行政院國家科學委員會專題研究計畫　期中進度報告

適用於 Wireless Mesh Network 繞徑通訊協定及維運系統的設計實作與效能量測(1/2)

計畫類別： 個別型計畫
計畫編號： NSC94-2213-E-009-107-
執行期間： 94 年 08 月 01 日至 95 年 07 月 31 日
執行單位： 國立交通大學資訊工程學系(所)

計畫主持人： 王協源

報告類型： 精簡報告
報告附件： 出席國際會議研究心得報告及發表論文
處理方式： 本計畫可公開查詢

中　華　民　國 95 年 4 月 25 日

Implementing and Evaluating Three Routing Protocols in Dual-Radio-Dual-Mode IEEE 802.11(b) Wireless Mesh Networks

S.Y. Wang, David Tao, C.C. Lin, and M.W. Hsu
shieyuan@csie.nctu.edu.tw
Department of Computer Science
National Chiao Tung University
Hsinchu, Taiwan

Abstract

In a dual-radio-dual-mode IEEE 802.11(b) wireless mesh network (WMN), each mesh access point has two IEEE 802.11(b) interfaces. One interface operates in infrastructure mode to serve IEEE 802.11(b) client devices so that they can readily use the mesh network without any software modification or installation. The other interface operates in ad-hoc mode to forward packets among mesh access points until they reach an Internet gateway or the access point with which the destination client device is associated. Compared with a single-radio WMN, such a WMN offers ease of use and higher network throughputs for client devices.

In this paper, we implement and evaluate the performance of three routing protocols (OSPF, AODV, and STP) operating in dual-radio-dual-mode IEEE 802.11(b) WMNs. Our simulation results show that among the tested protocols, OSPF provides the highest system total througput and results in the largest number of stable connections in such a WMN.

I. Introduction

As wireless networks become popular, users start to demand higher quality of services such as higher bandwidth, greater coverage, and higher survivability. Wireless mesh network (WMN) [1] is an emerging broadband Internet access technology. Due to its potential to meet these demands at low costs, WMN has attracted the interests of many researchers these days.

A WMN is composed of multiple mesh access points (mesh APs) and mesh clients. A mesh client represents a wireless device (e.g., a notebook computer equipped with an IEEE 802.11(a/b/g) [2] interface) by which a user accesses the Internet via the WMN. On the other hand, a mesh AP acts as a router (if the mesh network functions as a layer-3 network) or a switch (if the mesh network functions as a layer-2 network) to wirelessly forward mesh clients' packets among mesh APs until they reach an Internet gateway or the mesh AP with which the destination client is associated. Since only a few mesh APs need to connect to the Internet to act as Internet gateways, the wiring cost for a WMN can be significantly saved.

According to [1], WMNs can be classified according to their types as follows:

- Infrastructure/Backbone WMNs: In an infrastructure WMN, mesh APs form an infrastructure for mesh clients that connect to them. This type of WMN forms a wireless backbone for mesh clients.
- Client WMNs: A client WMN provides peer-to-peer networking services among client devices and is practically the same as a Mobile Ad-hoc Network (MANET).
- Hybrid WMNs: This type of WMN combines infrastructure and client WMNs.

Infrastructure WMNs are the most popular type in recent years. They are designed for large-scale deployments (e.g., a city) and thus performance is a top concern. In a single-radio infrastructure WMN, each AP uses only one interface to receive packets from clients, send packets to clients, and forward clients' packets among APs. Because all APs and clients need to contend for the bandwidth of a single channel, the system total throughput of the WMN is low. To enhance performance, some researchers and companies have proposed to use multi-radio APs in an infrastructure WMN [3], [4], [5].

In a dual-radio-dual-mode infrastructure WMN, each AP uses two interfaces. One interface operates in infrastructure mode to exchange packets with clients while the other interface operates in ad-hoc mode to forward clients' packets among APs. Since the ad-hoc and infrastructure mode interfaces can be set to use different channels, the network capacity of an infrastructure WMN can be greatly improved. Another advantage of this dual-radio-dual-mode approach is that clients can connect to mesh APs using the standard IEEE 802.11 association procedure without any software modification or installation. Because of these advantages, many commercial WMN products adopt this dual-radio-dual-mode architecture. Due to the importance of this architecture, this paper studies the performance of three routing protocols when they operate in such a WMN.

In an infrastructure WMN, APs need to run a routing protocol so that they know how to route clients' packets to an Internet gateway or to the AP with which the destination client is associated. The AP network in a WMN can be viewed as a Mobile Ad-hoc Network (MANET). Since WMNs share several features with MANET, the routing protocols developed for MANET [6] can be applied to WMNs. For example, Microsoft mesh networks [7] and RoofNet of MIT [8] are based on the Dynamic Source Routing (DSR) protocol [9], [10] and several commercial WMN products adopt the Ad hoc On-demand Distance Vector Routing (AODV) protocol [11] as their underlying routing protocols.

Although several sophisticated routing protocols developed for MANET can be used in WMNs, it is not clear whether they would perform better than the routing protocols developed for fixed Internet. WMNs are not exactly the same as MANET. In an infrastructure WMN, APs are fixed and only clients may move. When a client moves and changes its associated AP, the client location database in the WMN can be updated with the new AP. With this database, when a client wants to send packets to another client, it can look up the database to find the current AP with which the destination client is associated and sends packets to that AP. The routing paths among APs need not be affected by client movements. For this reason, the routing protocols developed for fixed networks may be already good enough for the AP network of an infrastructure WMN. MANET routing protocols generally assume that all network nodes are mobile. As such, most protocols aggressively broadcast control packets to quickly detect link breakage caused by node movements. However, in an infrastructure WMN where APs are fixed, the bandwidth consumed by these control packets will be wasted. For these concerns, some WMNs adopt the routing protocols developed for fixed networks. For example, Tropos Networks Inc. uses predictive wireless routing protocol (PWRP), which is analogous to OSPF [12], to support a WMN [13].

In the literature, the performances of routing protocols operating in dual-radio-dual-mode WMNs are rarely reported. In the paper, we use simulations to implement and evaluate the performance of three routing protocols that have been adopted by some commercial WMNs. These protocols are OSPF, AODV, and SPT (Spanning Tree Protocol) [14], respectively. By comparing their performances under various conditions, we reveal the advantages and disadvantages of these routing protocols when they operate in such a WMN. To mitigate performance bottleneck around the Internet gateway of a WMN, we implemented a multi-gateway WMN and studied the effectiveness of using multiple gateways on the system total throughput. We also implemented the expected transmission count (ETX) metric [15], which is used in RoofNet [8], in OSPF and studied how this metric can help OSPF in a WMN.

The rest of this paper is organized as follows: Section II discusses related work. Section III describes the routing protocols that we studied in the paper. Section IV presents the design and implementation of these routing protocols on the NCTUns 2.0 network simulator [16], [17]. Section V presents the simulation environment and evaluates the performance of these protocols under various conditions. Finally, we conclude the paper in Section VI.

II. Related Work

In a mesh AP, Wireless Mesh Routing (WMR) is the main component that decides how to route packets. Since a WMN is a type of ad hoc networks, the routing protocols developed for ad hoc networks

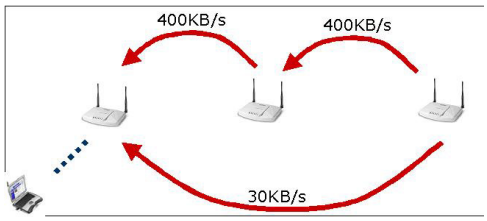

Fig. 1. A longer path (in hops) may provide a higher end-to-end throughput than a shorter path.

can be applied to WMNs. In multi-hop wireless networks, end-to-end throughputs are affected by the hop count of chosen routing paths, signal propagation loss, noise, and interference [18], [19], [20]. In [21], the authors studied the packet losses on a 38-node urban 802.11(b) mesh network and found that links with intermediate levels of losses are common. The performance and scalability of a WMN can be poor if the WMN adopts an inefficient routing protocol. The design of WMR is very important to the overall performance of a WMN.

Many existing ad hoc routing protocols use minimum hop-count as the performance metric to choose a routing path. Prior researches have shown that this kind of shortest-path routing may result in low performance in multi-hop wireless networks [22], [23]. As Fig. 1 shows, sometimes the shortest path (one hop in this example) may provide a lower end-to-end throughput than a longer path (two hops in this example) because the signal quality on the shortest path is bad. To address this problem, researchers have proposed several routing protocols that use different performance metrics.

In [24], the authors proposed the link quality source routing (LQSR) on the basis of DSR [9], [10], and conducted an evaluation of the performance of three link-quality metrics, which are ETX [15], per-hop RTT [25], and per-hop packet pair. They show that generally the ETX metric provides a better performance in stationary cases while the hop-count metric provides a better performance in mobile cases. The authors in [24] show that new metrics are needed for routing protocols to perform better in both mobile and stationary situations.

In order to improve the scalability of multi-hop WMNs or MANET, researchers begin to propose multi-channel wireless networks. The authors in [26], [27] focused on channel assignment problems. The authors in [3] proposed a new performance metric for multi-radio WMNs, which is called Expected Transmission Time/Weighted Cumulative ETT (ETT/WCETT). Other researchers proposed different approaches. The authors in [28] studied interface assignment and routing in multi-channel multi-interface MANET while the authors in [29] proposed a multi-channel routing protocol for multi-channel single-interface networks.

The focus of this paper is different from the focuses of these multi-radio papers. Most of these papers assume that each mesh AP has multiple ad-hoc mode interfaces to forward mesh clients' packets, and focus on how to assign different channels to different interfaces to achieve a better spatial reuse of wireless spectrum. In contrast, although in this paper each AP also has two interfaces (multi-radio), only one interface is set to the ad-hoc mode to forward clients' packets while the other is set to the infrastructure mode to serve normal IEEE 802.11 (b) clients. As discussed before, this dual-radio-dual-mode type of WMN is more practical for deployments. In addition, this paper studied the performance of three routing protocols in dual-radio-dual-mode WMNs, which has been rarely reported in the literature.

III. Types of Routing Protocols Studied

This paper studies the performance of three types of routing protocols proposed for WMNs. They are (1) ad hoc routing protocols, (2) spanning tree protocol, and (3) IP routing protocols, respectively. We briefly present their designs below.

A. Ad hoc Routing Protocols

Currently, there are two types of ad hoc routing protocols:
1. Proactive (table-driven): Proactive protocols (e.g., DSDV [30]) generate route control packets periodically between nodes. Every node needs to maintain a routing entry for every other node. Each time a periodic route control packet is received, a node recomputes the route derived from the control packet and updates its routing table if needed. The drawbacks of proactive routing protocols are listed as follows:
- In the low mobility environment, routes may not be changed over time and thus periodically broadcasting control packets will waste bandwidth.
- Nodes may maintain many routes that will never be used. The messages exchanged for these unused routes only waste network bandwidth.

2. Reactive (on-demand): Reactive protocols (e.g., DSR [9], [10], AODV [11]) trigger the routing path construction only when necessary. To send a packet, a node consults its routing table to find a valid route to the destination of the packet. If a valid route can be found, it sends out the packet. Otherwise, it initiates a route request process for the destination node. When receiving a response for the destination node, the source node generates a valid routing entry for this destination node. The validity of the route is determined by the lifetime specified for it. If the route is not used for some period of time, the route is considered to be no longer needed and is removed from the routing table. Contrary to proactive protocols, reactive protocols maintain routing information only when needing to transmit packets. It reduces unnecessary bandwidth overheads at the cost of spending more time on finding a route.

The Ad-Hoc on Demand Distance Vector (AODV) [11] routing protocol is a representative of the reactive protocols for ad hoc networks. In AODV, the protocol operation is performed based on the demands of packets. If no packet needs to be transmitted, no route will be maintained. The source node initiates the route discovery process by broadcasting the route request (RREQ) only when it tries to send a packet and there is no active route found in its routing table. Except the destination node, each node receiving the RREQ will re-broadcast it. The dissemination of RREQ works in the flooding manner until the destination node is reached. Upon receiving the RREQ, the destination node sends back a route reply (RREP) to the source node through the reverse path of RREQ. There is an alternative way to improve the response time of the route discovery process. If an intermediate node already has the routing information for the destination node, it can send back an RREP to the source node without further broadcasting the RREQ.

Currently, AODV is popular and is adopted by some commercial WMN products. As such, we select it as one of the three routing protocols studied in this paper.

B. Spanning Tree Protocol

The spanning tree protocol (STP) [14] is a loop-prevention method on LANs where multiple bridges (nowadays a bridge is called a switch) may be inter-connected and physical loops may be formed. The leader-election algorithm in STP selects a bridge on a LAN as the root bridge of the spanning tree. Each bridge running STP exchanges its local information in a format called Bridge Protocol Data Unit (BPDU). When the priorities of all bridges combined with their MAC addresses are exchanged over the whole network, the bridge with the highest ID is selected as the root bridge. All ports on the root bridge are known as designated ports. On a link segment, only one attached port can be designated and all others must be blocked. All designated ports are in what is known as the forwarding state. A port in the forwarding state is allowed to send and receive traffic. All of the other bridges are known as non-root bridges and they choose a port known as a root port to send and receive traffic to/from the root bridge. Using this method, redundant ports (links) are closed down and packets will not be endlessly spawned and trapped in loops. A closed port can be opened again if there is a change to the network topology and that port is needed for the new spanning tree.

A traditional fixed switch uses one of its ports to connect to another switch via a cable. In a WMN, however, a mesh AP uses its single ad-hoc mode interface to exchange packets with multiple neighboring APs. Therefore, a port in a WMN should be redefined to be the ad-hoc wireless connection that is used to exchange packets between two mesh APs.

Because the fixed AP network of a WMN functions like a layer-2 LAN with redundant (wireless) links, STP can be applied to these fixed APs to construct a spanning tree (packet forwarding paths) among them without any loop. Since STP can perform self-routing, self-organization, and self-healing, it is fault-tolerant and can cope with node mobility. STP has been widely implemented on switches. Many commercial WLAN AP products have included it as a standard feature so that these APs can readily form an IEEE 802.11 WDS (Wireless Distributed System) [2] when they are set to the bridging mode. Due to these reasons, we select STP as one the three routing protocols studied in this paper.

C. IP Routing Protocols

On the Internet, many routers are being controlled by IP routing protocols. Currently, the two most popular routing protocols controlling the routing of Internet traffic are OSPF [12] and RIP [31].

- Route Information Protocol (RIP) is based on the distance-vector algorithm. Routers broadcast their own routing tables periodically and calculate a shortest path based on the exchanged information to route packets. RIP is a simple routing protocol and is not suitable for large networks because it generates many control messages and thus wastes much network bandwidth.
- Open Shortest Path First (OSPF) is a routing protocol developed for IP networks by IETF. In the mid-1980s, because RIP was increasingly incapable of serving large and heterogeneous networks, OSPF was created to replace RIP. OSPF is a link-state routing protocol that relies on flooding of link-state advertisements (LSA) to all other routers within the same hierarchical area. As an OSPF router accumulates link-state information, it uses the Shortest Path First (SPF) algorithm to calculate the shortest paths to all other routers.

Due to the great success of OSPF on the Internet and the fact that the AP network of a WMN is a fixed network, it is natural for people to think that OSPF may also be able to provide good performances in WMNs. For this reason, some commercial WMN products use OSPF as the routing protocol running among mesh APs. Due to its importance, we include OSPF as one of the three routing protocols studied in this paper.

IV. Routing Protocol Design and Implementation

In this paper, we use simulations to compare the performances of these routing protocols. These protocols are implemented on the NCTUns 2.0 network simulator [16], [17]. NCTUns 2.0 is an extensible network simulator capable of simulating various protocols used in both wired and wireless IP networks. It provides a module-based platform for developers to develop their modules and integrate them into the simulator. By developing and linking modules on this platform, one can create a new device with a desired protocol stack. In this section, we present the detailed implementation of these routing protocols.

A. Two Types of WMN Usage

As described in IEEE 802.11, each mesh client is associated with a mesh AP via its infrastructure mode interface. As a mesh client moves, it may disassociate with the current mesh AP and reassociate with a new mesh AP. The Wireless Mesh Routing (WMR) module in a mesh AP knows which mesh clients are associated with the AP. For each associated mesh client, WMR records the IP and MAC addresses used by it. Two types of WMN usage are presented below:

1. From WMN to Internet:

In this type of usage, a mesh client wants to send packets to a host on the Internet. Example usages include setting up a TCP connection from a mesh client to a web server on the Internet to download

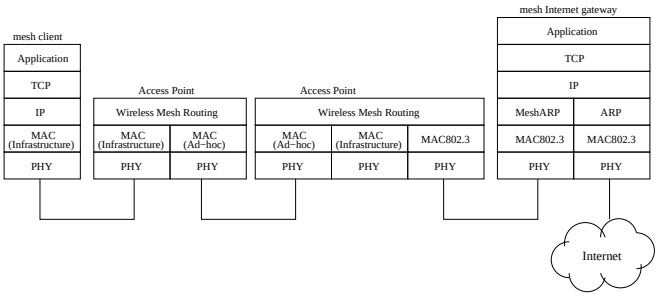

Fig. 2. The protocol stacks of the mesh client, mesh AP, and mesh Internet gateways (WMN-to-Internet)

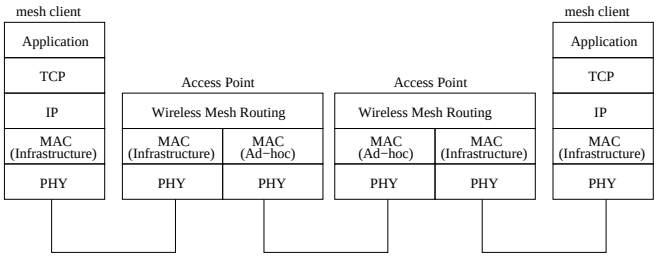

Fig. 3. The protocol stacks of the mesh client and mesh AP (WMN-to-WMN)

a web page. In this usage, the mesh client is given a gateway IP address so that its packets can be first sent to the mesh Internet gateway, which then further forwards the packets to the Internet. Fig.2 shows the protocol stacks of the mesh client, mesh AP, and mesh Internet gateways involved in this type of usage.

2. From WMN to WMN:

In this type of usage, a mesh client wants to send packets to another mesh client in the same WMN. Example usages include setting up a VoIP phone call between two mesh clients. In this usage, the source mesh client is given the IP address of the destination mesh client. Fig.3 shows the protocol stacks of the mesh client and mesh AP involved in this type of usage.

B. Address Resolution Protocol

A source mesh client needs to use the ARP protocol to find the MAC address used by the mesh Internet gateway or another mesh client. To do so, it broadcasts an ARP request via its infrastructure mode interface. This packet will be received by the mesh AP with which the source mesh client is associated. When a mesh AP receives such a packet asking for the IP-MAC address mapping of an interface, it first checks whether the specified IP address is used by one of its associated mesh clients. If this is the case, it broadcasts the ARP request via its infrastructure mode interface to all of its associated mesh clients. When the destination mesh client receives the ARP request, it sends back an ARP reply with its own MAC address via its infrastructure mode interface. Upon receiving the ARP reply, the mesh AP then sends back the ARP reply to the source mesh client via its ad-hoc mode interface. On the other hand, if the mesh AP finds that none of its associated mesh client uses the specified IP address, it broadcasts the ARP request via its ad-hoc mode interface to other neighboring mesh APs to continue the ARP process. Note that if a mesh AP connects to the mesh Internet gateway, it also sends the ARP request to the gateway. With this design, a mesh client can obtain the MAC address used by either the mesh Internet gateway or another mesh client.

C. Routing Procedures

In the following, we present the detailed routing procedures used in the three routing protocols. When a mesh AP runs the AODV, STP, or OSPF routing protocol, the function of its WMR module will implement the specified protocol. To save space, here we only present the routing procedures involved in the "From WMN to WMN" usage type.

1. AODV:

When a mesh AP receives a RREQ (route request) packet, its WMR first checks whether the destination mesh client is associated with the AP. If this is the case, the mesh AP sends the RREQ to the destination mesh client via its infrastructure mode interface. Upon receiving the RREQ, the destination mesh client sends back a RREP to the associated mesh AP, which then sends back the RREP toward the source mesh client via its ad-hoc mode interface. When the RREP travels its way toward the source mesh client through several intermediate mesh APs, the WMR modules running on these mesh APs will create a valid and active routing entry for the destination mesh client in their routing tables. Such a routing entry includes the (the MAC address of the destination mesh client, the MAC address of the ad-hoc mode interface of the next-hop mesh AP) mapping information.

When the RREP comes back to the source mesh client, the source mesh client can start sending its packets. First, the destination MAC address of the packet is filled in with the MAC address of the destination mesh client. Then the packet is sent out via the infrastructure mode interface. The packet will be received by the associated mesh AP via its infrastructure mode interface. The WMR module in the mesh AP checks whether the destination MAC address of the received packet matches the MAC address of any of its associated mesh clients. If this is the case, the packet is sent out through its infrastructure mode interface to reach the specified destination mesh client. Otherwise, the WMR module checks its routing table to find the MAC address of the next-hop mesh AP for this packet. Before forwarding the packet via its ad-hoc mode interface, the mesh AP encapsulates the packet with a tunneling header. It then fills in the next-hop MAC address as the destination MAC address of this tunneling header. The mesh AP then sends out the encapsulated packet via its ad-hoc mode interface, which will be received by other neighboring mesh APs via their ad-hoc mode interfaces. Based on the destination MAC address in the tunneling header, each of these mesh APs determines whether it should (1) drop the packet, (2) send the packet to the destination mesh client should it is associated with the AP, or (3) check its routing table to continue forwarding the packet to another mesh AP. In the second situation, the mesh AP strips off the tunneling header before sending the packet to the destination mesh client. In the third situation, the mesh AP replaces the tunneling header with a new one before forwarding the packet. The above steps are repeated until the packet finally reaches the destination mesh client.

2. STP:

The WMR in each mesh AP runs the spanning tree protocol to form a spanning tree among them. As described before, because in a WMN a mesh AP can connect to multiple neighboring mesh APs, we view the connection between two APs as a switch port and give it a port state. The WMR does not know how many ports it has in advance. Instead, when receiving a STP control packet from a neighboring AP, it checks whether this port (connection) has been created before. If not, it dynamically adds the port into the port list. When the spanning tree protocol has been running for a while among mesh APs, in each mesh AP some ports will be closed while some ports will be opened, and a spanning tree will be generated among mesh APs. From now on, mesh APs will forward packets only on this spanning tree to avoid the looping problem.

When a mesh client wants to send packets to another mesh client, it sends out the packets via its infrastructure mode interface. Upon receiving it, the associated mesh AP encapsulates the packet with a tunneling header and sends it out via an opened port on its ad-hoc mode interface. This port is chosen based on the result of running the learning-bridge protocol [32] among mesh APs. By default, every mesh AP runs the learning-bridge protocol. When a mesh AP does not know to which port to

forward a packet, it sends a copy of the packet to each of its ports except the one from which the packet is received. When a mesh AP receives a packet from a neighboring AP, it records the neighboring mesh AP as the next-hop mesh AP for the mesh client specified in the source MAC address of the packet. That is, it learns and installs routing entries when forwarding packets. The mesh AP then strips off the tunneling header from the packet. If it finds that the destination MAC address of the packet matches the MAC address of one of its associated clients, it sends out the packet via its infrastructure mode interface to the destination mesh client. When the packet is eventually received by the destination mesh client, the destination mesh client may send a reply packet back to the source mesh client via its infrastructure mode interface. Since the intermediate mesh APs have learned and installed the routing entries for the source mesh client, this reply packet will be forwarded hop-by-hop without flooding. When the reply packet is on its way back to the source mesh client, these intermediate mesh APs learn and install routing entries for the destination mesh client. Therefore, after a handshake between the source and destination mesh clients, the uncast routing path between them on the spanning tree can be directly used and blind-flooding is no longer needed.

3. OSPF:

The WMR in each mesh AP runs the OSPF protocol to build shortest routing paths to all other mesh APs. Each mesh AP periodically sends hello packets to its neighboring mesh APs so that network topology changes can be detected. If there is a topology change, an involved mesh AP will flood its updated LSAs to the whole network and all other mesh APs will use the LSAs to update their shortest path trees and routing tables. A mesh AP has two kinds of neighbors: mesh APs and mesh clients. A neighbor addition or removal of any kind will trigger the mesh AP to flood an updated LSA to the network. A mesh client does not send hello packets and it is viewed as a neighbor by only the mesh AP with which it is associated. If a mesh AP receives a LSA from a neighboring mesh AP and finds that this LSA contains a mesh client that is associated with it, it assumes that the mesh client has moved to the coverage area of the new mesh AP and removes the mesh client from its neighbor list. This neighbor removal operation will trigger the mesh AP to flood a LSA to the network to announce this topology change.

When a mesh AP floods a LSA, the MAC addresses of the mesh clients associated with it are included in the LSA. As such, the MAC address of every mesh client in a WMN is propagated throughout the network and known by every mesh AP. Since each mesh AP has a routing entry for every mesh client, a mesh client's packets can be forwarded hop-by-hop toward the mesh AP with which the destination mesh client is associated. To preserve the original source and destination MAC addresses, as in the AODV and STP approaches, intermediate mesh APs use a tunneling header to specify the next-hop mesh AP for a packet.

D. Multi-Gateway Wireless Mesh Networks

A WMN is commonly used as an Internet access network by which mesh clients send/receive information to/from the Internet. In a WMN, if only one mesh AP can connect to the mesh Internet gateway, all traffic flows will need to merge at that mesh AP. This will limit the aggregate throughput of all traffic flows in a WMN to only the bandwidth of an IEEE 802.11(b) interface.

To address this problem, we designed and implemented a multi-gateway WMN, where multiple mesh APs can connect to a mesh Internet gateway to share Internet traffic load among them. (These mesh APs are called the "gateway mesh AP" below.) Each of these gateway mesh APs uses a high-speed (e.g., 100 Mbps or Gbps) link to connect to the mesh Internet gateway to ensure that the used link is not be a performance bottleneck for the WMN. In our design, as Fig. 4 shows, when mesh clients broadcast their ARP requests to ask the MAC address of the gateway (they all use the same gateway IP address, which is 1.0.1.1 in this example), different mesh clients may get different gateway MAC addresses. By this design, Internet traffic generated by mesh clients can be directed to different gateway mesh APs without overloading a single gateway mesh AP. This will make a WMN more

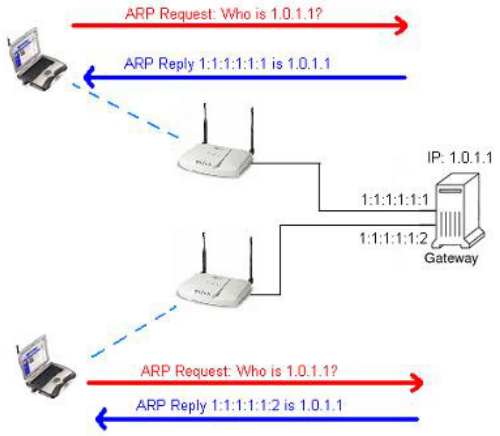

Fig. 4. The ARP request and reply procedure in a multi-gateway WMN

scalable with the number of mesh clients.

The internal design is presented as follows. After a mesh client sends out its ARP request, mesh APs flood the ARP request toward the mesh Internet gateway. Since the ARP request is a broadcast packet, it may be cloned by mesh APs and multiple copies of it may reach the mesh Internet gateway. To avoid wasting bandwidth, mesh APs use the sequence number carried in the ARP request to detect and avoid sending redundant packets. As such, for each ARP request initiated by a mesh client, only one copy of the ARP request will reach each of the interfaces of the mesh Internet gateway. When receiving an ARP request from an interface, the gateway performs two operations: (1) First, it puts the MAC address of this interface in the ARP reply and uses the same interface to send back the ARP reply. (2) Second, it creates a (mesh client source MAC address, interface ID) entry in its routing table. Fig. 5 shows the protocol stacks of gateway mesh APs and the mesh Internet gateway. The MeshARP module performs the two operations described above.

Executing the first operation enables mesh clients to use their nearby gateway mesh APs to send their traffic to the Internet. This allows the load of outoging Internet traffic to be balanced among different gateway mesh APs. As discussed previously, the same ARP request may reach the gateway more than once (one request from each of its interfaces) and this will cause the gateway to send back a reply for each of them. However, this behavior results in no problems. When the source mesh client receives multiple ARP replies each carrying a different gateway MAC address, it can simply use the MAC address carried in the first reply as the MAC address of the gateway. Normally, such a reply carries the MAC address of the gateway interface that has a lighter load, which is exactly the interface that should be used for load-balancing purposes. On the other hand, executing the second operation enables incoming Internet traffic to be routed to different mesh clients via different gateway mesh APs. This allows the load of incoming Internet traffic to be balanced among different gateway mesh APs.

E. OSPF with ETX support

ETX stands for the expected transmission count metric. It was designed for RoofNet [8] at MIT to enhance the performance of the ad hoc routing protocol. The ETX metric is used to help a routing protocol choose a routing path with a better end-to-end throughput. In the RoofNet, DSR is modified to work with the ETX metric and it is shown that using the ETX metric can improve the end-to-end throughput between a pair of nodes. To see how ETX can improve the OSPF routing protocol in WMNs, we implemented a version of OSPF with ETX support.

To support the ETX metric, each mesh AP counts the number of hello packets it receives during a period and uses it to calculate the delivery ratio of each neighbor. The mesh AP stores these delivery

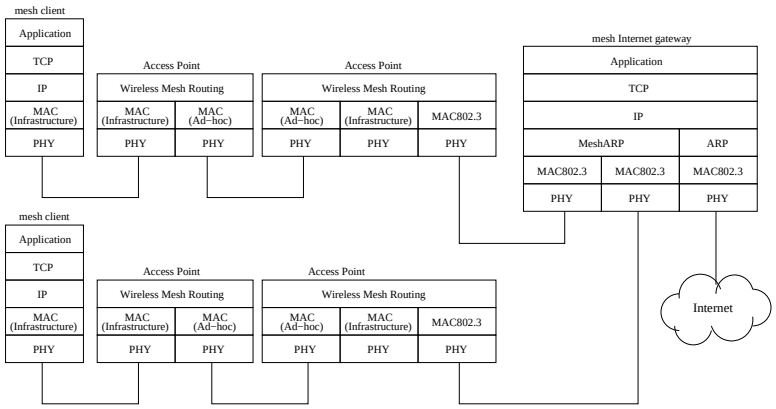

Fig. 5. The internal protocol stack of a multi-gateway WMN

ratios in its LSA packets to inform other mesh APs of the delivery ratios of its neighbors. These delivery ratios are used as link weights when a mesh AP computes the shortest paths to all other mesh APs. By this method, a smallest ETX path tree can be constructed, which can be used to find a high-throughput routing path between a pair of mesh clients.

V. Performance Evaluation

In this section, we evaluate the performance of WMNs using simulations. All simulations are performed on the NCTUns 2.0 network simulator. Each case is simulated 20 times with different random mesh client locations and their average is reported. The total simulation time for each case is 200 seconds, but we just took the last 100 seconds to avoid the influence of the startup behavior of TCP traffic flows. As Fig. 6 shows, 25 mesh APs are deployed in the field. They are placed in a 5x5 grid and the distance between two vertical/horizontal neighboring mesh APs is set to be 200 meters.

Each mesh AP has two IEEE 802.11(b) interfaces. One operates in the ad-hoc mode for forwarding packets among mesh APs while the other operates in the infrastructure mode for serving mesh clients. The transmission and interference ranges of these wireless interfaces are set to be 250/550 meters, respectively, which are used in the ns-2 network simulator [33] and commonly used by researchers. These interfaces use different channels to avoid interferences. During simulation, OSPF, STP, or AODV is run among mesh APs. The mesh AP at the center of the field connects to a mesh Internet gateway using a 100 Mbps link. There are 25 mesh clients scattered at random locations in the field. Each mesh client has an infrastructure mode IEEE 802.11(b) interface and uses it to send/receive packets to/from its associated mesh AP. To make Fig. 6 easy to read, a mesh client is drawn at a location close to a mesh AP. However, in simulations since mesh clients are randomly placed in the field, the client-AP association relationship is not always one-to-one as shown in the figure. Instead, multiple mesh clients may be associated with a mesh AP and a mesh AP may have no mesh client associated with it.

A. One-to-Multi Downlink Traffic Configuration

Nowadays, many Internet applications such as FTP, HTTP, email, etc. rely on Transmission Control Protocol (TCP) to reliably transport data across heterogeneous networks. Internet users usually download data from the Internet through an Internet gateway. Thus, we study a downlink TCP traffic case here. There is a TCP receiver (rtcp) running on each mesh client. Twenty five TCP senders (stcp) run on the mesh Internet gateway and each one greedily sends traffic to the TCP receiver running on each mesh client. In total, there are 25 greedy TCP traffic flows competing for the bandwidth of the WMN. In this case, all mesh clients are fixed.

Fig. 7 shows the system total throughput of the three protocols in the WMN. We see that these

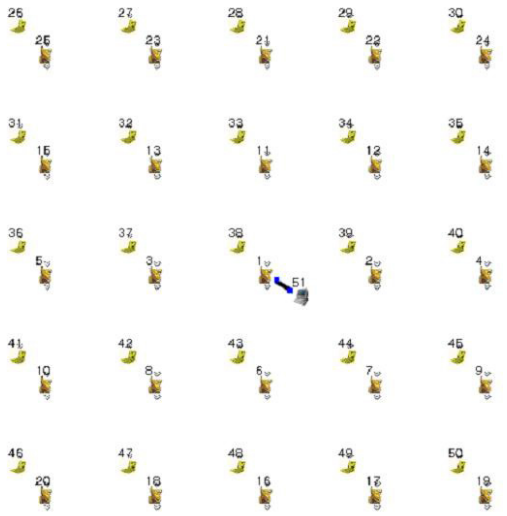

Fig. 6. The network topology of the 5x5 grid WMN

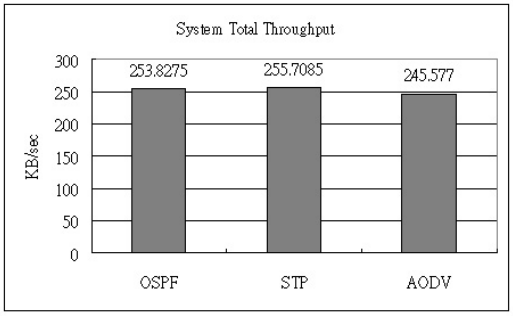

Fig. 7. The system total throughput in the one-to-multi downlink traffic case

routing protocols provide about the same system total throughput. However, Fig. 8 shows that although the gateway can establish a TCP connection to every mesh client, AODV results in the least number of stable connections in the WMN. In our definition, if the TCP receiver of a connection receives no data (0 KB/sec) for more than one half of the simulated period, the connection is viewed as a unstable connection; otherwise, it is viewed as a stable connection. A unstable TCP connection is the result of excessively triggering TCP congestion control on the connection, which prevents the TCP sender from sending out data for a long period of time. Due to the design of TCP congestion control, a packet loss will trigger TCP congestion control and many packet losses may result in a long transmission timeout.

To answer why AODV results in the least number of stable TCP connections in a WMN, we studied its protocol design and the effect of its parameters. We found that in a highly-utilized WMN, packet collisions happen quite frequently and in this situation TCP connections may constantly timeout for a long period of time. According to the design of AODV, if there is no traffic flowing on an established AODV routing path for a period of time (which is specified by the ACTIVE_ROUTE_TIMEOUT parameter), the source node will abandon the current path and re-flood the RREQ across the network to set up a new routing path. For this reason, when the TCP connection that uses the AODV routing path times out for a long period of time, AODV will abandon the used routing path and try to find a new one for the TCP connection. Flooding RREQ, however, consumes much network bandwidth and results in more packet collisions, which may cause the new routing path to be never set up. This will cause the TCP connection to timeout for even a longer period of time and eventualy make the TCP

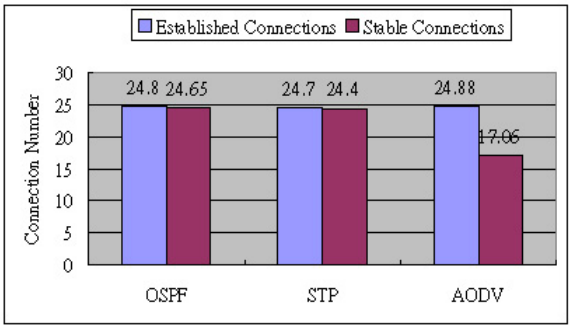

Fig. 8. The number of established and stable connections in the one-to-multi downlink traffic case

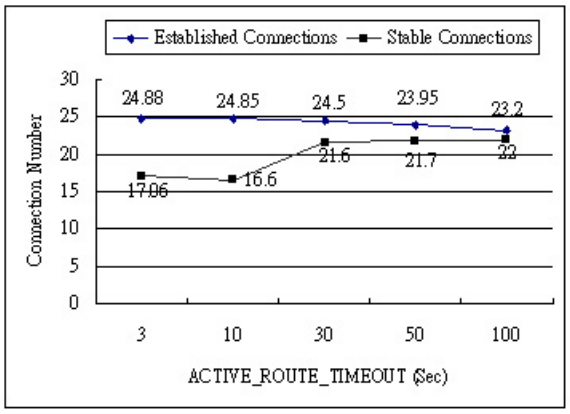

Fig. 9. The number of established and stable connections in AODV under different ACTIVE_ROUTE_TIMEOUT values

connection a unstable connection. The default value for this parameter is 3 seconds in AODV. We found that if we increase this value to a higher value such as 100 seconds, this problem can be mitigated and AODV can result in more stable connections in the WMN. Fig. 9 shows that the number of stable connections in the WMN increases as the ACTIVE_ROUTE_TIMEOUT value increases. Although increasing the value of this parameter can result in more stable connections, it also causes AODV to respond more slowly to node movements. As will be presented later, this will result in a lower number of stable connections when mesh clients move in a WMN.

Fig. 10 shows the relationship between the hop count of the established connections and their achieved throughput in the OSPF routing protocol. We can see that as the average hop count of a connection decreases, the achieved throughput of the connection increases. This phenomenon shows that in a WMN "short" TCP connections usually can achieve more bandwidth than "long" TCP connections. We also studied the 20 runs of the OSPF case and found that if in a run there are more connections with fewer hop counts, the system total throughput of the run will be higher. Fig. 11 shows the results. This phenomenon shows that using only the system total throughput as the sole performance metric to judge which routing protocol performs best is somewhat misleading. A high system total throughput can be easily achieved by letting several "short" TCP connection monopolize the system bandwidth of a WMN. The number of stable connections that can coexist to share the bandwidth of a WMN is an important performance metric that should also be considered.

B. Multi-to-Multi Peer Traffic Configuration

In recent years, peer-to-peer applications are becoming more and more popular. Some examples include VoIP and music/video download applications. Here we make a case to represent the usage of

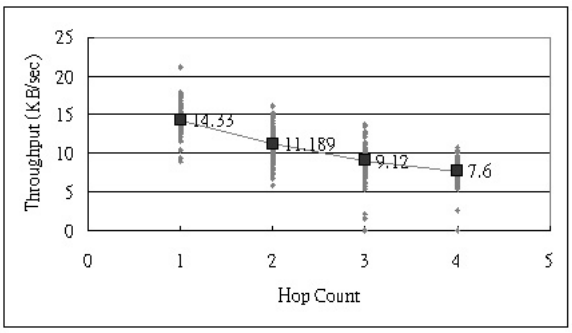

Fig. 10. The relationship between the average hop count and achieved throughput of connections

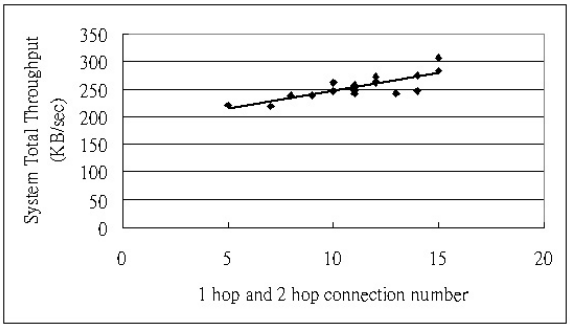

Fig. 11. The relationship between the number of "short" connections and the system total throughput

peer-to-peer applications. In this case, a TCP receiver (rtcp) and a TCP sender (stcp) run on each mesh client and each mesh client randomly sets up a greedy TCP connection to another mesh client. In total, there are 25 greedy TCP traffic flows in the system competing for the system bandwidth of the WMN. All mesh clients are fixed in this case.

Fig. 12 shows the system total throughput of the three routing protocols under this traffic pattern. We can see that the achieved system total throughputs are higher than those reported in Fig. 7. Note that in the previous downlink TCP traffic case, all TCP traffic flows need to merge at the single gateway. As such, they are bottlenecked at a single point in the WMN. In this multi-to-multi peer traffic case, however, all TCP traffic flows need not merge at the single gateway. Instead, they can choose the best shortest routing paths in the WMN for transporting their data. Clearly, this freedom will improve the efficiency of wireless bandwidth usage in the WMN. Fig. 12 shows that OSPF provides a higher system total throughput than AODV and STP. The reason why OSPF outperforms STP is that in STP a routing path between two mesh clients may not be the shortest one because it must reside on the spanning tree. As such, wireless bandwidth of the WMN is not efficiently utilized. The reason why OSPF outperforms AODV has been explained before. It is because in AODV the active route timeout of a routing path is constantly triggered, which causes the RREQ to be flooded to the network, wasting the wireless bandwidth of the WMN.

Fig. 13 shows the number of established and stable connections in these routing protocols. We can see that OSPF results in more stable connections than AODV and STP. The reason for AODV has been explained above. Here we explain the reason for STP. It is clear that, because packets can only be transported on the spanning tree, STP may waste wireless bandwidth due to the use of non-shortest-path route between a pair of mesh clients. For example, our simulation results show that on average a packet needs to traverse 3.99 hops to reach its destination client in STP while this number can be reduced to only 3.45 hops in OSPF. Due to this reason, given the same level of client traffic load, the level of congestion in STP is more severe than in OSPF. This means that more packets will be dropped

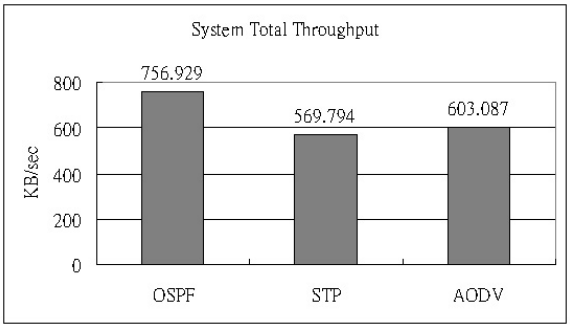

Fig. 12. The system total throughput in the multi-to-multi peer traffic case

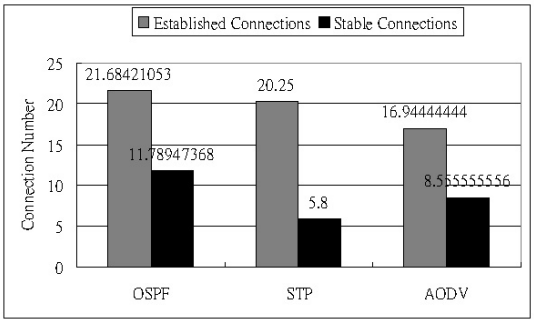

Fig. 13. The number of established and stable connections in the multi-to-multi peer traffic case

in STP, which include the control packets of STP. As a result, the spanning tree constructed in STP may need to be constantly repaired or changed. However, this will cause TCP connections to timeout more often and make them unstable connections.

C. Mobility Condition

In this case, we study the performance of the three routing protocols when mesh clients move. The settings are the same as those in the downlink TCP traffic case except that now all mesh clients move at 1 m/sec speed based on the random-waypoint mobility model. Fig. 14 shows the system total throughput of the three routing protocols in this mobility condition. We can see that they provide almost the same performance. Fig. 15 shows the number of established and stable connections in these three routing protocols. We see that AODV results in the least number of stable connections among the three routing protocols. This phenomenon can be explained as follows. When a mesh client changes its associated mesh AP from the old AP to the new AP, its routing path breaks. However, the mesh APs on the routing path will need to wait for a long time before detecting such an event. Since the waiting time is long and a mesh client constantly changes its associated mesh AP while it moves, the disruption to a mesh client's connection is too long and too often, which makes many mesh clients' TCP connections unstable.

In OSPF, when the new mesh AP gets the IEEE 802.11(b) association control packet from the mesh client, it broadcasts a LSA to inform other mesh APs that the mesh client now is associated with itself. If the old and new mesh APs are within each other's wireless transmission range, the LSA can reach the old mesh AP very quickly. This enables the old mesh AP to promptly forward the mesh client's packets to the new mesh AP and shortens the period of disruption to the mesh client. Fig. 16 shows the details.

In STP, a similar mechanism is used to deal with node mobility. In STP, when the new mesh AP gets an association packet, it broadcasts a control packet up the spanning tree to inform upper-level

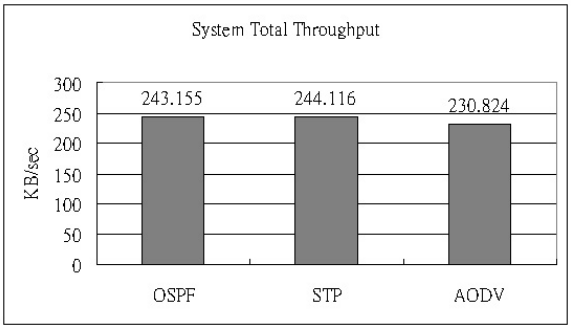

Fig. 14. The system total throughput under the mobility condition

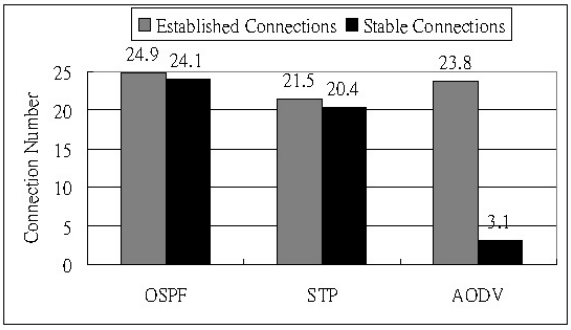

Fig. 15. The number of established and stable connections under the mobility condition

mesh APs of this association change. When the control packet reaches an appropriate layer in the spanning tree, the packets destinated to the mesh client will be directed toward the right branch of the spanning tree, which ends the period of disruption to the mesh client. Fig. 17 shows the details.

From above discussions, we can see that in dealing with node mobility, OSPF performs better than STP, which in turn performs better than AODV. Their performance in mobility conditions are reflected in the number of stable connections in the WMN.

D. Multi-Gateway WMN

Here we study the system total throughput of a WMN when it has one, two, three, and four mesh APs connecting to the mesh Internet gateway, respectively. The chosen mesh APs are the ones on the corners of the grid WMN and Fig. 18 shows the network topology of a two-gateway WMN. In the WMN, mesh APs run the OSPF routing protocol. The traffic settings of these cases are the same as those in the downlink TCP traffic case except that now multiple mesh APs connect to the mesh Internet gateway rather than just one. Fig. 19 shows the simulation results of these cases. We see that using more gateways in a WMN can significantly improve the system total throughput when most traffic in the WMN are Internet traffic. These results show that enough gateways should be deployed in a WMN to make its performance scalable with the number of mesh clients.

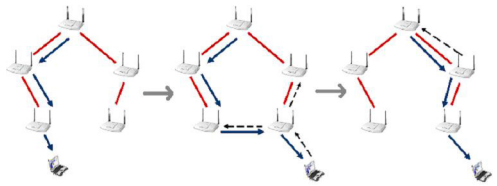

Fig. 16. The handling of node movement in OSPF

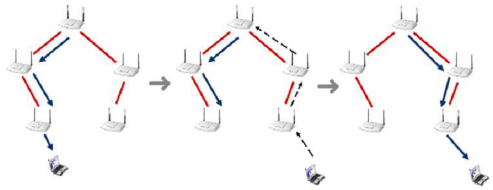

Fig. 17. The handling of node movement in STP

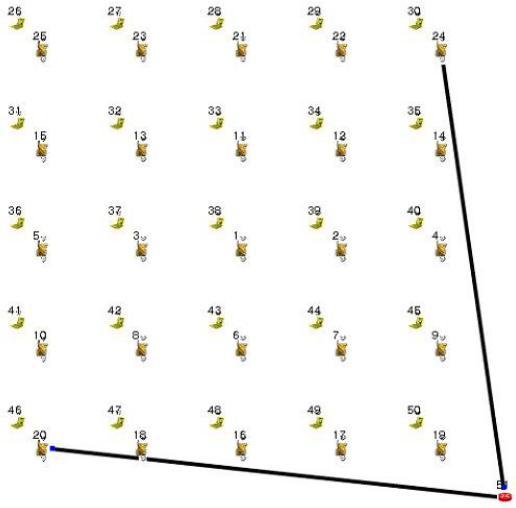

Fig. 18. The network topology of a two-gateway 5x5 grid WMN

E. OSPF with ETX

Here we study the effect of ETX when it is combined with the OSPF routing protocol. The traffic settings of the studied cases are the same as those in the downlink TCP traffic case. To let ETX show its capability in harsh wireless environments, a more realistic signal propagation model (the two-ray ground model) and bit-error model (the BER vs. power model) were used in the wireless PHY module during simulation. The received power of a packet is calculated based on the distance between the source and destination nodes. It is then added with a random fading with a variance of 10 dbm. Based on the combined power, the BER (bit-error-rate) for the received packet is calculated. The received packet is then dropped with a probability based on the calculated BER.

Fig. 20 shows that OSPF with ETX generates a lower system total throughput than OSPF in such a harsh environment while Fig. 21 shows that OSPF with ETX allows more stable connections to

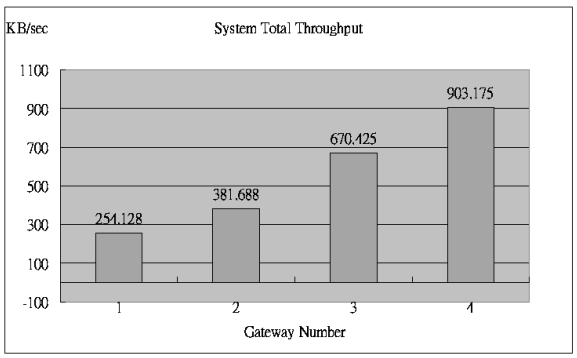

Fig. 19. The system total throughput of a multi-gateway WMN with different number of gateways

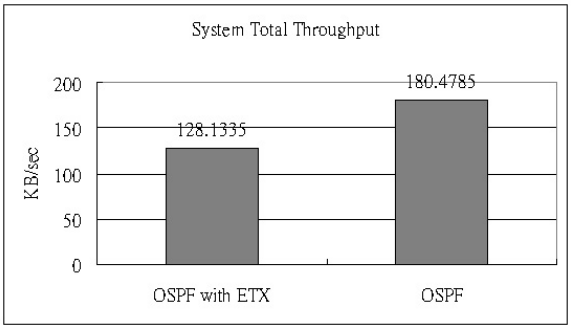

Fig. 20. The system total throughput under OSPF and OSPF with ETX in a harsh wireless environment

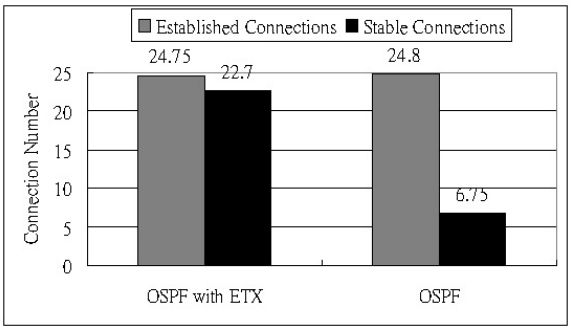

Fig. 21. The number of established and stable connections under OSPF and OSPF with ETX in a harsh wireless environment

coexist than OSPF in such a harsh environment. We found that in a harsh environment the shortest paths selected by OSPF are usually with high BERs and thus fragile. As such, most TCP connections with high hop counts are broken (in the TCP timeout state) and contribute little to the system total throughput. Instead, the system total throughput is mostly contributed by those TCP connections with low (2 or 3) hop counts. In other words, a few "short" TCP connections (initiated by those mesh clients that are close to the gateway mesh APs) monopolize the system bandwidth of the WMN and block out many "long" TCP connections (initiated by those mesh clients that are far away from the gateway mesh APs). Because the RTTs and hop counts of these "short" TCP connections are small, the TCP congestion control of these connections allow them to rapidly pump their data into the WMN even when they experience some packet losses.

In contrast, using the ETX metric in OSPF helps OSPF choose a low-BER and higher-throughput routing path for a connection. As such, "long" TCP connections become more robust and more of them can achieve a high throughput. However, the cost of this more even sharing of system bandwidth among "short" and "long" TCP connections is the reduced system total throughput. This is because the "long" TCP connections in OSPF with ETX cannot react to packet losses as rapidly as the "short" TCP connections in OSPF.

F. One-Radio WMN vs. Dual-Radio WMN

One advantage of dual-radio-dual-mode WMNs over one-radio WMNs is that client-AP traffic and AP-AP traffic can be transported over different frequency channels at the same to increase the system total throughput. To see how the second radio improves the system total throughput, we conducted two tests. In the first test, the ad-hoc mode and infrastructure mode interfaces of a mesh AP are set to use different channels. In contrast, in the second test, these two interfaces are set to use the same frequency channel. Because IEEE 802.11(b) MAC employs a carrier-sense multiple access (CSMA) mechanism to avoid collisions, multiple transmissions cannot be conducted at the same time on the

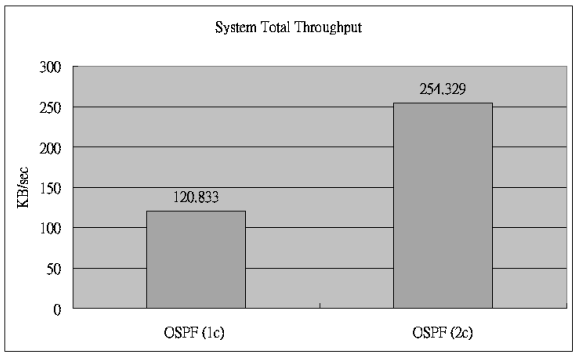

Fig. 22. The system total throughput of one-radio and dual-radio WMNs

same channel. As such, although the configuration of the WMN in the second test is not exactly the same as that of a single-radio WMN, their effects are the same — client-AP traffic and AP-AP traffic cannot be transported at the same time in a WMN.

Fig. 22 shows the system total throughputs achieved in the first and second tests when the OSPF routing protocol is used. The notation OSPF (1c) represents the case in which a mesh AP runs OSPF and its two interfaces are set to use the same channel. Similarly, OSPF (2c) represents the case in which a mesh AP runs OSPF and its two interfaces are set to use different channels. From this figure, we see that using two radios can improve the system total throughput significantly.

VI. Conclusion

In this paper, we studied the performance of three routing protocols (AODV, STP, and OSPF) operating in dual-radio-dual-mode infrastructure wireless mesh networks (WMN). Each of these protocols represents a different and popular approach that has been adopted by some commercial WMN products. AODV represents the ad hoc routing protocols developed for mobile ad hoc networks, STP represents the traditional routing approach developed for fixed bridges/switches, and OSPF represents the traditional routing protocols developed for the Internet. Because of their respective importance, in this paper we studied and compared their performances when they are applied to WMNs.

Our simulation results show that OSPF outperforms the others in dual-radio-dual-mode WMNs. Compared with AODV and STP, it provides higher system total throughputs, allows more stable connections to coexist in a WMN, and responds more quickly to movement of mesh clients. Our results show that with the ETX metric support, the system total throughput of a WMN can be more evenly shared by mesh clients rather than monopolized by just a few mesh clients. In this paper, we also implemented and studied the performance of multi-gateway WMNs. Our simulation results show that, when traffic in a WMN are mostly Internet traffic, a multi-gateway WMN can provide a much higher system total throughput than a single-gateway WMN. This suggests that one should deploy enough mesh gateways to make a WMN scalable with the number of mesh clients.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This research was supported in part by MOE Program for promoting Academic Excellence of Universities under the grant number 91-E-FA06-4-4, NSC under the grant number 94-2213-E-009-107, and III under the grant number 94C501.

References

[1] I. F. Akyildiz, X. Wang, and W. Wang. Wireless mesh networks: A survey. *Computer Networks Journal (Elsevier)*, 2005.
[2] IEEE Std. 802.11-1999. Part II: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. 1999.
[3] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. *ACM MobiCom*, September 2004.

[4] BelAir Networks. Mesh networking options. *available from http://www.belairnetworks.com/about_belair/challenge_mesh.cfm/.*

[5] MeshDynamics Structured Mesh. *available from http://www.meshdynamics.com/.*

[6] E. Royer and C. Toh. A review of current routing protocols for ad-hoc mobile wireless networks. *IEEE Personal Communication Magazine*, 6(2):46–55, 1999.

[7] Microsoft Research. Mesh networking. *available from http://research.microsoft.com/mesh/.*

[8] D. Aguayo, J. Bicket, S. Biswas, D. D. Couto, and R. Morris. MIT roofnet implementation. *available from http://pdos.lcs.mit.edu/roofnet/design/.*

[9] D. Maltz, D. Johnson, and Y. Hu. The dynamic source routing protocol for mobile ad hoc networks (DSR). *Internet draft*, April 2003.

[10] D. Maltz, D. Johnson, and Y. Hu. Dynamic source routing in ad hoc wireless networks. *Internet draft*, December 1998.

[11] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc on-demand distance vector (AODV) routing. *IETF RFC 3561*, July 2003.

[12] J. Moy. OSPF version 2. *IETF RFC 2328*, April 1998.

[13] Tropos Networks. *available from http://www.tropos.com/technology/whitepaper.html/.*

[14] IEEE 802.1D. *ANSI/IEEE Std 802.1D*, 1998 Edition

[15] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the 9th ACM International Conference, on Mobile Computing and Networking (MobiCom)*, San Diego, California, September 2003.

[16] S. Wang, C. Chou, C. Huang, C. Hwang, Z. Yang, C. Chiou, and C. Lin. The design and implementation of the NCTUns 1.0 network simulator. *Computer Networks*, 42(2):175–197, June 2003.

[17] NCTUns 2.0. *available from http://nsl10.csie.nctu.edu.tw/*, 2005.

[18] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. *ACM Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 66–80, September 2003.

[19] L. Huang and T. Lai. On the scalability of IEEE 802.11 ad hoc networks. *ACM International Symposium on Mobile Ad Hoc Networking and Computer (MobiHoc)*, pages 173–182, 2002.

[20] P. Gupta and P. Kumar. The capacity of wireless networks. *ACM Annual International Conference on Mobile Computing and Networking (MOBICOM)*, March 2000.

[21] D. Aguayo, J.Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. *SIGCOMM*, August 2004.

[22] K.W.Chin, J. Judge, A. Williams, and R. Kermode. Implementation experience with MANET routing protocols. *ACM CCR*, November 2002.

[23] B. Awerbuch, D. Holmer, and H. Rubens. High throughput route selection in multi-rate ad hoc wireless networks. *Technical report*, 2003.

[24] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. *ACM SIGCOMM Conference*, pages 133–144, August 2004.

[25] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *BroadNets*, 2004.

[26] A. Raniwala, K. Gopalan, and T. Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *MC2R(Vol.8)*, pages 50–65, 2004.

[27] A. Raniwala and T. Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. *proceeding of IEEE INFOCOM*, 2005.

[28] P. Kyasanur and N. Vaidya. Routing and interface assignment in multi-channel multi-interface wireless networks. *proceeding of IEEE WCNC*, 2005.

[29] N. V. J. So. A routing protocal for utilizing multiple channels in multi-hop wireless networks with a single transceiver. *UIUC Technical Report*, 2004.

[30] C. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. *ACM SIGCOMM Conference (SIGCOMM '94)*, pages 234–244, August 1993.

[31] G. Malkin. RIP version 2. *IETF RFC 2453*, November 1998.

[32] IEEE 802.1. *ANSI/IEEE Std 802.1*, 1998 Edition

[33] K. Fall and K. Varadhan, Eds. ns manual/ns Notes and Documentation. http://www.isi.edu/nsnam/.