *fully-connected* state in which it can send and receive data to and from the Internet.

Figure 3 depicts the sequence of messages discussed above, from route discovery to registration acknowledgment. None of the four messages used to establish the routes (RDIS, RADV, RREG, and RACK) are *flooded* into the network. This is in clear contrast to other MANET routing protocols where the overhead/node ratio increases with the size of the network. Thus, we expect that the overhead of MRP will be significantly lower than that of existing MANET protocols, especially for large networks.

For wireless networks, the failure to forward a packet to the next hop may represent a route failure, but, more often than not, it might be simply the effect of interference or a temporary fading effect. Therefore, over-reacting to a packet loss (i.e., entering the disconnected state and re-initiating the route discovery process) may be detrimental to the protocol's performance. Hence, upon suspecting a link loss, instead of entering the disconnected state, the node enters a temporary state named *verify-link* state.

While in *verify-link* state, the node probes the availability of the suspected link by using special route-check packets (RCHKs) that are unicasted to the next hop. If a reply is received (we reused RADV for the reply), the node again becomes fully-connected; otherwise, it enters the disconnected state. If a node loses connection to the gateway, all of its children will lose connection to the gateway. Even if it discovers a new route to the gateway, the children will not be able to use this new route to *receive data from* the gateway (the reverse routes will be broken). Hence, when a node loses connection to the gateway, it will send a route error (RERR) message to all of its children. Each of those children will enter the disconnected state and re-initiate the route discovery process. A detailed analysis on the effect of the *verify-link* state on MRP's performance is presented later in this section.

Figure 4 depicts the finite state machine for the MRP protocol we just described. Since, in some ways, it resembles the on-demand protocols of MANET (route discovery initiated by the joining node and route errors initiated by packet loss), we will call this version *MRP on-demand* (MRP-OD).

Due to the tree structure of the routing graph (both to and from the gateways), and the use of hop-

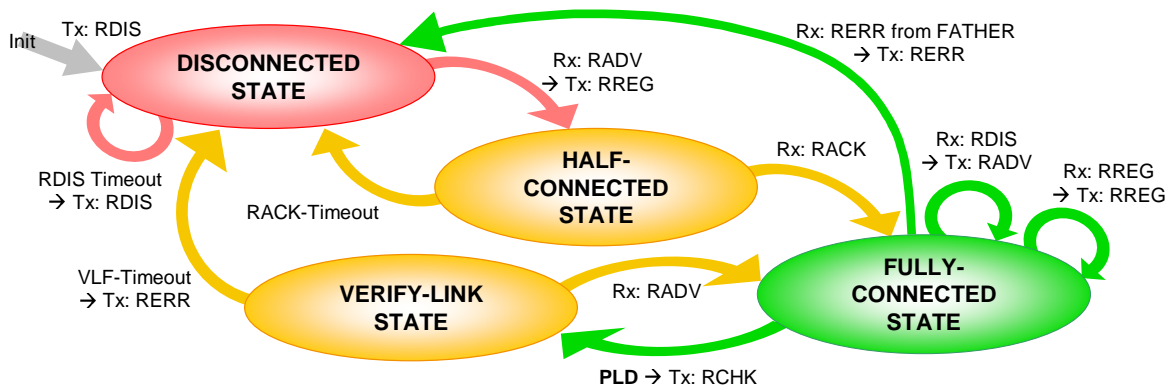


Fig. 4. Finite state machine for MRP-OD.

counts, the protocol is guaranteed to be loop free (it will be loop free as long as the routing metric is strictly positive).

B. MRP Beacon Mode

MRP-OD is fully functional; but, when a user node does not send any packets, it cannot detect that its route is no longer valid. Any packets from the gateway will be unable to reach the disconnected node.

A second version of MRP uses beacons to advertise routes: each node in fully-connected state (including the gateway) periodically sends beacons advertising the available routes. Each beacon acts as a gratuitous route advertisement (we used the same RADV packets as for MRP-OD). We will call this version *MRP beacon mode* (MRP-B). Notice that the MRP beacons are different from the 802.11 beacons in that they occur at different time intervals, are originated and consumed by the MRP rather than the MAC layer, and contain routing information instead of MAC layer management information.

In MRP-B, a node that wishes to join the network does not have to send an RDIS. Instead, it simply listens to the neighboring nodes slightly more than the beacon period and collects beacons (route advertisements). The route selection and registration proceeds identical to MRP-OD (the same states are used).

For detecting route disconnections, in addition to relying on packet forwarding failure, MRP-B utilizes the beacon messages originally intended for route setup. In MRP-B, each node monitors beacons sent from their parents. If a predefined number of beacons are missing, the node will enter the disconnected state and re-initiate the route discovery process. Also, as in the case of MRP-OD, if there is a packet forwarding failure, the node will enter the verify-link state and either become fully connected (if the RCHK is acknowledged) or will disconnect (if no reply is received). If link detection based only on beacons is employed, a real link breakage will be detected only after a considerable delay (several beacon periods).

The state diagram of MRP-B is practically identical to MRP-OD except for changing the triggers in some of the transitions, adding a periodic beaconing while in fully-connected state and a direct transition (triggered by beacon loss detection) from the fully-connected state to the disconnected state.

There is a trade-off between the beacon interval and the performance of the protocol: the more often a beacon is sent, the faster a node will be able to join the network and detect a route disconnection. However, it also increases the overall overhead of the routing protocol. Depending on the required performance, a reasonable range of beacon periods can be from tens of milliseconds to a few seconds.

We expect the MRP-B to exhibit better delay performance than MRP-OD because MRP-B is able to detect a route disconnection sooner than MRP-OD. This is true especially when traffic load is low and as a result, detecting forward failure does not promptly reflect a route change. On the other hand, MRP-B will have a higher overhead (due to the beacon messages).

C. Hybrid MRP

The two MRP versions we discussed are not mutually exclusive. They can be easily combined into a hybrid protocol (MRP-H). We simply use the same states as in MRP-OD and MRP-B and combine the state transitions from the two versions.

In MRP-H, the joining node broadcasts route discoveries (RDISs) and waits for route advertisements (RADVs) for a time equal to the minimum between the random delay of MRP-OD and the beacon period of MRP-B. The received RADVs include those generated in response to the RDISs, as well as the beacons. The joining node then selects the route and registers with the gateway. A route error can be discovered by either missing beacons or packet forwarding errors. The state diagram of MRP-H is similar to the one corresponding to MRP-OD (Fig. 4), except that there are multiple triggers for the state transitions.

We expect MRP-H to perform better than either MRP-B or MRP-OD as MRP-H detects link failures faster than either of the two versions and is able to discover better routes faster: the pool of available routes at the moment of the route selection will be larger due to the redundancy offered by the beacons and triggered answers. (This feature only makes a difference if some of the RDIS or RADV packets are lost due to noise/interference.)

Since none of the three versions of the protocol uses flooding to establish or repair routes, all three should scale well to large networks.

With respect to energy consumption, a mobile user with limited energy resources can minimize the energy expended for routing by employing MRP-OD (where no periodic beacons are employed). The rest of the network may run the more energy-expensive MRP-H or MRP-B. Thus, for mobile nodes, MRP can minimize both the energy expended for *routing*, as well as for *forwarding* (by selecting routes that avoid the mobile nodes).

If the nodes in a large mesh network are powered up simultaneously (e.g., after a power outage), the first tier of nodes (with direct connectivity to the gateways) will discover routes first, then, the second tier (nodes two hops away from the gateway), and the process will continue from the center to the edges of the network.

The design of MRP is based on the assumption that most data flows are established between client nodes and the gateway. Routes between two arbitrary client nodes may be sub-optimal. In particular, two client nodes (unless they are parents of one another), will always communicate through their common parent (which can be several hops away from either of them, and possibly the gateway); however, client-to-client communication is assumed to be infrequent in WMNs (for most applications).

D. Verify-Link State

The verify-link state is designed to reduce the uncertainty introduced by wireless links and to avoid unnecessary disconnections and route re-establishments. We define a *false disconnection* as the case where an MRP node erroneously enters a disconnected state due to temporarily disabled links (e.g., caused by fading or interference). We define a *true disconnection* as the case where a node enters a disconnected state due to permanent link breakage (e.g., caused by mobility or node failure).

False disconnections and consequent route re-establishment may cause significant performance degradation. Frequent re-routing will also reduce the available bandwidth for end users due to increased routing overhead. Therefore, one of the design goals of MRP is to suppress false disconnections and to identify and promptly respond to true disconnections. To this end, we introduced the verify-link state whose state transition and packet exchange (i.e., RCHK and RADV) were detailed in Section III-A.

One of the key parameters in the verify-link state mechanism is the verify-link-failure (VLF)-timeout value, which controls how long a node should stay in verify-link state before it transitions to disconnected state (assuming that no reply was received). The choice of a VLF-timeout value is a trade-off: a large VLF-timeout will reduce false disconnections but causes delay in true disconnection situations. On the other hand, a small VLF-timeout will ensure prompt handover in a true disconnection but increases the number of false disconnections. To study the efficiency of the verify-link state and to identify a working range of VLF-timeout values, we investigated the behavior of MRP through simulation experiments over a wide range of VLF-timeout values.

Since traffic load intensity and mobility intensity are expected to affect false disconnections and true disconnections, respectively, we tested VLF-timeout values for different traffic and mobility scenarios. Three performance metrics were used to measure how correctly and efficiently routes are maintained when different VLF-timeout values are used with the forward failure detection: routing overhead, packet delivery ratio (PDR), and end-to-end delay. The metrics are defined as follows:

$$Routing\ Overhead \equiv \frac{\sum_{n=1}^N \sum_{p=1}^{P_n^{cs}} b_{n,p}}{T_{sim} \times N} \quad (bps/node), \quad (1)$$

$$PDR \equiv \frac{\sum_{n=1}^N P_n^{dr}}{\sum_{n=1}^N P_n^{ds}} \quad (\%), \quad (2)$$

$$Delay \equiv \frac{\sum_{n=1}^N \sum_{p=1}^{P_n^{dr}} D_{n,p}}{\sum_{n=1}^N P_n^{dr}} \quad (ms/packet), \quad (3)$$

where:

N is the total number of nodes,

P_n^{cs} is the total number of broadcast/unicast routing control packets sent (either generated or forwarded) by the n^{th} node (including beacon packets),

$b_{n,p}$ is the number of bits in the p^{th} packet received by the n^{th} node,

T_{sim} is the total simulation time, and

P_n^{dr} is the total number of unicast data packets received by the n^{th} node,

P_n^{ds} is the total number of unicast data packets sent by the n^{th} node, and

$D_{n,p}$ is the end-to-end delay experienced by the p^{th} packet received by the n^{th} node.

We use Poisson traffic with data packets serving as probes. Every user node has two flows: one toward the gateway and the other one from the gateway. Packet size is 1500 bytes. The traffic intensity is varied by changing the average packet interval ranging from 1.5 s to 4.0 s (in this case, the number of mobile nodes was fixed at 16). Mobility intensity is controlled by the number of mobile nodes from 0 to 49 (in this case, the packet interval was fixed to 3.0 s). Randomly-deployed mobile nodes move according to the random waypoint mobility model [72] with zero pause time, a maximum speed of 10 m/s, and a minimum speed of 1 m/s (to avoid the drawbacks highlighted in [73]). The mobile nodes move inside a rectangular area with a single gateway in the center and 144 uniformly deployed fixed nodes. Total simulation time is 400 s. The size of the rectangular area is determined by the number of fixed nodes to ensure constant node density (approximately 25 nodes per km²).

Figure 5 shows the performance of MRP as a function of VLF-timeout under different traffic intensity. Only the result for the MRP-H is presented to avoid cluttering. The VLF-timeout was increased from 6 ms to 12 s. The results for the first (0 s) and the last VLF-timeout value (400 s) are not shown to scale. They are included to provide insight on the MRP's behavior for extreme values. The VLF-timeout of 0 s represents the case where the verify-link state is completely disabled; in this case, one forwarding

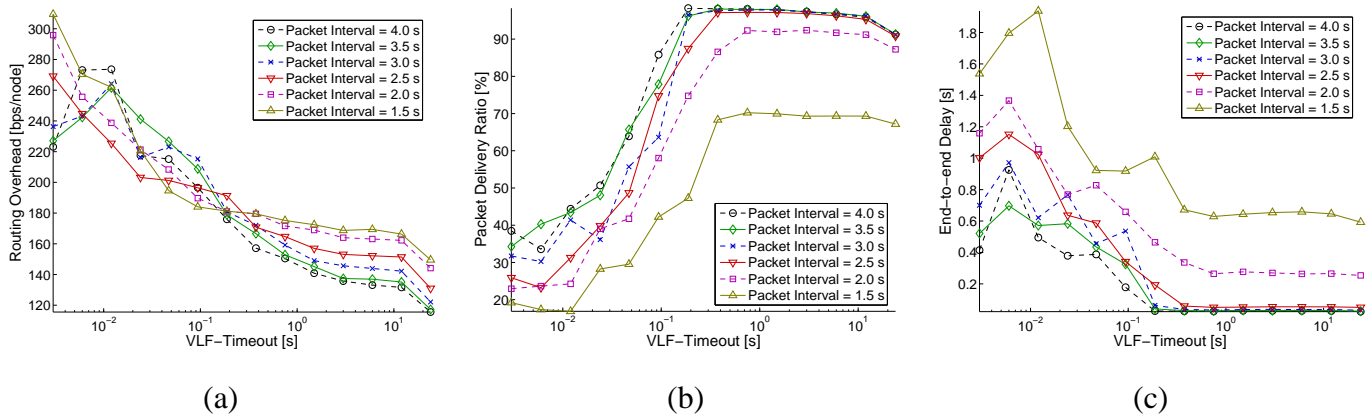


Fig. 5. The effect of verify-link state as a function of VLF-timeout and different average packet intervals: (a) routing overhead, (b) packet delivery ratio and, (c) end-to-end delay.

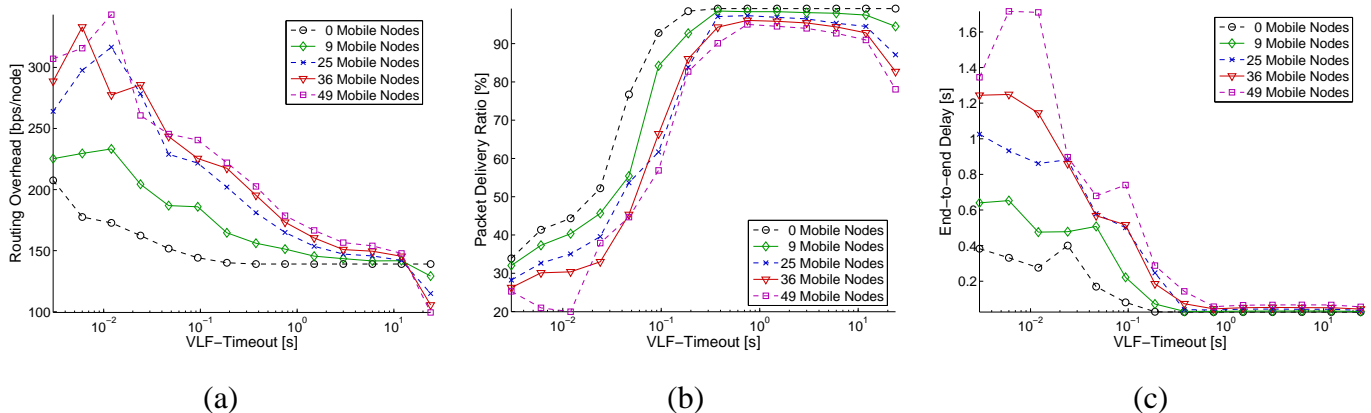


Fig. 6. The effect of verify-link state as a function of VLF-timeout and mobility intensity: (a) routing overhead, (b) packet delivery ratio and (c) end-to-end delay.

failure is interpreted as a route failure and, thus, it triggers a route rediscovery. The VLF-timeout of 400 s represents the case where packet forward failures are always ignored throughout the simulation time.

For small VLF-timeout values, the curves show low PDR and high end-to-end delay due to false disconnections. A false disconnection will trigger a route rediscovery that implies the exchange of several control packets (RDIS, RADV, RREG, and RACK). If the falsely disconnected node has children nodes, RERR packets are generated which leads to a cascade of false disconnections and route re-establishments.

As the VLF-timeout is increased, the overall performance is improved due to the reduced probability of false disconnections. However, if VLF-timeout is too large, true disconnections (from mobility) are temporarily ignored, leading to loss of packets for the mobile nodes.

Figure 6 shows the performance of MRP-H as a function of VLF-timeout and mobility intensity. When the number of mobile nodes is increased, the loss from ignored true disconnection increases.

Considering the results in Figures 5 and 6, the VLF-timeout should be at least 1 s (as for lower values, the false disconnections considerably lower the PDR in Fig. 5(b)). Figure 6(b) shows that the

PDR decreases steadily as the VLF increases over 1 s, as true disconnections are identified too late. In our simulations, we chose a VLF-timeout of 3 seconds. The performance could likely be improved if the value of the VLF-timeout is determined adaptively; but, we do not explore that option in this paper.

IV. PERFORMANCE EVALUATION

In this section, we present the results of the simulations which we performed using QualNet [74] to evaluate the performance of MRP. QualNet is the commercial version of GloMoSim [75] (which uses Parsec [76], a parallel discrete-event simulator). We chose QualNet because it provides accurate models (high level of detail) at each layer, including the physical and MAC layers. Results indicate that the accuracy of the models has a significant impact on the simulation results for wireless networks [77], [78].

A. Performance Metrics

We compared the performance of MRP with other routing protocols in terms of five metrics: routing overhead, packet delivery ratio (PDR), throughput, end-to-end delay, and average hop-count.

We employ Internet traffic as well as Poisson traffic in the performance evaluation; we use network throughput as a performance metric for Internet traffic and PDR as a performance metric for Poisson traffic. In the Internet traffic scenarios, the applications use mostly TCP as a transport layer, and the PDR does not capture the effect of packet loss on the TCP throughput.

The routing overhead, PDR, and the end-to-end delay are defined in Section III-D. We define the network throughput and hop-count as:

$$Network\ Throughput \equiv \frac{\sum_{n=1}^N \sum_{p=1}^{P_n^{dr}} b_{n,p}}{T_{sim}} \quad (bps), \quad (4)$$

$$Average\ Hop - Count \equiv \frac{\sum_{n=1}^N \sum_{p=1}^{P_n^{dr}} h_{n,p}}{\sum_{n=1}^N P_n^{dr}} \quad (hops\ per\ packet), \quad (5)$$

where N , P_n^{dr} , $b_{n,p}$, T_{sim} , n and p are defined in Section III-D, and $h_{n,p}$ is the number of hops traversed by the p^{th} packet of the n^{th} node.

We believe that the chosen metrics can be traded against each other (e.g., higher overhead for lower delays, lower throughput for lower delays, etc.), and a fair comparison has to consider all of the metrics.

B. Simulation Setup

For the physical and MAC layers, we used the default IEEE 802.11b in physical and MAC models (in RTS/CTS mode). The error rate in Qualnet (version 3.8) gradually decreases from approximately 0% at 270 m to almost 100% at 300 m. The physical layer uses a two-ray path-loss, constant shadowing model without fading.

For the higher layers, we used the standard TCP/IP protocol stack, and we implemented MRP as described in Section III. For MRP-B, we chose a beacon interval of 1 s. For MRP-OD and MRP-H, we chose the small random interval for sending route requests of 2 s. For MRP-H, we chose the beacon interval of 2 s (we shall see that the performance is comparable with MRP-B with 1 s beacon interval).

In order to assess the contribution of preferring stability to hop counts to the overall performance of the protocol, we simulate MRP-H with and without route stability and we labeled the corresponding results with MRP-H and MRP-S respectively.

Simulations are run for two different traffic types: Poisson (Section IV-C) and Internet traffic (Section IV-D). For the Poisson traffic, user data packets serve as probes, and the performance metrics indicate how correctly and efficiently routes are maintained by different protocols. On the other hand, the Internet traffic model is used to evaluate the performance of the protocol in a more realistic environment. The Internet traffic is created by mixing several popular Internet applications based on real Internet traffic measurement. The traffic composition trend captured from the real Internet traffic [79] was accurately reproduced in our simulation using multiple application protocol models provided by QualNet [74]: HTTP, TELNET, FTP, DNS, and VBR models. Traffic composition ratio was reproduced at all detail levels: flow, packet, and byte.

Simulations were performed to evaluate the influence of the following parameters:

- traffic load,
- network size,
- degree of mobility,
- perturbation, and
- intra-mesh traffic ratio,

where the perturbation represents the degree of randomness in geographical node distribution, and the intra-mesh traffic ratio indicates the percentage of the user traffic flows whose both ends are user nodes (as opposed to the typical mesh network case where one end is a gateway.)

To avoid the large number of graphs that would result if we performed every experiment for every possible setting, we chose a *base case* and varied one parameter at a time. For each parameter variation, we present the four metrics defined in Section IV-A for all of the protocols under consideration. For the

base case, we used the following scenario:

- The Poisson traffic has one incoming and one outgoing flow for every user with a mean inter-arrival rate of one packet every 1.5 seconds.
- The Internet traffic has (on average) one flow for every user (in the proportion reported in [79], i.e., 75% HTTP, 5% FTP, etc.).
- The stationary network consists of 144 nodes uniformly distributed in a rectangular area of 2.4×2.4 km^2 .
- There are 36 mobile nodes deployed in random locations. They move according to the random waypoint mobility model [72] with zero pause time, a maximum speed of 10 m/s, and a minimum speed of 1 m/s.
- There is a single gateway at the center of the network.
- The stationary nodes are deployed with zero perturbation (i.e., they are deployed in a grid of size 250 m).
- There is no intra-mesh traffic.
- Total simulation time is 400 s.

To evaluate the performance of MRP, we chose six well-known routing protocols representative of different classes of routing protocols: AODV [28], [46], DSR [29], LANMAR [51], [52], OLSR [40], RIP2 [45], and ZRP [47]. AODV and OLSR are standardized by the MANET workgroup of IETF. AODV and DSR are widely known *on-demand (reactive)* MANET routing protocols. LANMAR is designed for routing in large scale ad-hoc networks with group mobility. In QualNet implementation, LANMAR uses Fisheye State Routing (FSR) protocol for local scope routing. OLSR is based on the Inria implementation [80], and it is an optimized *link-state*, table-driven protocol designed for MANETs. RIP2 is a table-driven *distance-vector* protocol originally designed for wired Internet. The RIP2 model in QualNet follows the Cisco implementation. ZRP is a hybrid ad-hoc routing protocol where proactive and reactive approaches are employed for intra- and inter-zone routing, respectively.

For each scenario, we simulated the network thirty times with different random seeds (resulting in different initial placements and mobility for the mobile nodes and in different inter-arrival times of the offered load). In every graph, we present the average of the thirty experiments. To avoid cluttering the figures, confidence intervals are not presented.

C. Poisson Traffic

In this section, simulation results for Poisson traffic are examined for five different scenarios where traffic load, network size, mobility, perturbation, and intra-mesh traffic ratio are varied from the base case.

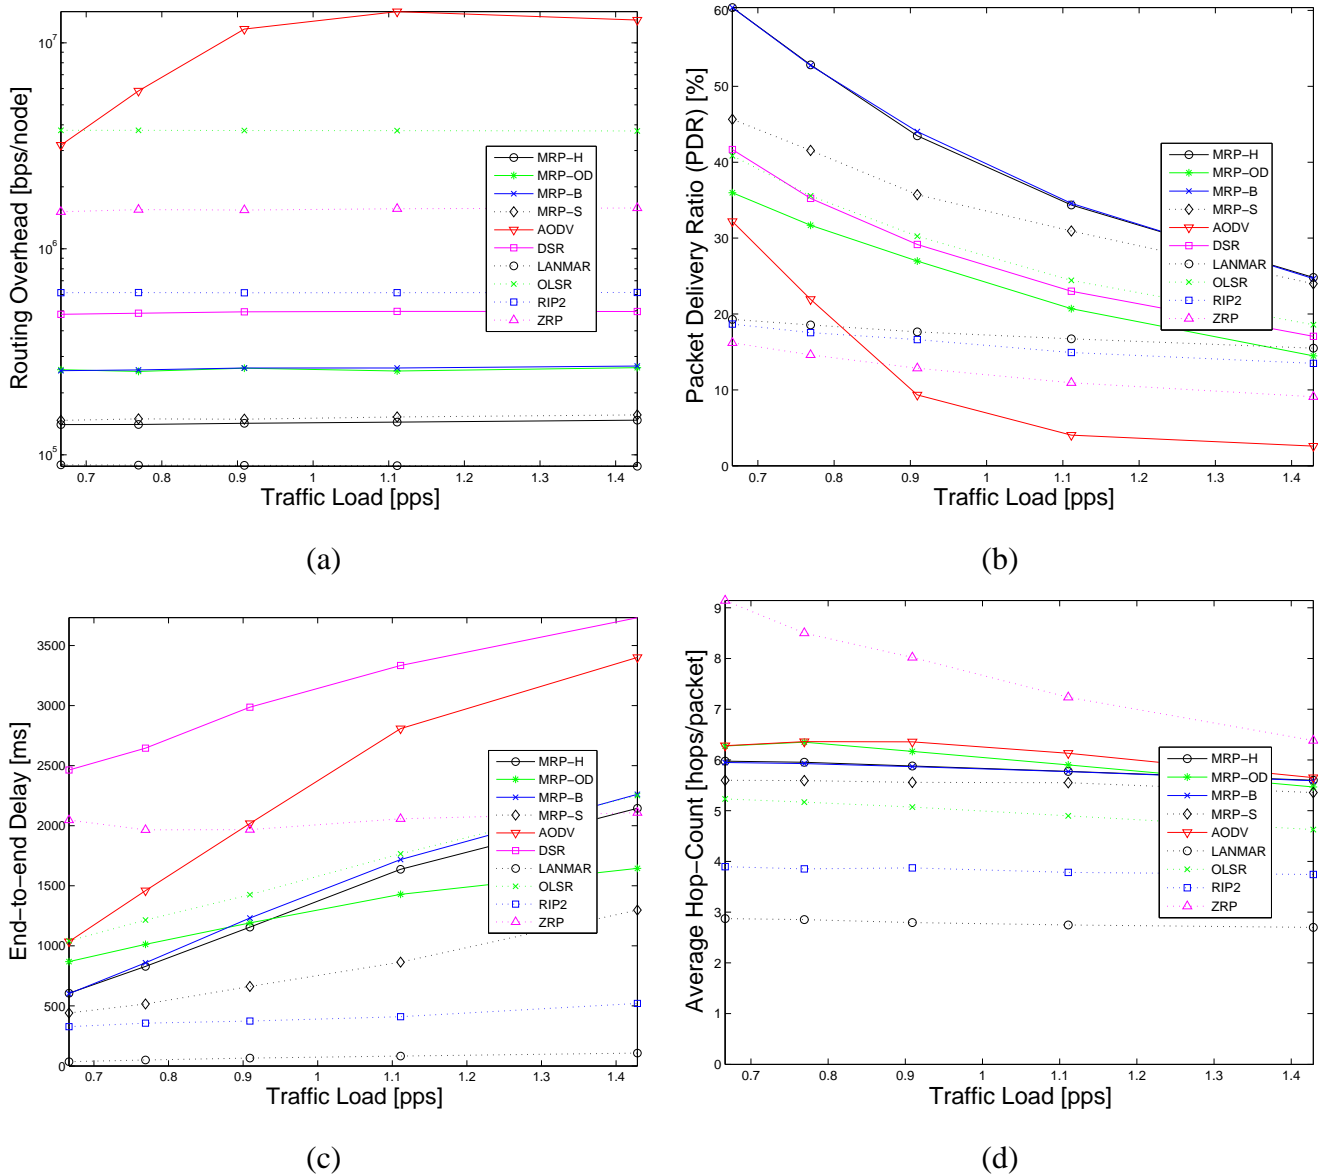


Fig. 7. The performance of the routing protocols for Poisson traffic (a) routing overhead, (b) packet delivery ratio, (c) end-to-end delay, and (d) average hop-count as a function of traffic load.

Except for the scenario of varied intra-mesh traffic, all nodes have packet streams flowing both to and from the gateway. The same traffic load is applied in both directions.

1) *Traffic Load:* We increase the load of the Poisson traffic by reducing the mean packet inter-arrival time from 1.5 s to 0.67 s (i.e., from 0.67 to 1.5 packets per second).

Figure 7 depicts the simulation results for the overhead, PDR, delay, and hop-count, respectively. In each graph, curves for nine routing protocols (including four versions of MRP) are plotted as the offered load increases.

Figure 7(a) shows that the routing overhead of AODV rises as traffic load is increased. The main reason behind this increase is the corresponding increase in lost packets (that triggers the route discovery

process). DSR shows lower overhead; however, the overhead of source route header is not included in the overhead calculation. In contrast, MRP and all the proactive protocols show overhead immunity to the traffic load. All versions of MRP and LANMAR form the group of the lowest overhead of all considered protocols. MRP-H has a very low overhead, due to the larger beacon period (twice as large as MRP-B).

Figure 7(b) shows that, as the offered traffic load intensifies, the PDR of AODV, DSR, OLSR, and ZRP drops faster than MRP. Although MRP-S shows smaller PDR than the other versions of MRP, it still outperforms the other protocols. The drop in PDR is due to the packet loss in the queues, as well as lost routes due to the routing protocols' attempt to restore failed routes (or what are considered to be failed routes due to packet drops). RIP2 shows the lowest PDR for the highest traffic load, indicating the unsuitability of wireline-oriented protocol for ad-hoc or mesh networks.

The delay increase in almost all of the protocols shown in Fig. 7(c) is due to the larger queuing delays resulting from the increase in offered load. MRP-OD has a higher delay than MRP-B and MRP-H, as it discovers broken routes later than the other two. The low delays of RIP2 and LANMAR have to be considered with the very low PDR in Fig. 7(b): the delay statistics consider only the few packets that reach their destinations.


Figure 7(d) shows that the hop-count of each packet is almost insensitive to an increase in traffic load. Two on-demand protocols (AODV and MRP-OD) show a slight increase in the hop-count. The hop-count result of DSR is not included because measuring the TTL field of IP packets is not possible in QualNet due to its source routing implementation.

As expected, among the versions of MRP, MRP-H exhibits the best overall performance as it combines strengths from both MRP-B and MRP-OD. MRP-S shows smaller delays and hop-count at the expense of a decrease in PDR.


2) *Network Size*: In this scenario, the number of fixed nodes is increased while keeping the network density constant (i.e., increasing the network diameter). The four graphs for overhead, PDR, delay, and hop-count are shown in Fig. 8, as a function of the network size.

In Fig. 8(a), MRP shows almost constant overhead, while all of the other protocols show the overhead increasing with the network size. This is expected because for MRP, the overhead for each node does not increase with the total number of nodes (see Section III). As a result, MRP also maintains the highest PDR and the lowest delay. This supports the claim that, for WMNs, MRP scales better than the existing protocols.

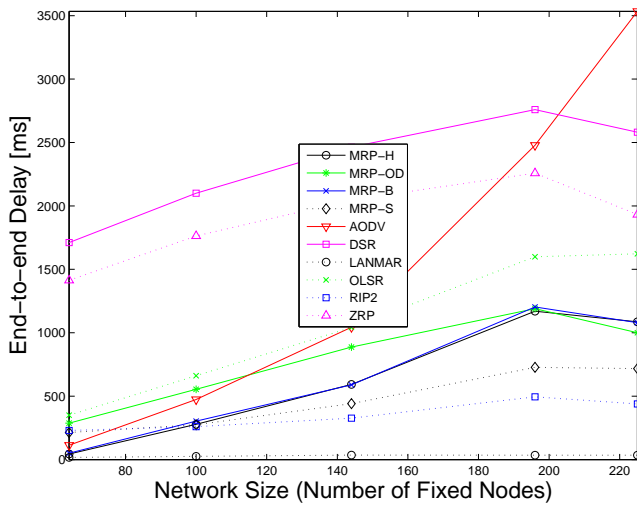
Fig. 8(b) shows that the PDR of all routing protocols decreases with the increase in network size. This is expected as the network capacity most likely decreases with the increase in network size. For large networks (i.e., over 100 fixed nodes), broken links due to mobility trigger route rediscoveries, which incur



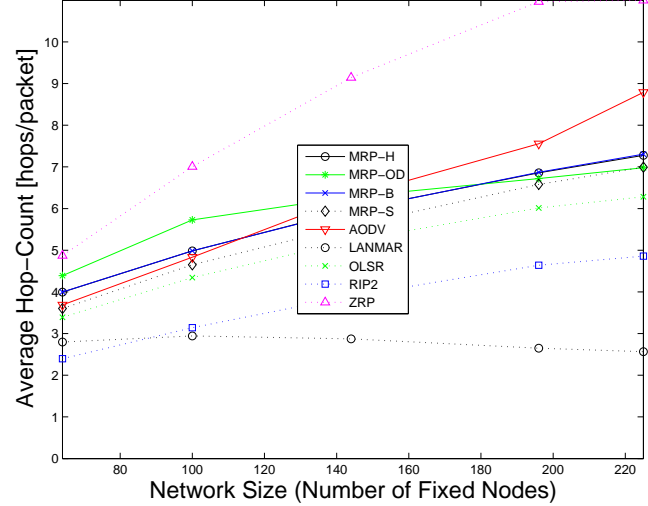
(a)



(b)



(c)



(d)

Fig. 8. The performance of the routing protocols for Poisson traffic (a) routing overhead, (b) packet delivery ratio, (c) end-to-end delay, and (d) average hop-count as a function of network size.

high overhead with flooding. Simulation results for varied mobility will be presented and discussed in detail in the following section.

Figure 8(c) shows high delay for DSR and ZRP compared to other protocols. All versions of MRP (and especially MRP-S) show good delay performance for scaled network size.

The hop-count shown in Fig. 8(d) directly reflects the increase in network size. The versions of MRP, AODV, and OLSR show hop-counts close to the ideal hop-count. ZRP shows longer than ideal routes indicating suboptimal routes, while LANMAR and RIP2 only route successfully to/from nodes closer to the gateway, and thus show shorter hop-counts.

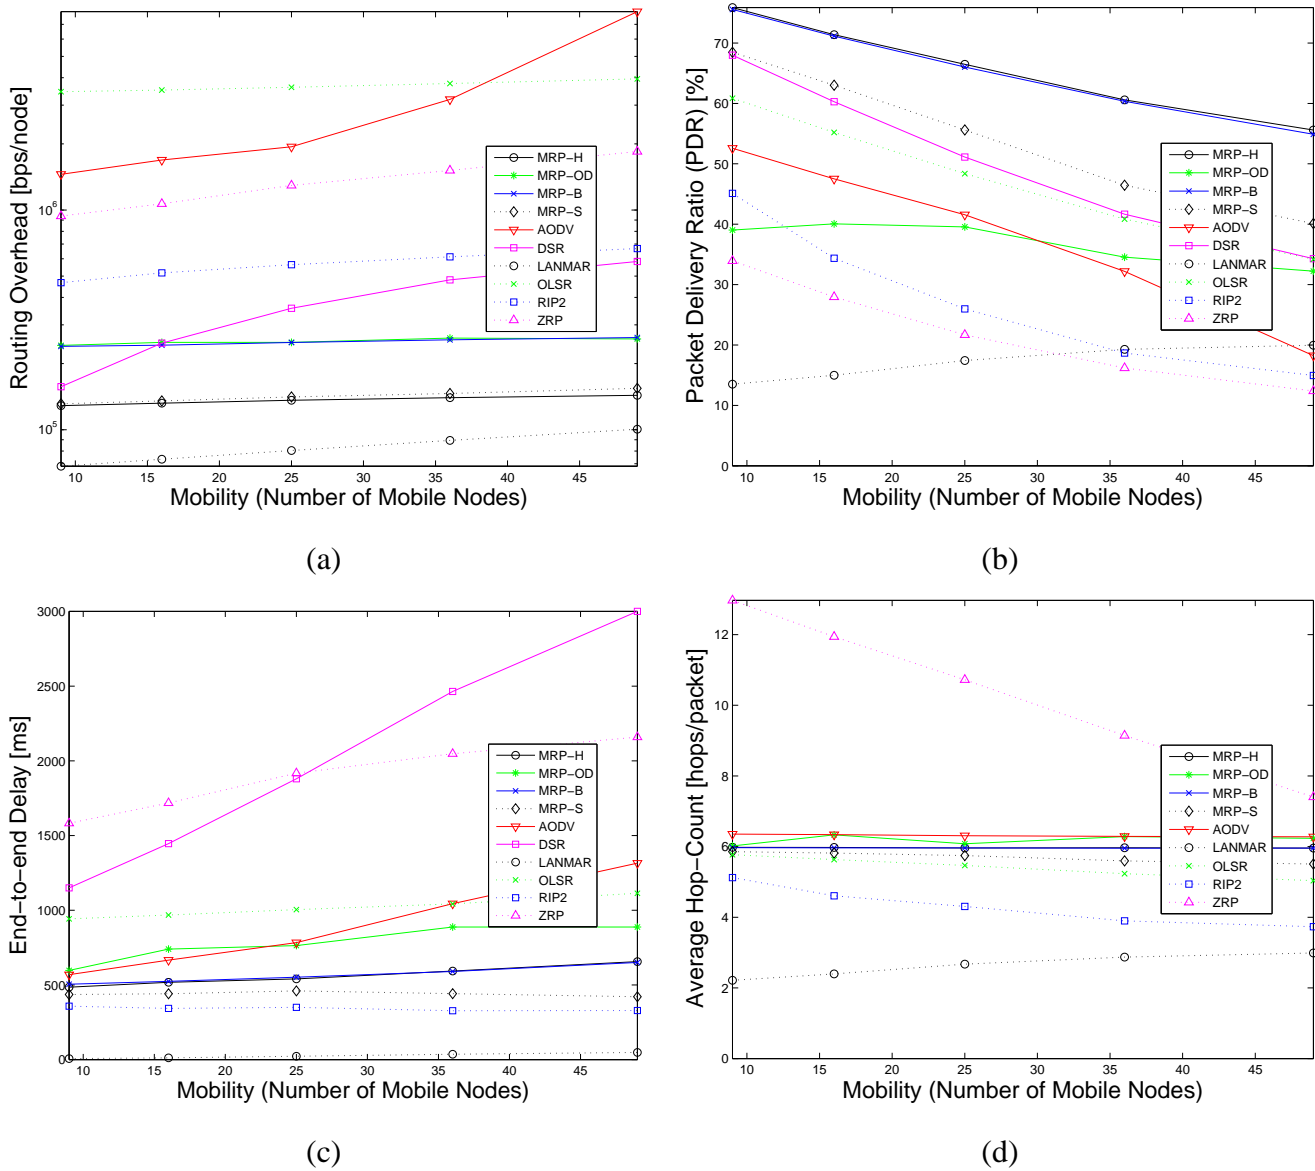


Fig. 9. The performance of the routing protocols for Poisson traffic (a) routing overhead, (b) packet delivery ratio, (c) end-to-end delay, and (d) average hop-count as a function of mobility.

3) *Number of Mobile Nodes*: In this scenario, we evaluate the performance of MRP and other protocols when the number of mobile nodes increases from 9 to 49. The results are shown in Fig. 9.

As in the cases of traffic load and network size, Fig. 9(a) shows that MRP shows a low and steady overhead as the number of mobile nodes is increased. In Fig. 9(b), MRP shows higher PDR than other protocols as the mobility intensifies. RIP2 shows a sharp drop in PDR as it was designed for stationary networks with sporadic topology changes. For a large number of mobile nodes, DSR shows rapidly increasing delay, as shown in Fig. 9(c). The result supports the claim that, for WMNs, MRP supports mobility more efficiently than existing protocols.

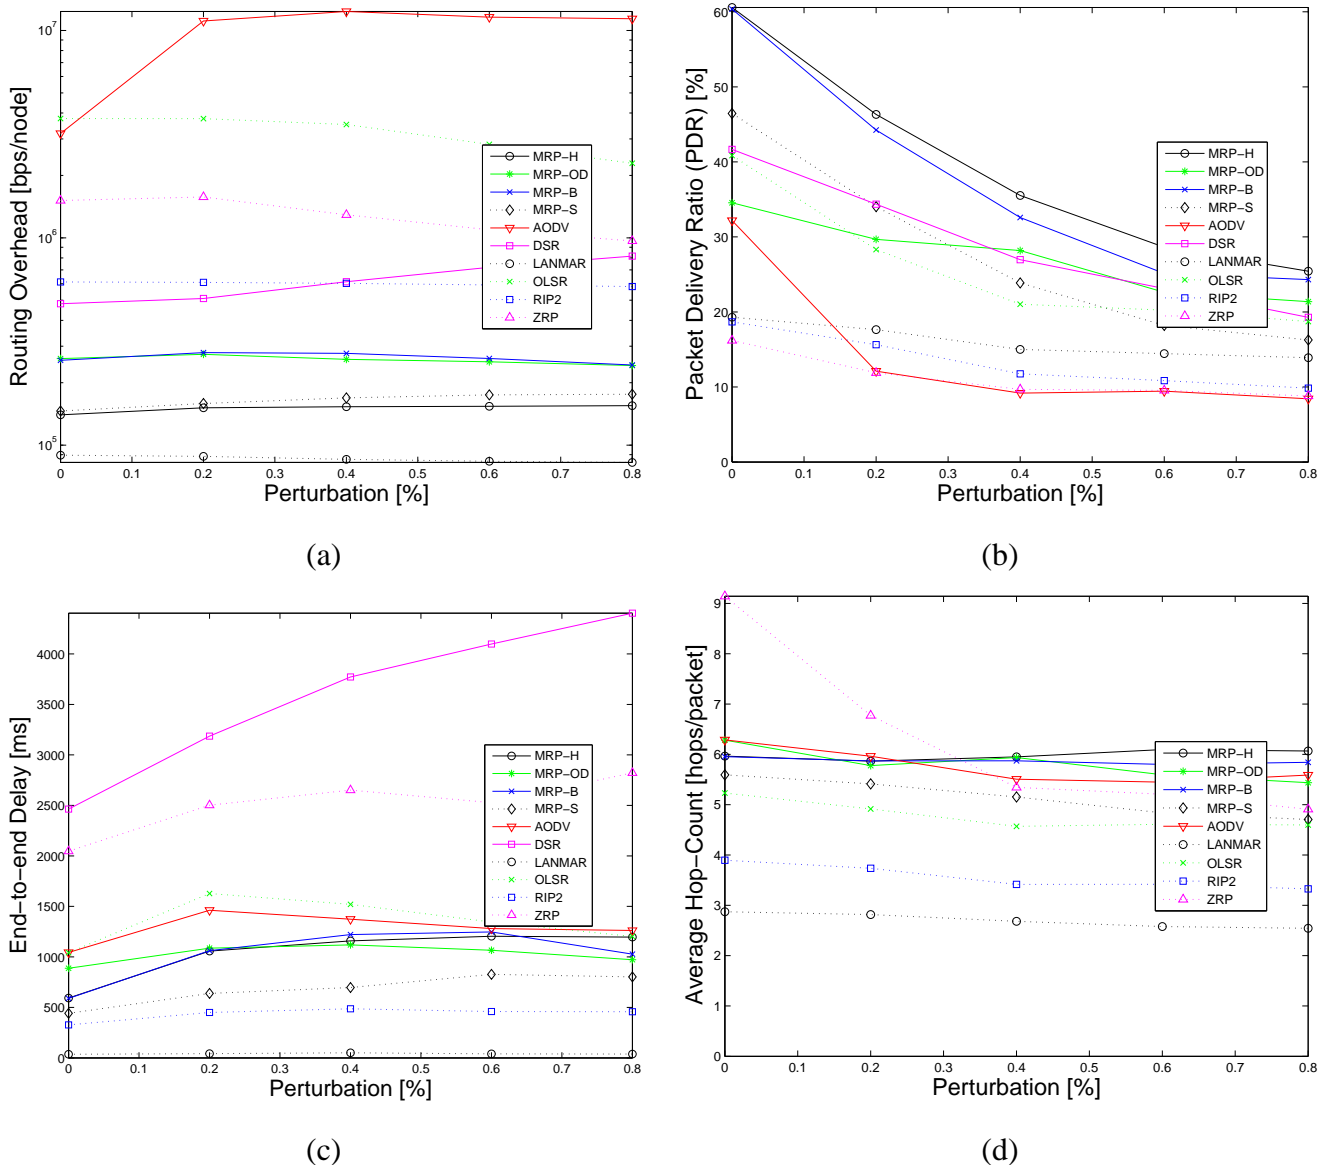


Fig. 10. The performance of the routing protocols for Poisson traffic (a) routing overhead, (b) packet delivery ratio, (c) end-to-end delay, and (d) average hop-count as a function of the perturbation in stationary nodes grid deployment.

4) *Perturbation*: The performance of MRP for increased randomness in the deployment of the stationary nodes is shown in Fig. 10. For this scenario, the stationary nodes are perturbed uniformly from their ideal grid positions. The degree of perturbation indicates by how much a node can deviate from its ideal position. A perturbation of 100% indicates that a node can be moved by up to one grid size in any direction.

As the perturbation increases, the network becomes less regular and hot spots are created. In Fig. 10(a), DSR and AODV show increased overhead, while all versions of MRP remain low and almost constant. OLSR and ZRP show a decrease in overhead; in OLSR, the multipoint relays (MPRs) benefit from non-uniform node distribution. Likewise, in ZRP, as perturbation increases, fewer routing zones are required to cover the user nodes, and thus, less inter-zone, on-demand routing is needed.

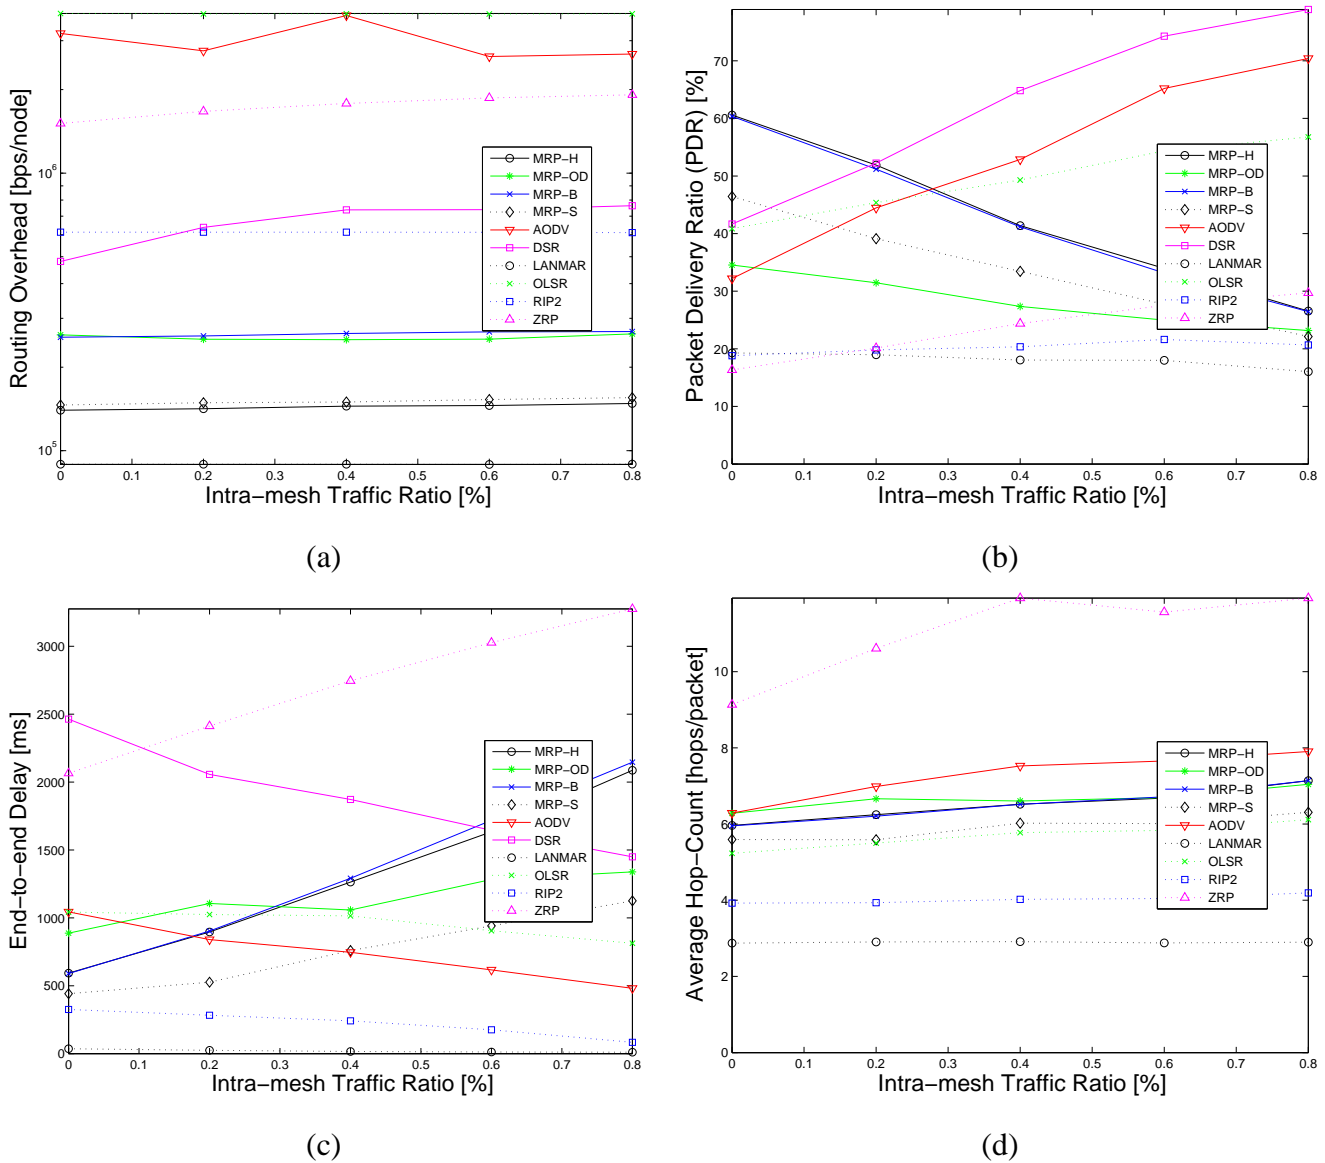


Fig. 11. The performance of the routing protocols for Poisson traffic (a) routing overhead, (b) packet delivery ratio, (c) end-to-end delay, and (d) average hop-count as a function of the intra-mesh traffic ratio.

Figures 10(b) and (c) show decreased PDR and increased delays as the network becomes less regular. MRP shows better PDR and delay performance than other protocols. Among the versions of MRP, MRP-H shows the highest PDR while MRP-S shows the lowest hop-count and delay.

Figure 10(d) shows that the hop-count for MRP increases slightly as the network becomes less regular; this is due to the route selection algorithm that favors stability to lower hop-counts.

5) *Intra-mesh Traffic*: In this scenario, the intra-mesh traffic ratio is increased from zero to 0.8. The ratio is calculated as the number of intra-mesh traffic flows over the total number of user traffic flows. For example, the ratio of 0.8 indicates that 80% of the user traffic flows are both originated and terminated within the mesh network (i.e., neither end of a traffic flow is a gateway.) The higher the intra-mesh traffic

ratio, the stronger the characteristics of *ad-hoc* networks rather than those of *mesh* networks. The results are presented in Fig. 11.

In Fig. 11(a), on-demand protocols such as AODV, DSR, and ZRP (ZRP has both on-demand and table-driven components) show a small increase in overhead as intra-mesh traffic ratio is increased. These protocols have to discover routes for destinations located farther away from the source due to increased intra-mesh traffic. Table-driven protocols show almost constant overhead.

Figure 11(b) shows increases in PDR for AODV, DSR, and ZRP; while for MRP, the PDR decreases as the intra-mesh traffic ratio increases. This result is expected as AODV, DSR, and ZRP are designed for generic ad-hoc networks where all of the user nodes are assumed to be uniformly likely to send data to any destination in the network. In contrast, routes in MRP are always routed through the common parent of the source and the destination (often the gateway); thus, the routes are sub-optimal. The extent of this effect is shown in Fig. 11(d).

D. Internet Traffic

In this section, simulation results for Internet traffic are examined for five different scenarios where traffic load, network size, mobility, perturbation, and intra-mesh traffic ratio are varied from the base case. The purpose of using Internet traffic is to evaluate the performance of MRP and other protocols in a more realistic environment. The Internet traffic model is created by mixing several popular Internet applications based on real Internet traffic measurements as depicted in Section IV-B.

1) *Traffic Load*: In this scenario, we vary the relative traffic load of Internet traffic from 0.2 to 1.4. The relative traffic load is controlled by changing the number of active users while maintaining the ratio among different applications. For a traffic load of 1.0, every node is likely to have one active application. If the load is greater than 1.0, some nodes are likely to have more than one active application.

Figure 12(a) shows that the routing overhead of on-demand protocols (AODV, DSR, and ZRP) increases with the traffic load. This result and its cause is similar to the case of Poisson traffic: more packets tend to trigger more route discoveries for those protocols. In contrast, MRP and all of the table-driven protocols show overhead immunity to the traffic load.

As the offered traffic increases, the network throughput approaches network capacity due to the feedback loop of TCP that adjusts the transmission rate to the available bandwidth. Figure 12(b) shows that AODV, DSR, and OLSR have lower saturation throughput than MRP. The lower throughput is due to the packet loss in the queues and failed routes due to packet drops. MRP-OD has lower throughput than MRP-H and MRP-B as it discovers broken routes later than the other two.

In Fig. 12(c), the delay increase for all protocols is due to the larger queuing delays resulting from increased offered load. DSR shows higher delay than other protocols. The low delay of LANMAR has

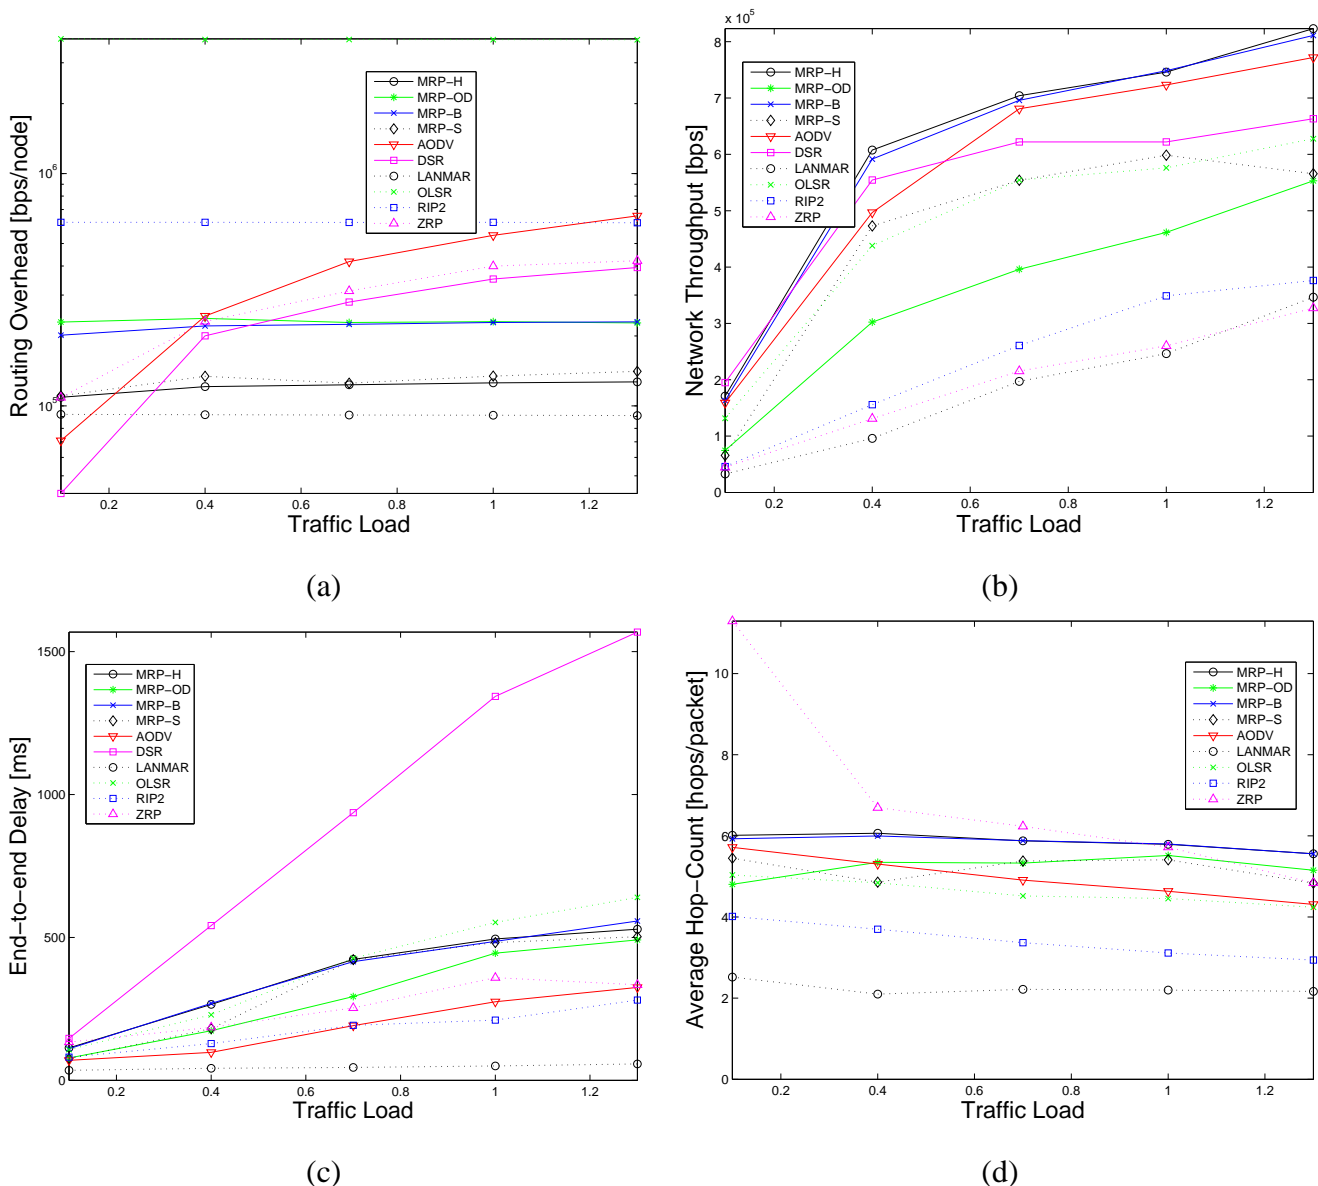


Fig. 12. The performance of the routing protocols for Internet traffic (a) routing overhead, (b) network throughput, (c) end-to-end delay, and (d) average hop-count as a function of traffic load.

to be considered in conjunction with its very low throughput (the delay statistics only consider the few packets delivered to their destinations). MRP-S has a higher delay than other versions of MRP due to the more frequent disconnections.

2) *Network Size*: We increase the network size in the same manner as in the Poisson traffic case, i.e., increase the number of fixed nodes while keeping the density constant. The results are shown in Fig. 13.

Figure 13(a) shows that the overhead of MRP stays almost constant, while that of other protocols (especially OLSR) rises as the network size increases. As in the case of Poisson traffic, this result is expected in MRP as the overhead for each node does not increase with the total number of nodes.

Figure 13(b) shows that MRP-H and MRP-B maintain the highest throughput. While packet loss simply

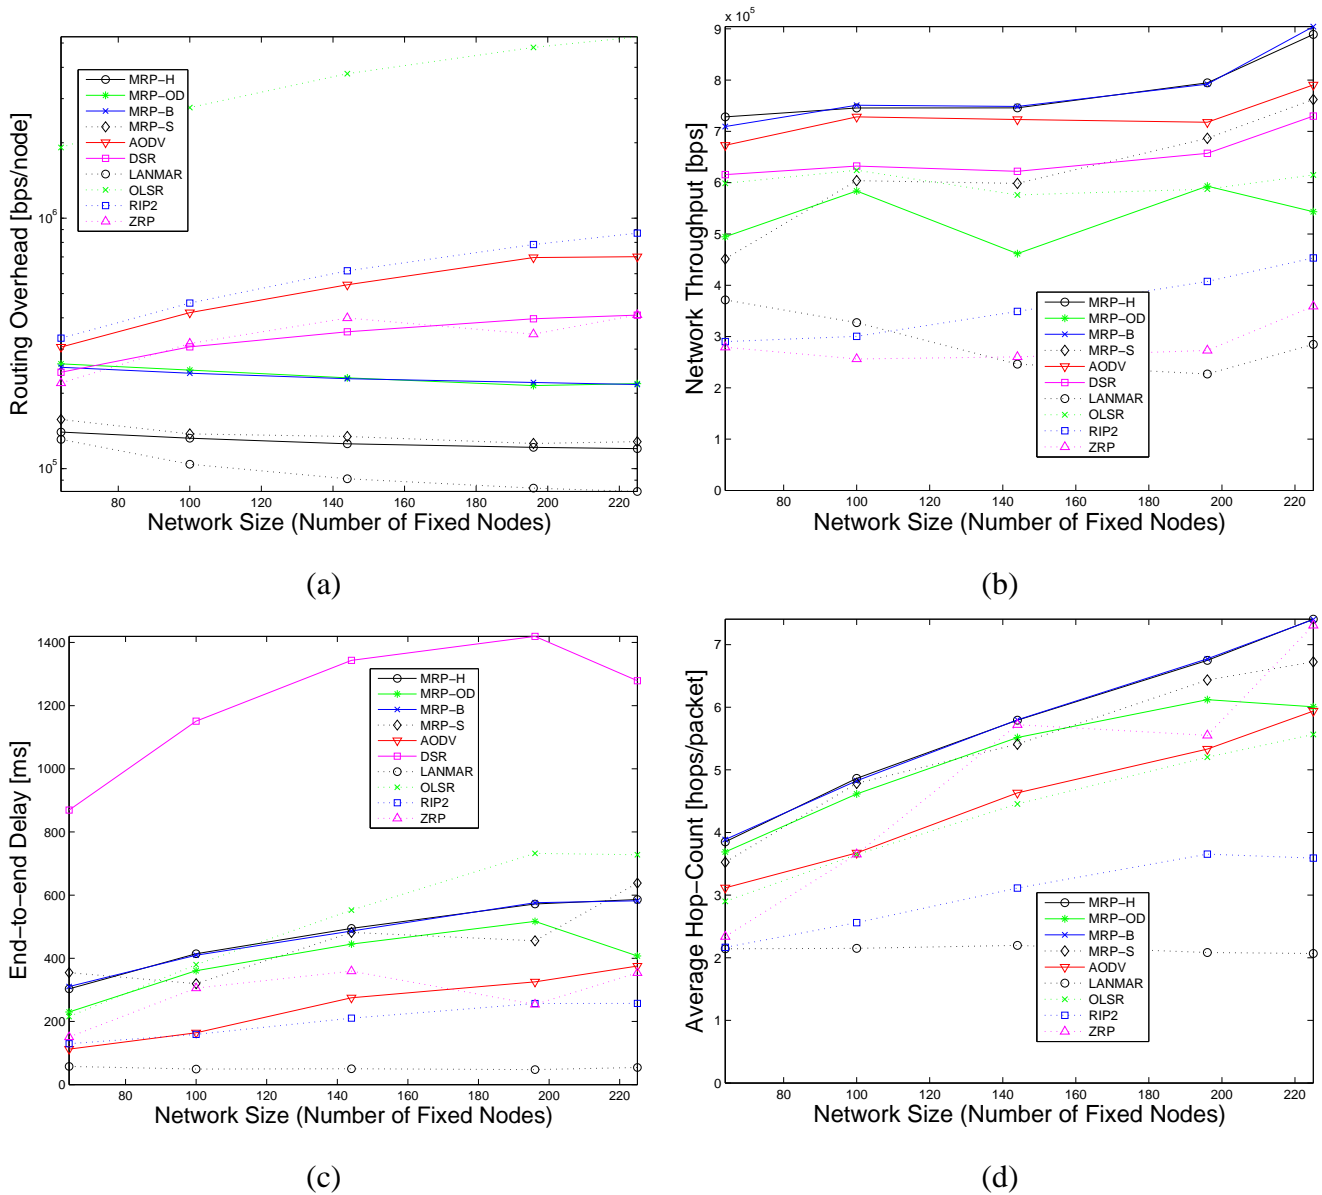


Fig. 13. The performance of the routing protocols for Internet traffic (a) routing overhead, (b) network throughput, (c) end-to-end delay, and (d) average hop-count as a function of network size.

reduces the PDR in Poisson traffic where UDP is used as its transport layer, the majority of the applications in Internet traffic rely on TCP. In response to packet loss and route rediscoveries, TCP congestion control reduces its transmission rate, thus decreasing the throughput.

In Figure 13(c), it is shown that the delay increases for all of the routing protocols as the network diameter increases.

The increase in hop-count of MRP in Fig. 13(d) is correlated to the increase in the network diameter. AODV has a lower hop-count than MRP, as AODV often chooses shorter (but possibly less reliable) routes, while MRP prefers longer routes through stationary nodes.

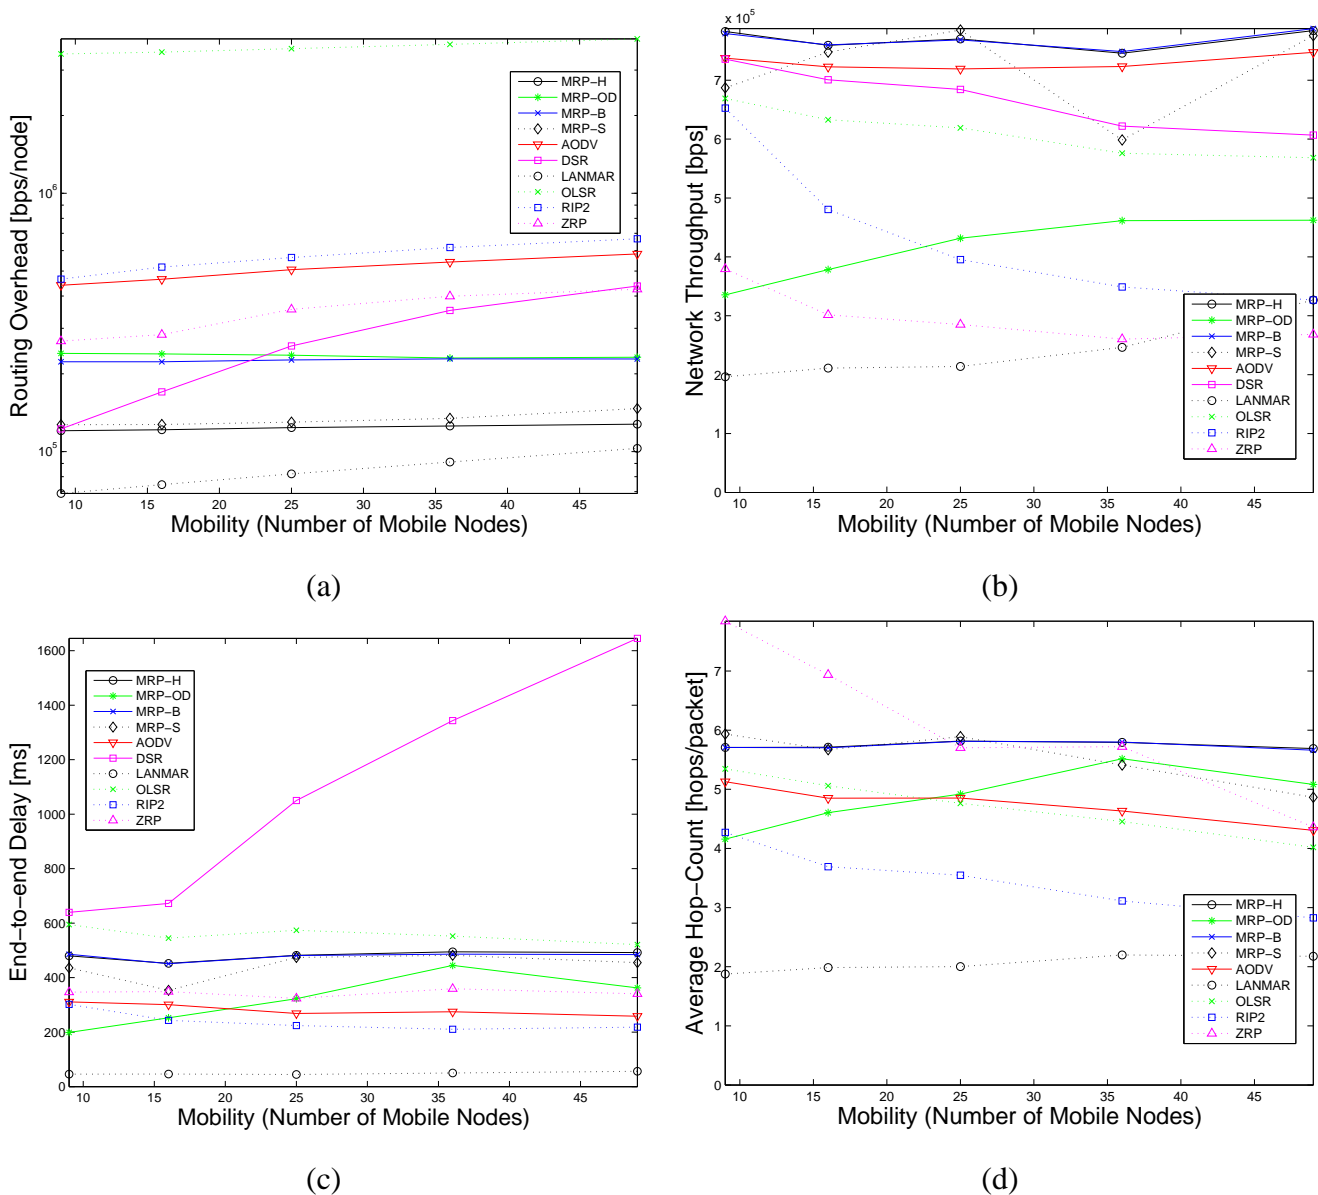


Fig. 14. The performance of the routing protocols for Internet traffic (a) routing overhead, (b) network throughput, (c) end-to-end delay, and (d) average hop-count as a function of mobility.

3) *Number of Mobile Nodes*: The number of mobile nodes is increased in the same manner as in the Poisson case. Figure 14(a) shows that MRP produces a low and almost constant overhead as the number of mobile nodes is increased.

In Fig. 14(b), MRP-H and MRP-B show the highest network throughput as the mobility intensifies. MRP-OD shows relatively low throughput because MRP-OD relies on packet loss to detect a broken link; multiple lost packets drastically affect the TCP's congestion control mechanism (that times-out multiple times before a route is rediscovered). Table-driven protocols such as RIP2, ZRP (ZRP is table-driven for intra-zone routing), and OLSR show decrease in the network throughput.

In Fig. 14(c), MRP shows higher delay than AODV. This result should be considered with AODV's

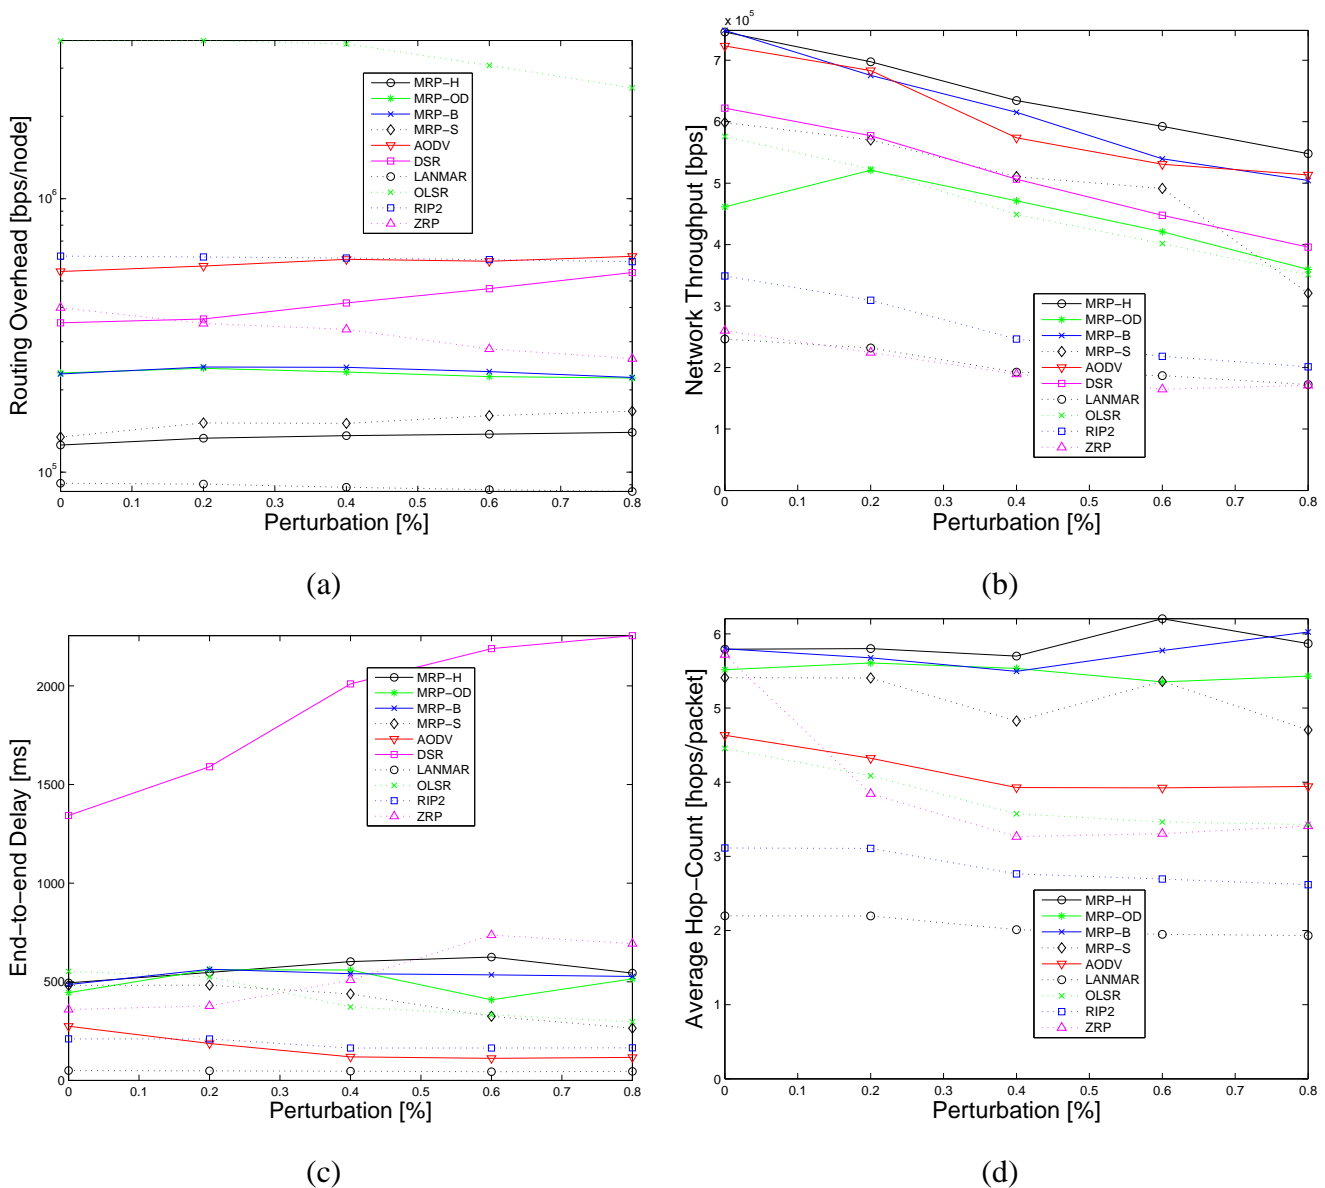


Fig. 15. The performance of the routing protocols for Internet traffic (a) routing overhead, (b) network throughput, (c) end-to-end delay, and (d) average hop-count as a function of the perturbation in stationary nodes grid deployment.

lower throughput and shorter hop-count in Figures 14(b) and (d), which indicates that AODV chooses shorter routes, but loses more packets than MRP due to less reliable routes. DSR shows higher delay than other protocols.

4) *Perturbation:* The deployment of the stationary nodes is perturbed in the same manner as in the Poisson case. The perturbation is increased from zero to 80%.

In Fig. 10(a), MRP shows almost constant overhead. Interestingly, for Internet traffic, the control overhead of AODV and DSR does not increase as in the case of Poisson traffic. As perturbation is increased, the overhead of OLSR decreases for the same reason as for the Poisson case (Section IV-C.4).

Figures 10(b) and (c) show decrease in the network throughput and increase in delay for all protocols

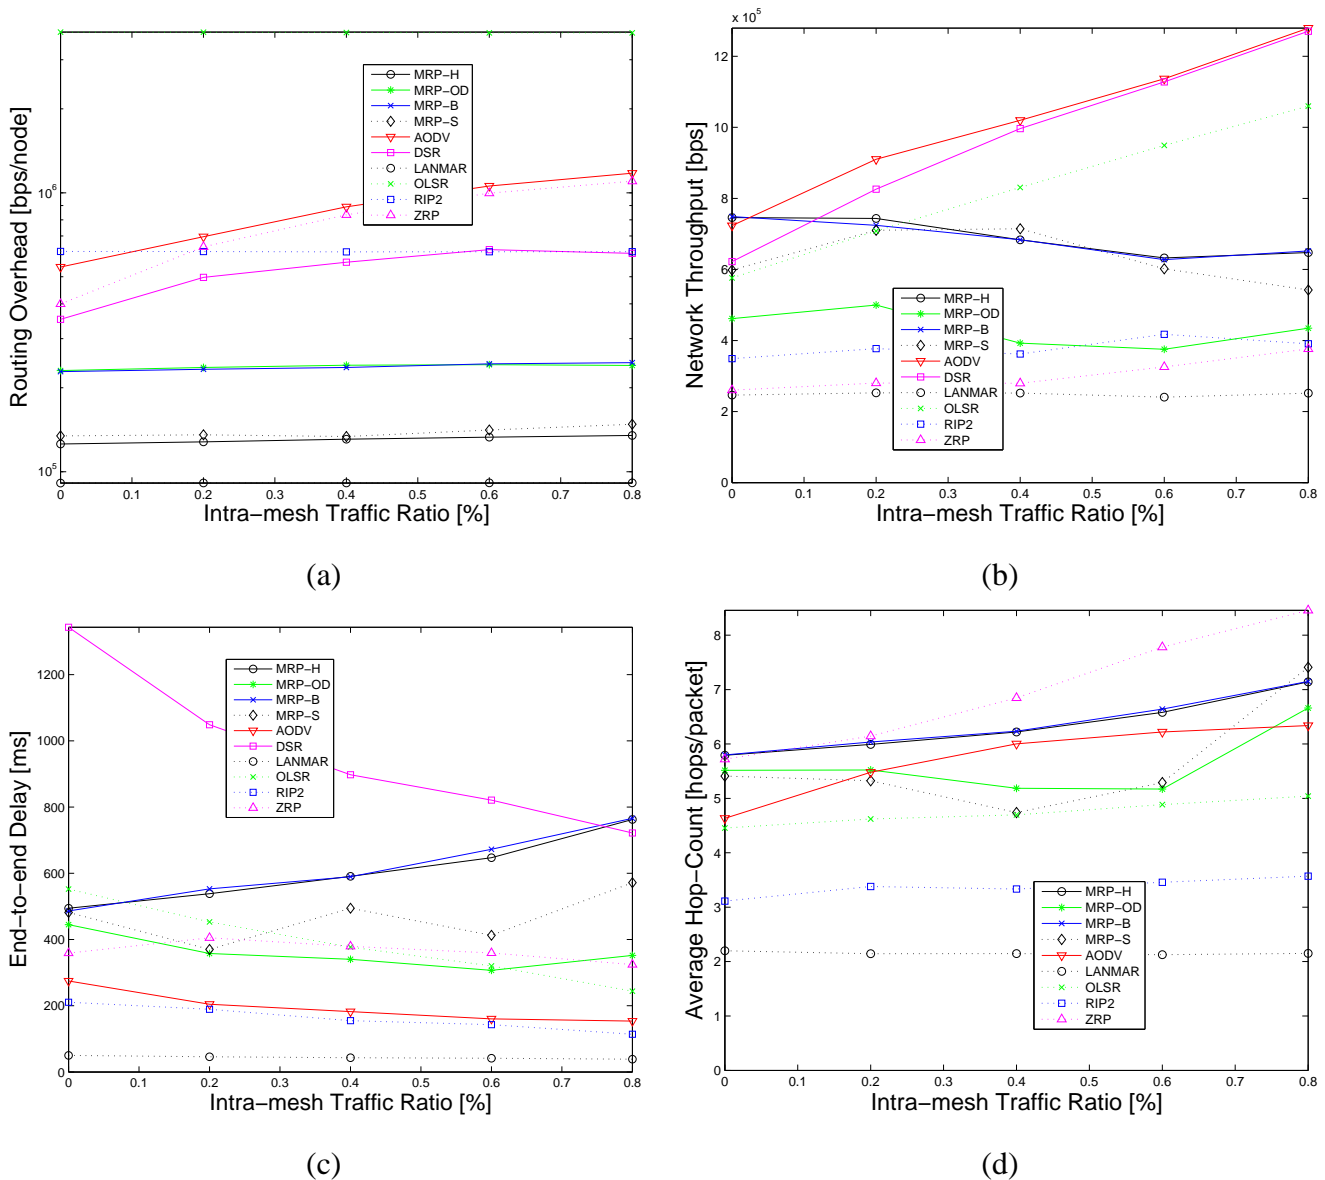


Fig. 16. The performance of the routing protocols for Internet traffic (a) routing overhead, (b) network throughput, (c) end-to-end delay, and (d) average hop-count as a function of the intra-mesh traffic ratio.

as grid is perturbed. MRP shows better throughput and reasonable delay performance compared to other protocols. Among three versions of MRP, MRP-H performs best.

Figure 10(d) shows that the hop-count of MRP slightly increases, while that of all others decreases. Again, this is the effect of the route selection procedure that prefers stable routes to lower hop-counts.

5) *Intra-mesh Traffic Ratio*: The intra-mesh traffic ratio is varied in the same manner as in the Poisson case. Intra-mesh traffic ratio is increased from zero to 0.8. The results are shown in Fig. 11.

In Fig. 11(a), on-demand protocols such as AODV, DSR, and ZRP (ZRP uses on-demand approach for inter-zone routing) show an increased overhead because these protocols have to perform route discovery for destinations located farther away. As expected, MRP and table-driven protocols show almost constant

overhead.

Figure 11(b) shows increased network throughput for AODV, DSR, OLSR, and ZRP, while MRP shows a decrease. This result is expected because the former protocols are designed for generic ad-hoc networks with destinations at arbitrary locations in the network. In contrast, routes in MRP always involve the common parent of the source and the destination. The resulting suboptimal routes are also reflected in the delay and hop-count performance shown in Figures 11(c) and (d), respectively.

The simulation results show that MRP performs well in WMNs that have the most user traffic to and from the gateway (and intra-mesh traffic is a small percentage of the overall traffic).

In this section, we showed that *for WMNs*, in most cases, MRP outperforms the other protocols for the considered metrics. MRP-H shows less overhead, higher PDR (for the Poisson traffic model), higher network throughput (for the Internet traffic model), and lower delay. MRP shows slightly higher hop-count than some of the other protocols as its route selection algorithm prefers stabler rather than smaller hop-count routes. MRP consistently performed better than others when traffic load, network size, mobility, and perturbation were increased. As MRP is optimized for WMNs, when intra-mesh traffic ratio is increased, MRP's advantage withers. Among the MRP versions, MRP-H showed the best performance, followed closely by MRP-B (although with twice the overhead of MRP-H). As expected, MRP-S consistently showed slightly improved delay and hop count at the expense of a decrease in PDR and a slight increase in overhead. MRP-OD showed relatively poor performance compared to the other versions.

The reduction in overhead and the use of stable routes are the two main sources of performance increase that are enabled by assumptions only valid for WMNs. The assumption that most flows are from and to the gateway, enables MRP to form and maintain routes only to the gateway (thus, reducing the overhead). The assumption of a fixed backbone enabled MRP to leverage the stable links of the backbone at a cost of a slight increase in the path length.

V. CONCLUSION

Wireless mesh networks are becoming increasingly popular as they have significant advantages over competing technologies. In this paper, we introduce a new routing protocol specifically designed for those networks. The design of the proposed routing protocol takes advantage of the particularities of WMNs, only maintaining *routing trees* to and from the gateways. Simulation results show that, for WMNs, in most situations, the proposed protocol outperforms existing Internet and MANET routing protocols.

We believe that the performance of the proposed protocol can be further improved if multiple routes are used (e.g., incoming and outgoing flows to the same or different gateways). We will explore this option in future work.

REFERENCES

- [1] M. L. Sichitiu, "Wireless mesh networks: Opportunities and challenge," in *Proc. of the Wireless World Congress*, (Palo Alto, Ca), May 2005.
- [2] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey," *Computer Networks*, 2005.
- [3] R. Bruno, M. Conti, and E. Gregori, "Mesh networks: Commodity multihop ad hoc networks," *IEEE Communications Magazine*, vol. 43, pp. 123–131, Mar. 2005.
- [4] "Belair Networks website." <http://www.belairnetworks.com>.
- [5] "Firetide website." <http://www.firetide.com>.
- [6] "Intel Wireless Mesh Networks website." <http://www.intel.com/update/contents/nc11032.htm>.
- [7] "Kiyon website." <http://www.kiyon.com>.
- [8] "Mesh Dynamics website." <http://www.meshdynamics.com>.
- [9] "Microsoft's Wireless Mesh Networks website." <http://research.microsoft.com/mesh/>.
- [10] "MeshNetworks website." <http://www.meshnetworks.com>.
- [11] "Nokia RoofTop website." <http://www.nwr.nokia.com>.
- [12] "Nortel Networks website." <http://www.nortelnetworks.com/solutions/wrlsmesh/>.
- [13] "PacketHop website." <http://www.packethop.com>.
- [14] "Radiant Networks website." <http://www.radiantnetworks.com>.
- [15] "SkyPilot network website." <http://www.skypilot.com>.
- [16] "Strix systems website." <http://www.strixsystems.com/>.
- [17] "Telabria website." <http://www.telabria.com/>.
- [18] "Tropos networks website." <http://www.troposnetworks.com/>.
- [19] D. Beyer, "Fundamental characteristics and benefits of wireless routing ("mesh") networks," in *Proc. of the International Technical Symposium of the Wireless Communications Association*, (San Jose, CA), Jan. 2002.
- [20] J. Jun and M. L. Sichitiu, "The nominal capacity of wireless mesh networks," *IEEE Wireless Communications Magazine, Special Issue on: Merging IP and Wireless Networks*, Oct 2003.
- [21] I. Chlamtac, M. Conti, and J. J.-N. Liu, "Mobile ad hoc networking: imperatives and challenges," *Ad hoc Networks*, vol. 1, no. 1, pp. 13–64, 2003.
- [22] E. Royer and C. Toh, "A review of current routing protocols for ad-hoc mobile wireless networks," *IEEE Personal Communications*, Apr. 1999.
- [23] J. Raju and J. Garcia-Luna-Aceves, "A comparison of on-demand and table-driven routing for ad-hoc wireless networks," in *Proc. of IEEE ICC*, June 2000.
- [24] S. R. Das, C. E. Perkins, and E. E. Royer, "Performance comparison of two on-demand routing protocols for ad hoc networks," in *Proc. of INFOCOM*, pp. 3–12, 2000.
- [25] C. Bhagwat, "Highly dynamic destination-sequenced distance vector routing (DSDV) for mobile computers," in *Proc. of ACM SIGCOMM*, pp. 234–244, Sept. 1994.
- [26] S. Murthy and J. Garcia-Luna-Aceves, "A routing protocol for packet radio networks," in *Proc. of Mobicom*, Nov. 1995.
- [27] C. Chiang, H. K. Wu, W. Liu, and M. Gerla, "Routing in clustered multihop mobile wireless networks with fading channel," in *Proc. of IEEE Singapore International Conference on Networks*, 1997.
- [28] C. Perkins, "Ad-hoc on-demand distance vector routing," in *Proc. of MILCOM*, Nov. 1997.
- [29] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing* (Imielinski and Korth, eds.), vol. 353, Kluwer Academic Publishers, 1996.
- [30] V. Park, M. Scott, and Corson, "A highly adaptive distributed routing algorithm for mobile wireless networks," in *Proc. of IEEE INFOCOM*, 1997.

- [31] R. Dube, C. Rais, K. Wang, and S. Tripathi, "Signal stability based adaptive routing (SSA) for ad hoc mobile networks," *IEEE Personal Communication*, Feb. 1997.
- [32] Y. B. Ko and N. H. Vaidya, "Location aided routing (LAR) in mobile ad hoc networks," *Wireless Networks*, vol. 6, pp. 307–321, Sept. 2000.
- [33] S. Basagni, I. Chlamtac, V. Syrotiuk, and B. Woodward, "A distance routing effect algorithm for mobility (DREAM)," in *Proc. of ACM Mobicom'98*, (Dallas, TX), pp. 76–84, Oct. 1998.
- [34] J. Li, J. Jannotti, D. S. J. DeCouto, D. R. Karger, and R. Morris, "A scalable location service for geographic ad-hoc routing," in *Proc. of ACM Mobile Communications Conference*, (Boston, MA), Aug. 2000.
- [35] K. Amouris, S. Papavassiliou, and M. Li, "A position-based multi-zone routing protocol for wide area mobile ad-hoc networks," in *Proc. of IEEE Vehicular Technology Conference*, (Houston, TX), pp. 1365–1369, 1999.
- [36] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: a scalable and robust communication paradigm for sensor networks," in *Mobile Computing and Networking*, pp. 56–67, 2000.
- [37] H. Xiaoyan, X. Kaixin, and M. Gerla, "Scalable routing protocols for mobile ad hoc networks," *IEEE Network*, vol. 16, July-Aug 2002.
- [38] M. G. G. Pei and T. Chen, "Fisheye state routing: A routing scheme for ad hoc wireless networks," in *Proc. of ICC 2000*, (New Orleans, LA), June 2000.
- [39] W. List and N. Vaidya, "A routing protocol for K-hop networks," in *Proc. of WCNC 2004*, Mar. 2004.
- [40] T. Clausen and P. Jacquet, "Optimized link state routing protocol (OLSR)." RFC 3626, Oct. 2003.
- [41] J. Moy, "OSPF version 2." RFC 2328, Apr. 1998.
- [42] D. Oran, "OSI IS-IS intra-domain routing protocol." RFC 1142, Feb. 1990.
- [43] Y. Rekhter and T. Li, "A border gateway protocol." RFC 1771 (BGP version 4), Mar. 1995.
- [44] C. Hedrick, "Routing information protocol." RFC 1058, June 1988.
- [45] G. Malkin, "RIP version 2 - carrying additional information." RFC 1388, Jan. 1993.
- [46] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc on-demand distance vector (AODV) routing." RFC 3561, July 2003.
- [47] Z. J. Haas and M. R. Pearlman, "The zone routing protocol (zrp) for ad hoc networks," internet-draft, IETF MANET Working Group, November 1997. Expiration: May, 1998.
- [48] Y. Ko and N. H. Vaidya, "GeoTORA: A protocol for geocasting in mobile ad hoc networks," in *Proc. of the 8th International Conference on Network Protocols (ICNP)*, (Osaka), Nov. 2000.
- [49] Y. Ko and N. Vaidya, "Geocasting in mobile ad hoc networks: Location-based multicast algorithms," in *Proc. of WMCSA*, (New Orleans), 1999.
- [50] R. Ogier, F. Templin, and M. Lewis, "Topology dissemination based on reverse-path forwarding (TBRPF)." RFC 3484, Feb. 2004.
- [51] G. Pei, M. Gerla, and X. Hong, "LANMAR: landmark routing for large scale wireless ad hoc networks with group mobility," in *Proc. of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, (Boston, MA), Aug. 2000.
- [52] M. Gerla, X. Hong, and G. Pei, "Landmark routing for large ad hoc wireless networks," in *Proc. of IEEE GLOBECOM 2000*, (San Francisco, CA), 2000.
- [53] D. Beyer, "Accomplishments of the DARPA SURAN program," in *Proc. IEEE MILCOM 90 Conference*, (Monterey, California), Oct. 1990.
- [54] J. Garcia-Luna-Aceves, C. Fullmer, E. Madruga, D. Beyer, and T. Frivold, "Wireless internet gateways (WINGs)," in *Proc. IEEE MILCOM '97*, (Monterey, California), Nov 1997.
- [55] S. Roy and J.J.Garcia-Luna-Aceves, "Node-centric hybrid routing for ad-hoc wireless extensions of the internet," in *Proc. of Globecom*, 2002.
- [56] S. Roy and J.J.Garcia-Luna-Aceves, "Using minimal source trees for on-demand routing in ad hoc networks," in *Proc. IEEE INFOCOM*, (Anchorage, Alaska), Apr. 2001.
- [57] K. Xu and M. Gerla, "A heterogeneous routing protocol based on a new stable clustering scheme," in *Proc. of the Military Communications Conference (MILCOM)*, (Anaheim, CA), Oct. 2002.

- [58] C. Tschudin and R. Gold, "LUNAR: Lightweight underlay network ad-hoc routing," tech. rep., University of Basel, Switzerland, Jan. 2002.
- [59] M. J. Miller, W. D. List, and N. H. Vaidya, "A hybrid network implementation to extend infrastructure reach," tech. rep., Coordinated Science Laboratory, University of Illinois at Urbana-Champaign., Jan. 2003.
- [60] J. Xi and C. Bettstetter, "Wireless multi-hop internet access: Gateway discovery, routing, and addressing," in *Proc. Intern. Conf. on Third Generation Wireless and Beyond (3Gwireless)*, (San Francisco, CA), May, 28–31 2002.
- [61] Y. Yuan, H. Yang, S. H. Y. Wong, S. Lu, and W. Arbaugh, "ROMER: resilient opportunistic mesh routing for wireless mesh networks," in *Proc. of the First IEEE Workshop on Wireless Mesh Networks (WIMESH'05)*, (Santa Clara, CA), Sept. 2005.
- [62] M. Mosko and J. Garcia-Luna-Aceves, "Multipath routing in wireless mesh networks," in *Proc. of the First IEEE Workshop on Wireless Mesh Networks (WIMESH'05)*, (Santa Clara, CA), Sept. 2005.
- [63] K. Ramachandran, M. M. Buddhikot, G. Chandranmenon, S. Miller, K. Almeroth, and E. Belding-Royer, "On the design and implementation of infrastructure mesh networks," in *First IEEE Workshop on Wireless Mesh Networks (WIMESH'05)*, Sept. 2005.
- [64] A. Raniwala and T. Chiueh, "Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network," in *Proc. of IEEE Infocom '05*, 2005.
- [65] J. Tang, G. Xue, and W. Zhang, "Interference-aware topology control and QoS routing in multi-channel wireless mesh networks," in *Proc. of Mobihoc'05*, May 2005.
- [66] A. P. Subramanian, M. M. Buddhikot, and S. Miller, "Interference aware routing in multi-radio wireless mesh networks," in *The second IEEE Workshop on Wireless Mesh Networks (WIMESH'06)*, Sept. 2006.
- [67] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," in *Proc. of SIGCOMM'04*, (Portland, OR), Aug 2004.
- [68] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. of Mobicom'04*, (Philadelphia, PA), sep 2004.
- [69] Y. Yang, J. Wang, and R. Kravets, "Designing routing metrics for mesh networks," in *First IEEE Workshop on Wireless Mesh Networks (WIMESH'05)*, Sept. 2005.
- [70] L. Iannone, R. Khalili, K. Salamatian, and S. Fdida, "Cross-layer routing in wireless mesh networks," in *Proc. of the 1st International Symposium in Wireless Communication Systems*, (Mauritius), Sept. 2004.
- [71] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol," in *Proc. of European Wireless 2002. Next Generation Wireless Networks: Technologies, Protocols, Services and Applications*, (Florence, Italy), pp. 156–162, Feb. 25-28 2002.
- [72] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communication & Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, no. 5, pp. 483–502, 2002.
- [73] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies – INFOCOM 2003*, pp. 1312–1321, Mar. 2003.
- [74] Scalable Network Technologies, Inc., "QualNet User's Manual, version 3.6." <http://www.scalable-networks.com/>, 2003.
- [75] "Glomosim website." <http://pcl.cs.ucla.edu/projects/glomosim/>.
- [76] R. Bagrodia, R. Meyer, M. Takai, C. Yu-An, Z. Xiang, J. Martin, and H. Y. Song, "Parsec: a parallel simulation environment for complex systems," *IEEE Computer*, vol. 31, Oct. 1998.
- [77] M. Takai, J. Martin, and R. Bagrodia, "Effects of wireless physical layer modeling in mobile ad hoc networks," in *Proceedings of MobiHoc 2001*, Oct. 2001.
- [78] J. Hsu, S. Bhatia, M. Takai, R. Bagrodia, and M. J. Acriche, "Performance of mobile ad hoc networking routing protocols in realistic scenarios," in *Proceedings of MILCOM 2003*, Oct. 2003.
- [79] K. Thompson, G. J. Miller, and R. Wilder, "Wide-area internet traffic patterns and characteristics," *IEEE Network*, vol. 11, Nov-Dec 1997.
- [80] "Inria OLSR." <http://menetou.inria.fr/olsr/>.