

# DMesh: Incorporating Practical Directional Antennas in Multichannel Wireless Mesh Networks

Saumitra M. Das, *Student Member, IEEE*, Himabindu Pucha, *Student Member, IEEE*, Dimitrios Koutsonikolas, *Student Member, IEEE*, Y. Charlie Hu, *Member, IEEE*, and Dimitrios Peroulis, *Member, IEEE*

**Abstract**—Wireless mesh networks (WMNs) have been proposed as an effective solution for ubiquitous last-mile broadband access. Three key factors that affect the usability of WMNs are high throughput, cost-effectiveness, and ease of deployability. In this paper, we propose DMesh, a WMN architecture that combines spatial separation from directional antennas with frequency separation from orthogonal channels to improve the throughput of WMNs. DMesh achieves this improvement without inhibiting cost-effectiveness and ease of deployability by utilizing *practical* directional antennas that are widely and cheaply available (e.g., patch and yagi) in contrast to costly and bulky smart beamforming directional antennas. Thus, the key challenge in DMesh is to exploit spatial separation from such practical directional antennas despite their lack of electronic steerability and interference nulling, as well as the presence of significant sidelobes and backlobes.

In this paper, we study how such practical directional antennas can improve the throughput of a WMN. Central to our architecture is a distributed, directional channel assignment algorithm for mesh routers that effectively exploits the spatial and frequency separation opportunities in a DMesh network. Simulation results show that DMesh improves the throughput of WMNs by up to 231% and reduces packet delay drastically compared to a multiradio multichannel omni antenna network. A DMesh implementation in our 16-node 802.11b WMN testbed using commercially available practical directional antennas provides transmission control protocol throughput gains ranging from 31% to 57%.

**Index Terms**—Directional antennas, multiple channels, wireless mesh networks (WMNs).

## I. INTRODUCTION

WIRELESS MESH NETWORKS (WMNs) are characterized by static mesh routers connected by wireless links to each other, each providing connectivity to end hosts. The most significant application of such networks is to provide broadband access where wired infrastructure is difficult or economically infeasible to deploy. Three key factors that affect the usability of WMNs are high throughput, cost-effectiveness, and ease of deployability. These characteristics allow WMNs to support a

large number of users, while being competitive with other access services.

There have been several significant efforts on improving the throughput of WMNs [1], [2]. These works aim to improve throughput using multiple radios that utilize multiple channels (available in the IEEE 802.11a/b/g standards) to separate the contending transmissions in the frequency domain. However, these works assume the use of omnidirectional antennas at mesh routers, where a transmission on a given channel requires all other nodes in range to remain silent or use alternative channels. Thus, although multiple channels can separate the transmissions in the frequency domain, the extent of such separation is potentially limited by the number of available channels.

- 1) There exists a bound on the number of channels possible due to spectrum regulation. For example, 802.11a has 12, while 802.11b has only three nonoverlapping channels.
- 2) Further, the number of available channels could potentially be even lower due to stricter spectrum regulation in some countries or channels being set aside to support other communication (e.g., among mesh routers, users and mesh routers, or other networks).
- 3) Finally, significant co-channel interference exists even among the nonoverlapping channels [2], which limits the extent of frequency separation among contending transmissions. Due to these limitations, there is a need to find other means of separating contending transmissions to improve the throughput of WMNs.

Compared with omni antennas, directional antennas offer *spatial separation* between contending transmissions and have the potential to further enhance the throughput of WMNs. In this paper, we propose DMesh, a WMN architecture that combines spatial separation from using directional antennas with frequency separation from using orthogonal channels to improve the throughput of WMNs. An important requirement for DMesh is to accomplish this throughput improvement without inhibiting the other two key WMN requirements: cost-effectiveness and ease of deployability. The high cost of *smart* beamforming directional antennas and their form factor make it difficult to achieve these two requirements. Thus, in DMesh, we focus our effort on incorporating *practical* directional antennas that are widely and cheaply available (e.g., patch and yagi). The key challenge in DMesh is to exploit spatial separation from such practical directional antennas despite their lack of electronic steerability and interference nulling, as well as the presence of significant sidelobes and backlobes. Further, enabling DMesh poses new design challenges, as it requires specialized protocols to enable routing, and channel

Manuscript received October 1, 2005; revised March 4, 2006 and May 1, 2006. This work was supported in part by the National Science Foundation under Grant ANI-0338856, in part by the Purdue Research Foundation, and in part by an equipment grant from Intel Corporation.

S. M. Das, H. Pucha, D. Koutsonikolas, and Y. C. Hu are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: smdas@purdue.edu; hpucha@purdue.edu; dkoutson@purdue.edu; ychu@purdue.edu).

D. Peroulis is with the Birk Nanotechnology Center, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: dperouli@purdue.edu).

Digital Object Identifier 10.1109/JSAC.2006.881631

assignment to exploit directionality. Central to our architecture is a distributed, directional channel assignment algorithm for mesh routers that exploits directional antennas and multiple channels to separate the contending transmissions in both *spatial* and *frequency* domains. Effectively, DMesh enables two degrees of separation between contending transmissions.

This paper makes the following contributions: 1) we propose and evaluate a WMN architecture that exploits both directional antennas for spatial separation and multiple orthogonal channels for frequency separation to provide significantly increased throughput; 2) we focus on cost-effective and deployable techniques that can easily be incorporated to extend current single-interface networks such as Roofnet [3]; 3) we describe and evaluate a distributed routing protocol along with associated directional channel assignment algorithms to fully exploit the proposed architecture, using a realistic and detailed antenna model; 4) we evaluate the proposed architecture and compare it to solutions that use omnidirectional antennas in a detailed simulator; and 5) we validate the performance of our architecture using experiments in a mesh network testbed with the same antennas modeled in the simulations.

Experimental results show that compared to a multiradio multichannel network using omni antennas, a DMesh network increases the average per source throughput from 128% for 25 traffic sources up to 231% for 50 traffic sources and provides up to ten times smaller packet delays in an 802.11a network with 50 mesh routers. The throughput improvement from DMesh is more pronounced when the number of available channels is limited. When only 6 instead of 12 channels in an 802.11a network are used, the throughput using DMesh is improved by up to 176% for 25 traffic sources with significantly lower delay in comparison to a multichannel omni network. DMesh also provides transmission control protocol (TCP) throughput gains between 30% and 68% in 802.11a and 35% in 802.11b networks. We implemented DMesh in our 16-node wireless mesh network testbed using commercially available practical directional antennas. Our testbed evaluation shows that DMesh increases TCP throughput in 802.11b WMNs between 31% and 57% compared with a multichannel omni WMN.

The rest of this paper is organized as follows. In Section II, we describe DMesh, our architecture for a wireless mesh network with directional antennas and multiple channels. Our evaluation methodology and results are detailed in Sections III and IV. An evaluation of the real-world performance of DMesh in a testbed is described in Section V. Section VI summarizes the related work and Section VII concludes this paper.

## II. DMESH: A DIRECTIONAL WMN

In this section, we first describe the architecture of DMesh followed by descriptions of the associated physical tree formation, routing protocol, and distributed directional channel assignment.

We consider a typical single-channel, single-interface WMN deployment with omnidirectional antennas in which mesh routers are placed on the rooftops of subscribers [3] or other infrastructure (e.g., streetlights) and are interconnected via 802.11 links. DMesh is then used to enhance the performance of this existing network. To enable DMesh, we assume that each mesh router can have up to  $k$  additional interfaces (multiple

radios) each with its own practical directional antenna. The practical directional antennas used in DMesh are nonsteerable and they always point to the direction toward which they were manually placed during the network deployment. Since mesh routers are not highly form factor-constrained like other wireless devices [e.g., personal digital assistants (PDAs)], the use of multiple radios and antennas is feasible. In fact, multiradio mesh routers are already commercially available [4]. Given this DMesh architecture, the omnidirectional interface of each mesh router is always available for providing robust connectivity if the directional neighbors fail. Further, this omni interface is also used as a CONTROL interface.

The common application scenario for WMNs is likely to be Internet access in which most traffic will flow to and from the gateway nodes which have wired connections to the Internet. These gateways will typically be deployed sparsely in the network due to their higher cost. The number and placement of these gateways are likely to be based on the area of deployment and availability of resources. Alternatively, in some scenarios, it may be possible to place such gateways intelligently [5].

To support the application scenario described above, we propose to build high throughput routing trees rooted at the gateways. DMesh consists of three main stages to construct such trees: 1) physical tree formation to provide gateway connectivity using the directional antennas; 2) routing state creation and maintenance to correctly deliver packets along the tree; and 3) channel assignment to separate spatially contending transmissions whenever possible to further increase the throughput.

Optimized link-state routing (OLSR) [6] is used as the multihop routing protocol in our single-radio single-channel omnidirectional mesh network testbed (described in Section V). DMesh extends the OLSR protocol to aid in physical formation of trees using practical directional antennas, set up and maintain corresponding routing state, and perform channel assignment. We call this extended protocol DOLSR (*directional OLSR*).

### A. Physical Tree Formation

The physical tree formation takes place as follows. When each mesh node boots up, the DOLSR daemon reads the configuration file (*olsr.conf*) to identify the interface attached to an omni antenna and starts running in the single-channel single-interface mode. The configuration file contains a list of interfaces, along with the type of antenna attached to the interface, as well as a PIP (pointing IP address) for each interface. The use of PIP is explained later in this section. In this mode, the node starts listening for *Host and Network Association* (HNA) messages on its omni interface. Nodes that have Internet access (i.e., a default route to a gateway or the gateway itself) advertise their connectivity over their omni interface using periodic HNA packets. The joining node picks the best next hop node based on metrics measured to nodes from which it receives HNA messages. The metric could be of various types, for example, hop count, remaining path bandwidth or a measurement based metric (e.g., ETX [7]). After the new node joins, it continues to receive HNA messages to reevaluate its choice of a next hop.

To incorporate directional antennas into this existing network, we modified the HNA message to also advertise whether a node is willing to host directional interfaces for its neighbors. A node

can host directional interfaces after it itself has Internet connectivity through a directional interface or if it is the gateway node. In addition, the node should have free interfaces to connect to new children. For example, the gateway node bootstraps DMesh by advertising its willingness to host directional interfaces in its HNA packets. The one-hop neighbors of the gateway receive these HNA packets.

When DOLSR receives HNA packets on the omni interface with willingness indicated, it also evaluates the best potential *directional* next-hop node (PARENT) using a suitable metric. For example, throughput measurements using the omni interface can be used to drive the PARENT selection. This choice of PARENT is then reported, directional antennas on the PARENT node and CHILD node (joining node) are installed pointing at each other, and corresponding entries in the *olsr.conf* on the PARENT and CHILD nodes are modified by offline means to include information about the new directional interface. Specifically, the PARENT node marks the interface as *DIR* and stores the IP address of the CHILD node's interface it is pointing to and *vice versa*. This IP address is referred to as the PIP and only packets received from the PIP configured for a particular interface are passed to the DOLSR protocol. The PIP is required to filter out packets received from the sidelobes from other nodes that the antenna is not pointing at. Directional interfaces are initially configured to operate on a default channel to enable connectivity. Note that in a cooperative network such as RoofNet, a CHILD could purchase antennas for its PARENT node that is willing to host an interface. On the other hand, in a commercial WMN, the Internet service provider (ISP) takes care of the installation.

After installation is complete and the updated configuration file is read in by the DOLSR daemon, the PARENT node sends out READY packets on the newly installed directional interface. The READY message acts as a trigger for the routing protocol on the CHILD node to set up the routing state (explained in the next section). Once the routing state is set up, the CHILD node also starts indicating willingness to host directional interfaces and more nodes can now join the tree. In this way, nodes install at least one-directional antenna on one of their interfaces pointing to a PARENT node, and a tree will be physically formed toward the gateway. In this paper, we evaluate our architecture over many different topologies generated by simulating the above tree formation model.

### B. Routing Protocol

The *physical* placement of directional antennas in DMesh naturally forms a tree structure. On top of this physical tree, DOLSR now sets up forwarding entries to route packets along the tree. The routing state is set up as follows. Let  $T_a$  be any PARENT node. As mentioned earlier,  $T_a$  periodically broadcasts READY messages on its newly installed directional interface. Only the corresponding CHILD node will receive the READY message due to filtering based on the PIP. On receiving READY messages from  $T_a$ , the CHILD node sends a JOIN message to  $T_a$ .  $T_a$  uses the IP address received from the JOIN packet to set up a forwarding entry toward the joining node. The joining node uses the PARENT's IP address (from the READY message) as a default route and inserts this information into its

routing table. Once this exchange is completed, the joining node becomes part of the tree and sends its own HNA packets on the omni interface and can potentially become a PARENT for other nodes. Note that nodes that have the gateway as a PARENT node are denoted SUPERPARENTs.

Once the JOIN packet is received, a ROUTE\_SETUP message initiated by  $T_a$  is multicast along the tree simply by having each node recursively send the message to each of their active interfaces. The multicast proceeds until the gateway is reached since the gateway does not rebroadcast the message to the other SUPERPARENTs. The ROUTE\_SETUP packet contains the IP address of the interface that the joining node used in the JOIN packet. This information is used to set up forwarding table entries in the routing tables of all nodes in the subtree including the gateway, to reach the newly joined node. This information also enables peer-to-peer routing (between mesh routers): a mesh router can route to all mesh routers in its own subtree rooted at its SUPERPARENT, and it can route to the mesh routers that are offsprings of other SUPERPARENTs by going through the gateway. A node only needs to consult its forwarding table to determine the next hop for any packet. A packet that is not matched with any forwarding entry is sent on the default path to the gateway, i.e., to the node's PARENT.

*Switching Parents:* In certain situations, a node may be required to switch parents, e.g., if a better next-hop is detected through metric measurements or when a new gateway is installed. For example, when a new gateway is installed, each gateway starts its own HNA, and to the mesh routers, this is treated like any other HNA: if the metric advertised from a different gateway is lower than the one the node is currently using, the node will find that to be a better choice. When a node changes its PARENT,<sup>1</sup> it multicasts a LEAVE message similar to the ROUTE\_SETUP message to invalidate old routing entries. When a node changes its PARENT, the subtree rooted at that node also moves to the new parent. The new PARENT multicasts a ROUTE\_SETUP message to refresh the routing entries at all nodes in the subtree of its SUPERPARENT. In this manner, the routing state is maintained despite physical topology changes. For simplicity of evaluation, in this paper, we consider only a single gateway.

An internal node reattaching to a new PARENT can cause cycles to form. To eliminate this, we restrict the choice of the new PARENT to a node that is not part of a node's own subtree as follows: every node keeps track of nodes in the subtree rooted at itself and does not choose any of these nodes as parents, and this is accomplished by each node storing a list of nodes that sent ROUTE\_SETUP messages through it. Also, in our architecture, multihoming or multiple parents are not supported and each node has only one PARENT.

*Failure Handling:* In DMesh, issues such as simultaneous joining and oscillations are rare since only a designated node (with matching PIP) can attach to an interface. If a node does not receive any READY messages from its directional interfaces, it defaults to using the omnidirectional interface which exists as a backup for robust connectivity.

<sup>1</sup>This is done by manual realignment of the antenna and software configuration changes.




Fig. 1. Channel assignment schemes. (a) Omnidirectional with OCA. (b) Directional with C-DCA. (c) Directional with A-DCA.

In summary, DOLSR aids in the formation of the physical DMesh tree, as well as sets up and maintains routing state in the presence of physical topology changes.

C. Distributed Directional Channel Assignment

Once the routing state is set up, one degree of separation has been achieved since more parallel transmissions can now occur in the network due to spatial separation. However, practical directional antennas cannot provide complete spatial isolation due to the presence of *sidelobes* and *backlobes*. Therefore, DMesh also uses directional channel assignment to achieve the second degree of separation by separating these spatially contending transmissions in the frequency domain to further increase the throughput of the mesh network.

The directional channel assignment is performed by a PARENT for its CHILD node when the node joins. As mentioned earlier, both interfaces are initially configured on the same default channel to enable connectivity. The PARENT selects a new channel using a channel assignment scheme and sends an ASSIGN message to the CHILD. Both PARENT and CHILD then use *iwconfig* to set their respective directional interfaces into the selected channel. Following this, bidirectional communication can occur. The channel assignment is periodically reevaluated every 300 s.

1) *Channel Assignment Schemes*: Our objective is to design a distributed directional channel assignment algorithm that provides good performance and is easy to implement and deploy. In the following, we discuss our two omnidirectional and four-directional channel assignment schemes using an example. Consider the scenario in Fig. 1(a) where there are five mesh routers (A–E) and two flows  $A \rightarrow B$  and  $C \rightarrow D$ . Assume that E has a flow terminating at some other node in the direction indicated.

*Omnino/No Channel Assignment (ONOCA)*: In this scheme, omnidirectional antennas are used without multichannel techniques. If A transmits to B omnidirectionally, then the channel is reserved in the entire region due to the ready-to-send (RTS) and clear-to-send (CTS) range of A and B, as shown in Fig. 1(a). Thus, only one flow can proceed. In this network, there is no spatial or frequency separation possible between contending transmissions in a single collision domain.

*Omnio/Channel Assignment (OCA)*: In this scheme, omnidirectional antennas are used with multichannel techniques. Each node first attempts to pick a unique unused channel for communicating to its PARENT. If the node cannot obtain an un-

used channel, it reuses the least loaded channel among those that have been selected by nodes in its interference range. This algorithm is similar to the one described in a previous work [1] and is implemented in a distributed manner by periodic exchange of channel usage information and least loaded channel selection. A least loaded channel is defined to be the one being used by the lowest number of flows. If the number of flows using a pair of channels is the same, then the traffic transmitted over the channels over the last information exchange period is used to determine the least loaded channel. Thus, OCA exploits the frequency separation offered by multiple channels to allow contending transmissions in the same collision domain. If we apply this to Fig. 1(a), assuming two channels exist in the network, both flows  $A \rightarrow B$  and  $C \rightarrow D$  can proceed using different channels  $C_0$  and  $C_1$ . However, E still cannot transmit, since it is in the contention range of the other two flows. E realizes this through exchange of channel usage information and picks the least loaded channel in its vicinity.

*Directional/No Channel Assignment (DNOCA)*: In this scheme, directional antennas are used without multichannel techniques. Under this scheme, flows  $A \rightarrow B$  and  $C \rightarrow D$  can proceed since their *cone of interference* (the area defined by the sector of the transmitting nodes) is separated spatially. However, E cannot transmit since it lies in the cone of interference of flow  $C \rightarrow D$  that uses the only available channel  $C_0$ .

*Directional/Channel Assignment (DCA)*: In DCA, directional antennas are used with multichannel techniques, and we aim to separate spatially (directionally) contending transmissions in the frequency domain. Two transmissions are defined to be *directionally* contending if either one of the nodes originating the transmission lies in the cone of interference of the other.

We propose and evaluate three directional channel assignment schemes. The first scheme is *conservative DCA (C-DCA)*. In C-DCA, given a set of channels  $C$ , a node X assigns a channel to a CHILD in two stages.

- 1) It first attempts to find a free channel  $C_i$ . A free channel for node X is defined as one that is not used by any node whose cone of interference contains X or its CHILD, i.e., no node in the vicinity of X and its CHILD is directionally contending on channel  $C_i$ .
- 2) If no such channel exists, the node X selects a channel  $C_i$  that is least loaded. If we use C-DCA in Fig. 1(b), E will select to transmit using  $C_1$ , effectively insulating its

transmission from its directionally contending transmissions. C-DCA is conservative because it always assigns a new channel to a transmission that lies in the cone of interference of another transmission. The C-DCA heuristic requires exchange of channel usage information in a local area to allow for selection of channels. For example, node E needs to receive channel usage information from C and D and perform geometric calculations to realize that it is in the cone of interference of D, while executing the C-DCA algorithm. Calculating the cone of interference requires knowledge of the positions of C and D which is easily available in a static mesh network [e.g., through a localization scheme or global positioning system (GPS)].

In a mesh network with no mobility, it is even relatively simple to just encode the location of the mesh router when deploying or installing it, and no GPS hardware needs to be permanently installed in the mesh router. We use this one-time encoding method in our testbed. Other deployed testbeds such as RoofNet also maintain the GPS coordinates of their mesh routers.

Note that C-DCA effectively takes into account the interference from sidelobes on a receiving node's antenna by assuming that the antenna can potentially receive interference from any direction (not just from the main lobe). In other words, C-DCA assumes the possibility of interference if a node lies in the transmit cone of neighboring nodes regardless of the orientation of the node's own antenna. For example, E could potentially use  $C_0$  if its flow is not directed toward C or D. However, this would make E prone to interference through power received on its sidelobes on  $C_0$ . C-DCA does not allow the selection of  $C_0$  to avoid such interference from sidelobes.

In addition, C-DCA takes into account the interference from the transmission power from a neighboring node antenna's sidelobes. Since C-DCA assumes a *flat-topped* transmit cone to make the geometric calculations simpler and applicable to any antenna, it can only anticipate interference caused by transmission through a main lobe of a neighboring node's antenna. To take into account the power from the sidelobes of a neighboring node's antenna, we incorporate measurement-based enhancements in the channel assignment as follows: Each node overhears transmissions on all its directional interfaces and reception of data on an interface from any node whose IP address is different from the PIP for that interface indicates interference on the channel being used on that interface.

Our second assignment scheme, *aggressive DCA* (A-DCA), assigns channels similar to C-DCA but with one important difference: X is considered to be in the cone of interference of another node Y, only if both X and Y lie in each other's cone of interference. Thus, in Fig. 1(b), if E transmits in the direction shown, although E lies in the cone of interference of D, D does not lie in the cone of interference of E, and thus node E using A-DCA will reuse channel  $C_0$ . Thus, A-DCA is more aggressive in identifying opportunities where channels can be reused based on directionality, and thus results in a reduced channel usage. This can result in more interference (e.g., if E has sidelobes it will receive interference power from D). To accommodate interference from other transmissions, A-DCA uses a "guard angle" to control the aggressiveness of channel assignment as follows: the beamwidth of all transmit and receive cones during the geo-

TABLE I  
ARCHITECTURE CHOICES IN A MULTIRADIAL MESH NETWORK

Antenna	Channel Assignment
Omni	ONOCA, OCA
Directional	DNOCA, C-DCA, A-DCA, M-DCA

metric interference calculations are increased by guard angle degrees. Thus, a higher guard angle makes A-DCA more conservative when reusing channels, while a lower guard angle makes A-DCA aggressive in reusing channels. Since the cone of interference is larger, nodes close to the transmission are likely to be separated in frequency. A typical guard angle value is  $30^\circ$ . Note that A-DCA does not use measurement-based enhancements, since that would limit its aggressiveness.

Finally, we evaluate a third assignment scheme, *measurement-based DCA* (M-DCA), which uses only measurements instead of geometry to infer interference between two nodes. M-DCA works similarly to the measurement-based enhancement in C-DCA, where each node overhears transmissions in all its directional interfaces, and uses this information to discover interfering nodes.

Table I summarizes the various WMN architecture choices considered in this paper. Using omnidirectional antennas, we can apply OCA and ONOCA schemes, whereas using directional antennas we can apply DNOCA, A-DCA, M-DCA, and C-DCA.

2) *Distributed Algorithm for OCA/C-DCA/A-DCA/M-DCA*: In order to select a channel in a distributed manner using local knowledge, a node needs to know the channel usage in its vicinity. This channel usage information is locally exchanged since the channel can be reused at a sufficient distance. Thus, each node periodically (every 60 s) uses the omni CONTROL interface to broadcast its own channel usage information. The broadcast uses sequence numbers to suppress duplicates and is rebroadcasted for two hops in order to cover the interference range of that node which is assumed to be twice that of the communication range. Note that in DMesh, the reachability of any node using the directional antennas is kept similar to the reachability achieved if the node were using an omnidirectional antenna by adjusting the transmission power on the directional interfaces. Also, the directional interfaces using simple non-steerable antennas cannot be used to broadcast channel usage information since they will not cover all possible nodes that need to receive this information.

In OCA/C-DCA/A-DCA, each node broadcasts a channel vector  $C$  that specifies whether it is using each channel  $C_i$ , and a rate vector  $R$  that specifies a time average of the amount of traffic it has transmitted on each channel  $C_i$  over the past 60 s. For A-DCA/C-DCA, nodes initially exchange their position information and cache it for future use since the mesh routers are static. In our testbed, this position information is encoded using GPS during installation. However, other localization techniques could be easily used for this purpose as well. For A-DCA/C-DCA, a destination vector  $D$  indicating the destination being communicated with on each channel  $C_i$  is included for directionality information. The final channel map maintained at every node is constructed from the individual state vectors ( $\langle C, R, D \rangle$ ) received from neighboring nodes. In

addition, each CHILD also sends its PARENT a list of channels (BCL) computed from its own channel map which contains all channels being used in the CHILD's neighborhood.

Once a node collects the channel usage information from its interfering neighbors (which may be more than one-hop away), it uses the assignment algorithms to assign channels to a new CHILD or periodically (every 300 s) reevaluate the channel assignments for its current children.

In summary, for all the schemes,<sup>2</sup> each node obtains a channel usage map which stores a list of node identifiers along with their associated channel and rate vectors. For A-DCA and C-DCA, a destination vector is also obtained. For M-DCA, there is no need of exchanging information between nodes. Each node monitors each of its directional interfaces for arrival of data from any node other than the child assigned to it. Reception of data from a non-child node indicates interference on the channel being used on that interface. This interfering node is then added to the channel usage map.

### III. EXPERIMENTAL METHODOLOGY

We use the QualNet simulator [8] to evaluate DMesh. QualNet has been widely used to study directional antennas and multichannel networks [9], [10], [11]. We model a node with multiple interfaces with support for dynamic channel assignment on each interface.

*MAC, Physical Layer, and Antennas:* We used IEEE 802.11a and 802.11b with autorate fallback. The simulator models OFDM and DSSS for 802.11a and 802.11b, respectively. Both MAC layers were verified to produce close to theoretically maximum throughput [12] and also model multirate operation based on internode distances. The two-ray path loss propagation model is used. We used real patterns for the directional antennas that faithfully model sidelobes. The patterns were taken from datasheets of commercially and cheaply available directional antennas from [13] and [14]. In fact, our testbed evaluation uses one of the antennas we model in the simulation.

*Topology and Traffic:* We simulated a static mesh network of 50 nodes placed randomly in an area of 1000 m  $\times$  1000 m. Each node is assumed to have one omni interface, one-directional interface to connect to a PARENT, and two additional directional interfaces to support up to two children. The directional interfaces are pointed to neighbors with a random error between 0–10° to simulate manual placement of antennas. The gateway is assumed to be at the center of the area. In order to isolate the effect of increased spatial reuse obtained from directional antennas from the range enhancement that also results, the reachability of any node using the directional antenna is kept similar to the reachability achieved if the node were using an omnidirectional antenna by adjusting the transmission power. We vary the number of sources and evaluate using both user datagram protocol (UDP) and TCP traffic. All the sources communicate with the gateway node to simulate an Internet access pattern. The rate for each source is picked randomly from 0 to  $\text{Rate}_{\max}$ .  $\text{Rate}_{\max}$  is varied from 100 kb/s to 5 Mb/s. A packet size of 1500 bytes is chosen.

*Metrics:* We use the following three metrics to evaluate the performance.

- 1) *Packet delivery ratio* (PDR)—the ratio of the overall number of successfully received data packets at the gateway to the number of data packets sent.
- 2) *Delay*—the average time between transmission and reception of data packets. This metric accounts for all possible delays caused by queuing at the interface queue, retransmission delays at the MAC, and propagation and transfer times.
- 3) *Average source throughput*—the total number of bytes successfully transferred from each source during the simulation duration, divided by the simulation duration and averaged over all sources. This metric is similar to PDR, except it also gives an indication of how much raw bandwidth is achieved. For all schemes, control overhead occurs only in exchange of channel usage information, and is negligible compared to the overall data traffic.

Therefore, in the results below, due to lack of space, we omit the control overhead for all schemes.

### IV. PERFORMANCE EVALUATION

In this section, we first evaluate the performance of various WMN architectures. We then perform controlled experiments by varying other parameters such as the number of available channels to study the tradeoffs involved in architecting a practical DMesh network. Before studying the various architectures, we evaluated several tree construction metrics (details in [15]) and chose the best tree construction metric found for use in all architectures studied. Specifically, an expected throughput (ETH) metric, in which a node probes the throughput (e.g., using the *netperf* tool) to its potential PARENTs (from whom it receives an HNA) and selects the PARENT with the highest throughput, was found to provide good performance through simple local measurements. We use this metric in our evaluation.

#### A. Overall Performance Comparison

In this section, we evaluate the performance gain from the different WMN architectures listed in Table I. The beamwidth of the directional antenna is assumed to be 45°.

Fig. 2 shows the performance of all the above mesh network architectures in terms of the packet delivery ratio, delay, and average source throughput. As expected, the PDR, delay and throughput of ONOCA are the worst among all architectures. This is because ONOCA does not exploit any degrees of separation, and is thus unable to support the simultaneous traffic load in the network.

DNOCA achieves better performance than ONOCA. DNOCA exploits spatial separation to support simultaneous transmissions and thus achieves higher throughput than ONOCA. At 500 kb/s traffic, the PDR of DNOCA is 48% better than that of ONOCA. However, spatial separation *alone* is not sufficient to support high throughput. The PDR of DNOCA drops drastically as traffic increases beyond 500 kb/s. This shows that practical directional antennas cannot be effectively used as complete spatial isolators. Two factors contribute to this poor performance: 1) the beamwidth of a practical directional antenna cannot be made arbitrarily small in order to completely separate all contending transmissions and 2) the presence of sidelobes limits the usability

<sup>2</sup>The pseudocode for A-DCA, C-DCA, and OCA is available in [15].

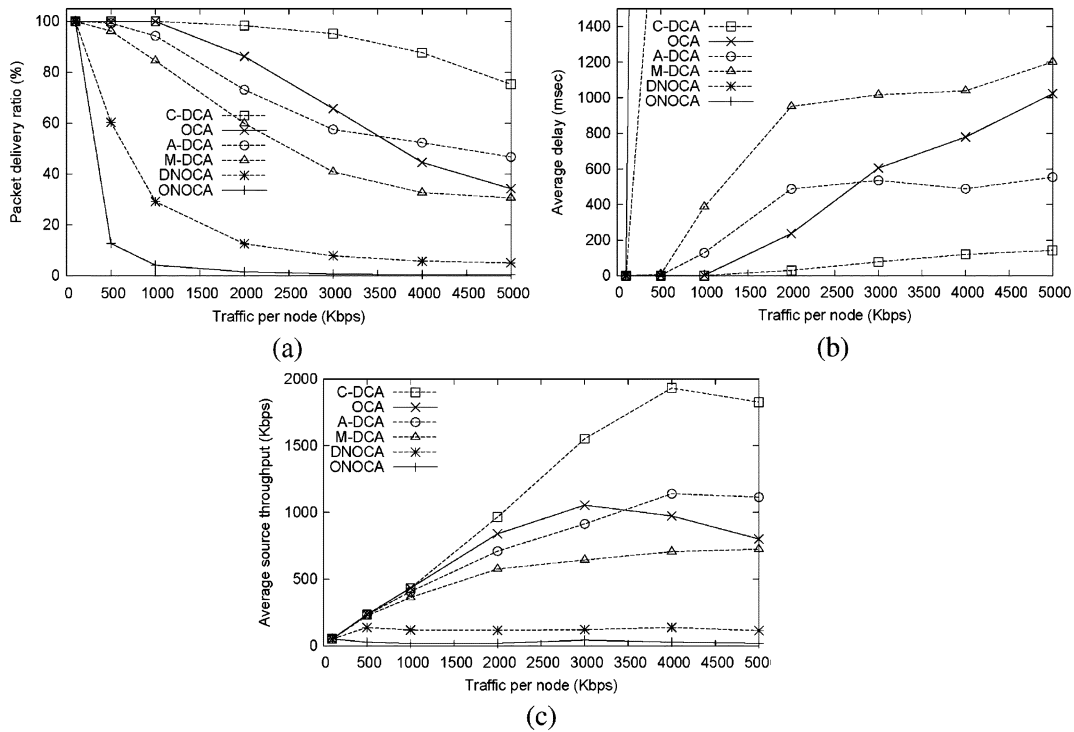


Fig. 2. Comparison of different mesh network architectures. (a) Packet delivery ratio (%). (b) Average delay (ms). (c) Average throughput (kb/s).

of multiple practical directional antennas at a node since one antenna will be able to hear the other due to reception from sidelobes. Thus, any pair of interfaces on the same node will always contend with each other. We also observed the sidelobe problem in our DMesh testbed and found that the problem was not mitigated even with greater physical separation between the antennas of a pair of interfaces.

OCA also achieves better performance than ONOCA. OCA exploits frequency separation to enable higher throughput than ONOCA. At 500 kb/s traffic, OCA achieves 88% higher PDR than ONOCA. However, separation by frequency alone is not sufficient for achieving the best performance. The PDR of OCA drops to only 40% and the delay drastically worsens as the traffic increases to 5 Mb/s. An important factor contributing to this degradation is that there is an upper bound on the number of distinct orthogonal channels available which limits the ability to completely isolate all contending transmissions by frequency. For example, an examination of the channel assignment in this scenario shows that OCA shares many channels among multiple flows. Moreover, as the traffic increases, the per channel load increases, thereby increasing the contention among flows and thus the throughput drops. Thus, although OCA delivers a fraction of the packets, these packets are delivered after significant amount of queuing and buffering leading to high average delays. Note that another significant factor that could restrict the performance of OCA is interchannel interference. Although the simulation evaluation of OCA assumes no inter-channel interference,<sup>3</sup> in our DMesh testbed, we observed that two interfaces maximally separated by orthogonal channels fail to achieve simultaneous throughput. In summary, the observed performance

<sup>3</sup>QualNet does not simulate interchannel interference and the code where this feature can be implemented is not open-source.

of both DNOCA and OCA suggests that an additional degree of separation can further improve the throughput.

Another interesting observation from the above experiments is that OCA significantly outperforms DNOCA. Although both schemes have a single degree of separation, separation by frequency alone is more effective than separation by space alone from using practical directional antennas. This effect was also observed during our testbed experiments. This is primarily because any pair of interfaces on a single node, when connected to practical directional antennas can always hear each other (due to sidelobes and backlobes) and it is difficult to separate them physically far enough apart from each other to obtain simultaneous throughput. On the other hand, any pair of interfaces on a single node separated by frequency try to filter out each others signal. Although this filtering is also not perfect, the separation achieved between contending signals from this technique is superior to the separation achieved using a pair of practical directional antennas.

C-DCA fully exploits two degrees of separation offered by directional antennas and multiple channels to maximally isolate contending transmissions. The PDR and throughput of C-DCA are the highest and the delay the lowest among all schemes across all traffic scenarios. In particular, the throughput gains achieved in comparison to OCA are up to 128% with up to ten times lower delay. On one hand, C-DCA avoids the sidelobe problems with multiple interfaces in DNOCA through frequency separation. On the other hand, C-DCA reduces the high channel load problem in OCA through spatial separation. Finally, C-DCA can account for interference from sidelobes of nearby transmitting antennas through measurements.

As described in Section II-C1, C-DCA is conservative in reusing channels. In contrast, the A-DCA algorithm aggres-




Fig. 3. Channel assignment issues in A-DCA and M-DCA. (a) A-DCA side-lobe interference. (b) M-DCA estimation problem.

sively exploits spatial separation to reuse channels which can potentially reduce the channel usage. However, the results show that A-DCA does not offer significant benefits. The performance of A-DCA is always lower than that of C-DCA and in fact comparable to that of OCA. In fact, A-DCA outperforms OCA only when the traffic exceeds 3 Mb/s when the average channel load in OCA increases significantly. The performance degradation of A-DCA is because A-DCA assumes a flat-topped pattern during geometrical calculations in channel assignment. This pattern has only one main lobe, and the power is kept very low for all other angles outside the main lobe. Since our simulations use real antenna patterns (i.e., with sidelobes and backlobes causing interference), while A-DCA optimistically assumes no interference from sidelobes in order to maximize channel reuse, it does not always assign noninterfering channels correctly. An example of this is shown in Fig. 3(a) in which node C using A-DCA reuses channel  $C_0$  to transmit to node D since the nearby node E is directed away from C. However, sidelobes and backlobes in the antennas on C and E using the same channel  $C_0$  can cause interference, degrading the throughput. We also performed simulations varying the guard angle in A-DCA to tune its aggressiveness. The results, reported in detail in [15], show that highly aggressive channel reuse exploiting fine-grained spatial separation is not feasible with practical directional antennas. The effectiveness of A-DCA with smart antennas that can tightly control the sidelobes is a focus of our future work.

Finally, to evaluate the accuracy of geometrical techniques in C-DCA/A-DCA, we compared them with a pure measurement-based channel assignment scheme, M-DCA. M-DCA accounts for sidelobes since it uses real measurements to infer interference between two nodes, instead of geometry. Surprisingly, Fig. 2 shows that M-DCA achieves worse performance than A-DCA and OCA. The reason for this is that in M-DCA, each node checks if it interferes only with nodes from which it can *receive* packets, that is only from nodes in its communication range, and not in its range of interference. However, even nodes in twice the communication range may interfere because of carrier sensing. For example, in Fig. 3(b), node A can assign channels to reduce interference from nodes B and C within its communication range  $R$ . However, node A cannot identify node I as an interferer (due to it not being able to receive packets from I) and is likely to assign the same channel node I is using, resulting in interference. Such inaccuracies in channel assignment significantly degrade the performance of M-DCA.

Note that such far away nodes could be identified using two-hop packet exchanges and interference estimated using geometric calculations. However, the objective of M-DCA is to study the effectiveness of a pure measurement-based algorithm.

In summary, exploiting two degrees of separation by integrating practical directional antennas with channel assignment significantly improves the performance of WMNs. Further, a combination of measurement-based and geometric techniques to assess interference for channel assignment, as in C-DCA, provides the best performance.

Finally, we also compared how the algorithms fare as the traffic load is increased by increasing the number of traffic sources from 25 to 50. In this case, we found that the gap between C-DCA and OCA becomes even larger. For example, when the maximum traffic is 2 Mb/s, the difference between the throughput of the two schemes is about 15% with 25 sources, but increases to 231% when all 50 nodes are sources. This is because when the number of sources increases, the total traffic in the network increases, and the channels become loaded faster. Hence, OCA has to reuse loaded channels, and this results in performance degradation. On the other hand, C-DCA can still exploit the second degree of separation, achieving higher throughput.

### B. Impact of Available Physical Channels

In this section, we study the impact of the number of available channels on the performance of WMNs. As the number of available channels is reduced, both OCA and C-DCA have to resort to reusing channels. This increases the average contention per channel.

Fig. 4 depicts the performance of C-DCA and OCA, as the number of available channels is varied from 3 to 12. Two observations can be made from the results. First, both OCA and C-DCA experience performance degradation as the number of available channels is reduced, since both OCA and C-DCA have to resort to reusing channels. This increases the average contention per channel and worsens performance. Second, C-DCA performance is less affected than OCA as the available channels are reduced. This is because OCA performance is directly dependent on the number of channels as it has only one degree of separation between contending transmissions. In contrast, C-DCA requires less than 12 channels (11 channels on average) to completely isolate contending flows in most of the scenarios we studied. Even when C-DCA has to reuse channels, the average load per channel is also lower. Thus, in comparison to the 12 channel case, the gains of C-DCA against OCA in PDR, throughput, and delay increase to 173%, 176%, and 676%, respectively, when the number of channels is reduced to six. In summary, C-DCA is more effective in utilizing scarce spectrum resources than OCA.

### C. Additional Results

Apart from the evaluation in this section, we studied additional issues in architecting DMesh which are detailed in [15]. We studied the impact of the antenna beamwidth and found that lower beamwidth antennas (typically larger in size) provide better performance as the number of available channels reduces or the density of nodes in a collision domain increases.



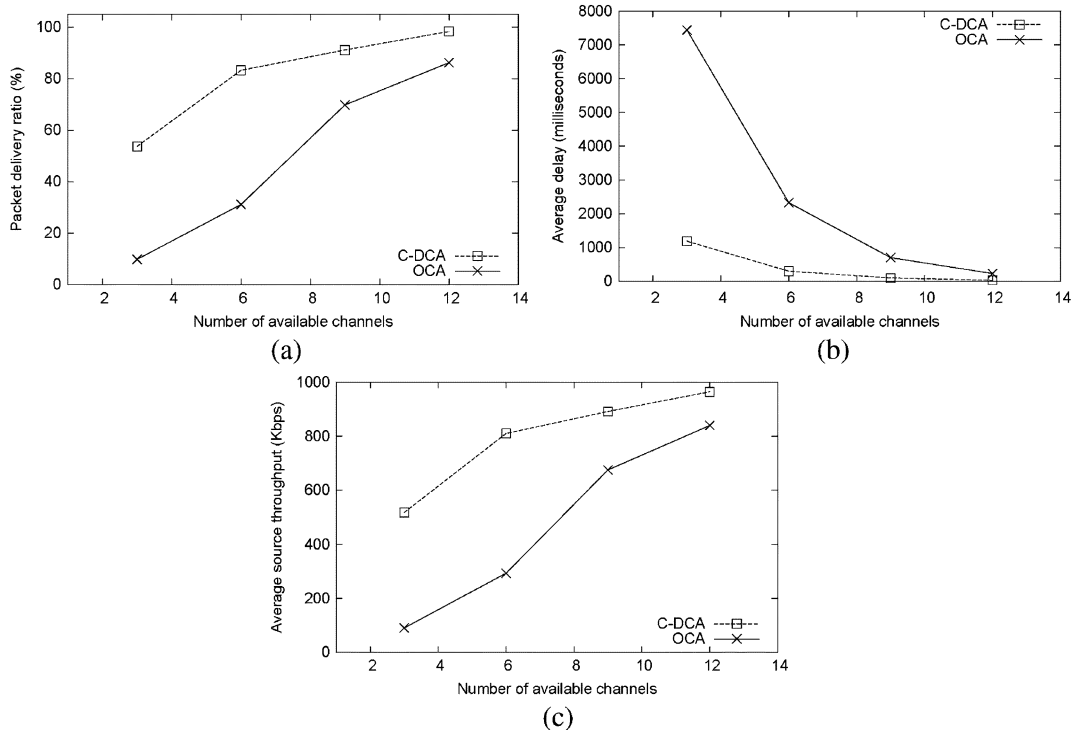


Fig. 4. Impact of a finite number of available channels. Maximum traffic per node is 3 Mb/s. Beamwidth used is  $45^\circ$ . (a) Packet delivery ratio (%). (b) Average delay (ms). (c) Throughput (kb/s).

We also compared C-DCA and OCA in 802.11b networks and found that C-DCA outperforms OCA by about 80%. Finally, we also studied TCP performance and found that DMesh provides TCP throughput gains between 30% and 68% in 802.11a and 35% in 802.11b networks.

## V. TESTBED EVALUATION

In order to verify the usefulness and performance of the DMesh architecture, we incorporated practical directional antennas in MAP (Mesh@Purdue) [16], a wireless mesh network testbed at Purdue. We use C-DCA as the channel assignment algorithm for DMesh. We also compare DMesh with OMesh, the original MAP architecture with only omnidirectional antennas and running the OCA channel assignment algorithm.

### A. Setup

Our testbed consists of 16 wireless mesh routers (small form factor desktops) spread out across three buildings [EE, MSEE and VC (Visitor Center)]. We use Senao Engenius 2511 802.11b wireless cards on the mesh routers. Each radio is attached to a 2 dBi rubber duck omnidirectional antenna with a low loss pigtail to provide flexibility in antenna placement. Each mesh router runs Linux kernel 2.4.20-8 and the open-source *hostap* drivers are used to enable the wireless cards. IP addresses are statically assigned. The wireless cards we use can support a wide range of power settings (up to 200 mW). As in the simulations, the power was adjusted when the cards were used with directional antennas to keep the range similar to the omnidirectional antennas. We used directional antennas in the outdoor part of our network to interconnect the three buildings together. The outdoor portion of our testbed is depicted in Fig. 5.




Fig. 5. Purdue MAP testbed schematic (top view). Only outdoor mesh routers are depicted.

Two of the buildings (EE and MSEE) have two outdoor mesh routers each equipped with two or three radios, while the VC building has a single outdoor mesh router with two radios. These outdoor nodes provide high bandwidth interconnection across the indoor networks in each building. We used directional antennas on these outdoor nodes and compared their performance to using omnidirectional antennas. The directional antennas are the same as those modeled in the simulation. Specifically, we use a 802.11b  $45^\circ$  beamwidth antenna (model HG2412Y from [13]). Note that each mesh router also has an additional omnidirectional control interface.

### B. Evaluation

In this section, we demonstrate the advantages of the DMesh architecture using testbed experiments.

1) *Exploiting Multiple Channels*: The first challenge in the evaluation was to obtain simultaneous throughput in the multi-channel omnidirectional scenario. As a baseline, we found that a single flow between any two nodes without any other interfering flows achieved a TCP throughput in the range of 4.5–5 Mb/s using *netperf*. We then ran two simultaneous flows on interfaces

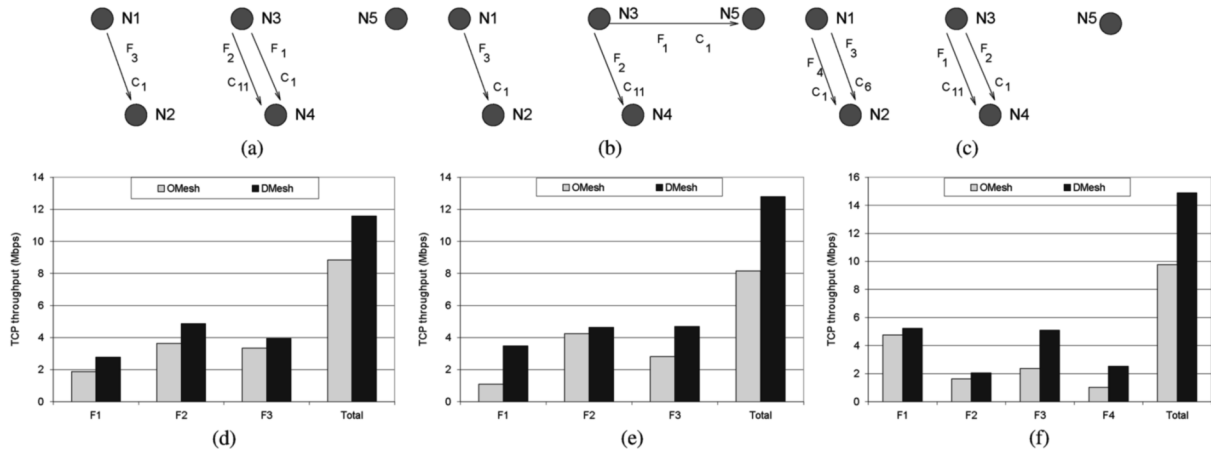


Fig. 6. Testbed results. OMesh is a WMN with omni antennas running OCA, while DMesh is a WMN with practical directional antennas running C-DCA. (a) S1: scenario setup. (b) S2: scenario setup. (c) S3: scenario setup. (d) S1: two channels, three flows. (e) S2: two channels, three flows. (f) S3: three channels, four flows.

$I_1$  and  $I_2$  on nodes  $N_1$  to  $N_2$  (Fig. 5) on maximally separated channels (1 and 11). In this experiment, we found that either one radio achieved close to the full throughput in the single flow case (4.5 Mb/s) with the other card being totally starved and achieving less than 0.5 Mb/s, or both radios achieved around 2–3 Mb/s each. Thus, even though we used external antennas with pigtailed, simultaneous throughput was not possible due to interference among the two radios. We then purchased 12 inch extension cables for each radio and further separated the antennas on both  $N_1$  and  $N_2$ . This resulted in almost simultaneous throughput (with some loss) being achieved of around 4.3 Mb/s each. Thus, antenna separation is critical to enabling multichannel simultaneous throughput and is required on both the receiver and transmitter. In conclusion, to exploit even maximally frequency separated transmissions required antenna isolation. We also noticed that using power control to have a higher power on one interface adversely affected the other interface despite antenna separation. We thus use similar transmission powers on all interfaces of a single mesh router.

2) *Evaluating the Gain From DMesh:* We used three scenarios depicted in Fig. 6(a)–(c), each in a different topology and traffic setting, to demonstrate the throughput gains from DMesh, compared with OMesh. In scenario S1, we assume that the number of channels allowed to be used is two. This is because the Purdue Wireless LAN (Airlink) operates on the remaining orthogonal 802.11b channel in the area, and we did not want to disrupt the other users’ performance. There are three flows, as shown in Fig. 6(a). Both OMesh and DMesh automatically assign separate channels on  $N_3$  based on the OCA and C-DCA algorithms. Node  $N_1$  chooses channel 1 to assign to its flow  $F_3$ . The performance of the flows in this scenario for OMesh and DMesh are shown in Fig. 6(d). The results show that for OMesh, the total throughput is 8.8 Mb/s. This occurs because  $F_1$  and  $F_3$  contend with each other causing a reduction in their individual throughputs. We also observe that unfairness exists, since  $F_3$  achieves a higher fraction of the bandwidth than  $F_1$ . Interestingly, the throughput of  $F_2$  also reduces from its maximum individual throughput although it is on a separate channel. In contrast, using DMesh,  $F_2$  achieves close its individual maximum throughput. The directivity reduces the interchannel inter-

ference from other nodes. In addition,  $F_1$  and  $F_3$  improve their throughputs due to increased spatial separation. Overall DMesh provides a 31% increase in TCP throughput. This result is close to the 35% gain observed overall in simulation results for TCP performance.

Note that although DMesh improves the throughput, the practical directional antennas do not completely isolate the transmissions spatially. For example, even in DMesh,  $F_1$  and  $F_3$  do not obtain their maximum individual throughput which is on average 4.5 Mb/s. Thus, the gain from DMesh can be attributed to a reduction in interference power levels since the interference between  $F_1$  and  $F_3$  is limited to the power received from their sidelobes. This reduced interference is enough to provide significant throughput gains to justify the use of such practical directional antennas in WMNs.

The pattern of the antenna used in this experiment has a very low gain in the backlobes in comparison to the sidelobes and the main lobe. To verify whether further gains are possible if the back lobe is facing the interfering flow, we used scenario 2, Fig. 6(b). Similar to the previous scenario, we again assume that only two channels are available. In this scenario, although  $F_1$  and  $F_3$  are still on the same channel, the antenna used for  $F_1$  is facing the other way with its back lobe towards the antenna for  $F_3$ . The results in Fig. 6(e) show that OMesh performs slightly worse than in scenario 1 and achieves an aggregate throughput gain of around 8 Mb/s. One possible explanation is that  $F_1$  now operates over a longer distance [and thus lower signal-to-noise ratio (SNR)] and is starved by  $F_3$  more in comparison to scenario 1. As expected, DMesh now performs better, providing a 57% TCP throughput gain over OMesh, since the transmitting antenna for  $F_1$  only interferes using the backlobe with  $F_3$ . The receiving antenna for  $F_1$  whose main lobe is towards  $F_3$  could potentially interfere with  $F_3$ . However, the power received and thus the interference are lowered since this antenna is 380 ft away and only transmits short TCP ACK and MAC layer ACK packets. Note that despite the low gain backlobes, the transmitting antenna for  $F_1$  is not totally isolated from the antennas involved in  $F_3$ . Nodes  $N_1$ ,  $N_2$ , and  $N_3$  can still ping each other on the directional interfaces although with a packet loss rate of close to 50%.

In the third scenario Fig. 6(c), we used all three 802.11b orthogonal channels in experiments performed late at night so as not to interfere with Airlink users. There are four flows depicted in the figure. Again, we find that OMesh provides an aggregate throughput of less than 10 Mb/s Fig. 6(f). Flows  $F_2$  and  $F_4$  contend since OMesh runs out of channels. Although DMesh also runs out of channels, it significantly reduces interference power between  $F_2$  and  $F_4$  and provides an aggregate throughput of close to 15 Mb/s resulting in a 50% TCP throughput gain.

In summary, the testbed evaluation shows that DMesh is a viable and useful architecture to improve the performance of WMNs. Although practical directional antennas do not provide *complete* spatial isolation, the reduced interference power provides significant throughput gains. In particular, the reduced interference is useful for multiple-channel networks to obtain close to their maximum throughput potential. Costlier array-based antennas could potentially improve the spatial isolation with tightly controlled sidelobes and backlobes. However, such antennas are still not cost-effective. The cost-benefit tradeoffs in using more advanced antennas in comparison to cheaper versions is an interesting topic of our future research. As a side note, our experiences show that using practical directional antennas was not inconvenient. Their sizes are reasonable and orienting the antennas perfectly towards each other was not critical to performance.

## VI. RELATED WORK

Many design issues of WMNs have been studied [2], [17]–[19] and many companies are offering products for deploying WMNs [4], [20]. The use of multiple radios has been proposed as a means to increase the throughput of wireless networks [1], [2], [18], [21], [22]. Multiple beamformed antennas were previously proposed in a position paper [17] for the TAPs architecture. In contrast to DMesh, TAPs proposes to use significantly costlier smart beamforming antennas but may provide higher spatial reuse. We position DMesh as an architecture to quickly and cheaply deploy high throughput mesh networks with minimal cost using commodity antennas combined with channel assignment. The work in [23] and [24] uses multiple directional antennas. However, their technique does not exploit frequency separation and is designed for situations where only a single channel is available.

Many MAC layer solutions [25]–[29], as well as routing layer solutions [30] have been proposed to exploit multiple channels in wireless networks. The work in [1], [11], [18], and [22] applies multichannel techniques to mesh networks. A recent work [1] proposes a WMN architecture with multiple interfaces and utilizes multiple channels to improve throughput. However, their architecture provides only one degree of separation (frequency) as it uses omnidirectional antennas. In contrast, DMesh combines cheap directional antennas with multiple channels to further improve the performance of WMNs. Two recent works [1], [31] study *joint* routing and channel assignment for multi-radio, multichannel omnidirectional WMNs. In contrast to these approaches, DMesh performs directional channel assignment and decouples routing from channel assignment (similar to in [2]). The work in [2] computes routes for a network in which channel assignment has been previously done; while DMesh computes channel assignments for a routing tree that has been physically formed due to placement of antennas.

The use of beamforming antennas have been proposed for mobile ad hoc networks [9], [10], [32]–[34]. Work has also been done on combining channel assignment and sectorized separation of users in cellular networks (e.g., [35]). However, the multihop routing and nonexistence of infrastructure devices such as base stations make WMNs very different from cellular networks. Thus, cellular network techniques cannot be directly applied to WMNs.

## VII. CONCLUSION

In this paper, we proposed DMesh, a novel architecture for improving the performance of wireless mesh networks. Such an architecture exploits multiradio, multichannel nodes in the mesh network, where each interface is equipped with a practical directional antenna. We also proposed a distributed algorithm to perform routing and directional channel assignment in the DMesh architecture. By exploiting the spatial separation offered by directional antennas and the frequency separation offered by multiple nonoverlapping channels, our proposed architecture allows more concurrent transmissions than an omnidirectional, multiradio, multichannel mesh network, and as a result achieves higher throughput. In particular, simulation results and evaluation on a mesh network testbed show that, compared with the omnidirectional, multiradio, multichannel mesh network, our proposed architecture improves packet delivery ratio and throughput and drastically lowers average per-packet delay. In conclusion, DMesh provides higher throughput for WMNs while remaining cost-effective and easy to deploy. DMesh can be used to naturally extend current widely deployed single radio mesh networks.

## ACKNOWLEDGMENT

The authors thank the anonymous reviewers for their helpful comments.

## REFERENCES

- [1] A. Raniwala and T. Chiueh, "Architectures and algorithms for an IEEE 802.11-based multi-channel wireless mesh network," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 2223–2234.
- [2] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks?" in *Proc. ACM MobiCom*, Sep. 2004, pp. 114–128.
- [3] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proc. SIGCOMM*, Aug. 2004, pp. 121–132.
- [4] Bel Air Networks. [Online]. Available: <http://www.belairnetworks.com>
- [5] R. Chandra, L. Qiu, K. Jain, and M. Mahdian, "Optimizing the placement of Internet taps in wireless neighborhood networks," in *Proc. ICNP*, Oct. 2004, pp. 271–282.
- [6] T. Clausen, P. Jacquet, C. Adjih, A. Laouiti, P. Minet, P. Muhlethaler, A. Qayyum, and L. Viennot, "Optimized link state routing protocol (OLSR)," RFC 3626, Oct. 2003.
- [7] D. S. J. D. Couto, D. Aguayo, J. C. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. ACM MobiCom*, 2003, pp. 134–146.
- [8] QualNet. [Online]. Available: <http://www.scalable-networks.com>
- [9] M. Takai, J. Martin, A. Ren, and R. Bagrodia, "Directional virtual carrier sensing for directional antennas in mobile ad hoc networks," in *Proc. ACM MobiHoc*, Jun. 2002, pp. 183–193.
- [10] R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya, "Using directional antennas for medium access control in ad hoc networks," in *Proc. ACM MobiCom*, Sep. 2002, pp. 59–70.
- [11] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted seeded channel hopping for capacity improvement in IEEE 802.11 ad-hoc wireless networks," in *Proc. ACM MobiCom*, Sep. 2004, pp. 216–230.

- [12] J. Jun, P. Peddabachagari, and M. Sichertiu, "Theoretical maximum throughput of IEEE 802.11 and its applications," in *Proc. NCA*, 2003, pp. 249–256.
- [13] *HyperLink-Technologies*, [Online]. Available: <http://www.hyperlink-tech.com>
- [14] MAXRAD. [Online]. Available: <http://www.maxrad.com>
- [15] S. M. Das, H. Pucha, D. Koutsonikolas, Y. C. Hu, and D. Peroulis, "DMesh: Incorporating practical directional antennas in wireless mesh networks," Purdue Univ., West Lafayette, IN, Aug. 2005, Tech. Rep., TR-ECE-05-16.
- [16] Mesh@Purdue. [Online]. Available: <http://www.engineering.purdue.edu/MESH>.
- [17] R. Karrer, A. Sabharwal, and E. Knightly, "Enabling large-scale wireless broadband: The case for TAPs," *ACM CCR*, vol. 34, no. 1, pp. 27–32, 2004.
- [18] A. Raniwala, K. Gopalan, and T. Chiueh, "Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks," *ACM MC2R*, vol. 8, no. 2, pp. 50–65, Apr. 2004.
- [19] V. Gamberoza, B. Sadeghi, and E. W. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *Proc. ACM MobiCom*, 2004, pp. 287–301.
- [20] Mesh Networks. [Online]. Available: <http://www.meshnetworks.com>
- [21] P. Bahl, A. Adya, J. Padhye, and A. Wolman, "Reconsidering wireless systems with multiple radios," *ACM MC2R*, pp. 39–46, Oct. 2004.
- [22] P. Kyasanur and N. H. Vaidya, "Routing and interface assignment in multi-channel multi-interface wireless networks," Univ. Illinois at Urbana-Champaign (UIUC), Urbana, IL, Tech. Rep., Oct. 2004.
- [23] B. Raman and K. Chebrolu, "Revisiting MAC design for an 802.11-based mesh network," in *Proc. ACM HotNets*, Nov. 2004, to be published.
- [24] —, "Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks," in *Proc. ACM MobiCom*, Aug. 2005, pp. 156–169.
- [25] Y. Liu and E. Knightly, "Opportunistic fair scheduling over multiple wireless channels," in *Proc. IEEE INFOCOM*, 2003, pp. 1106–1115.
- [26] J. So and N. Vaidya, "Multi-channel MAC for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver," in *Proc. ACM MobiHoc*, 2004, pp. 222–233.
- [27] A. Tzamaloukas and J. J. Garcia-Luna-Aceves, "A receiver-initiated collision-avoidance protocol for multi-channel networks," in *Proc. IEEE INFOCOM*, 2001, pp. 189–198.
- [28] A. Nasipuri and S. Das, "A multichannel CSMA MAC protocol for mobile multihop networks," in *Proc. IEEE WCNC*, 1999, pp. 1402–1406.
- [29] S.-L. Wu, C.-Y. Lin, Y.-C. Tseng, and J.-P. Sheu, "A new multi-channel MAC protocol with on-demand channel assignment for multi-hop mobile ad hoc networks," in *Proc. I-SPAN*, 2000, p. 232.
- [30] J. So and N. H. Vaidya, "A routing protocol for utilizing multiple channels in multi-hop wireless networks with a single transceiver," Univ. Illinois at Urbana-Champaign (UIUC), Urbana, IL, Tech. Rep., Oct. 2004.
- [31] M. Alicherry, R. Bhatia, and L. E. Li, "Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks," in *Proc. ACM MobiCom*, Aug. 2005, pp. 58–72.
- [32] Y. Ko and N. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *Proc. IEEE INFOCOM*, Mar. 2000, pp. 13–21.
- [33] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," in *Proc. ACM MobiHoc*, Oct. 2001, pp. 95–105.
- [34] A. Nasipuri, J. Mandava, H. Manchala, and R. Hiromoto, "On-demand routing using directional antennas in mobile ad hoc networks," in *Proc. IEEE ICCN*, Oct. 2000, pp. 535–541.
- [35] J.-L. Pan, S. S. Rappaport, and P. M. Djuric, "Multibeam cellular communication systems with dynamic channel assignment across multiple sectors," *Wireless Netw.*, vol. 5, no. 4, pp. 267–278, 1999.



**Saumitra M. Das** (S'99) received the B.Eng. degree from the University of Bombay, Bombay, India, and the M.S. degree from Carnegie Mellon University, Pittsburgh, PA. He is currently working towards the Ph.D. degree at the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN.

His research interests include cross-layer system design for multihop wireless networks, scalable routing strategies in wireless ad hoc networks, and mobile robotics.



**Himabindu Pucha** (S'00) received the B.Eng. degree from the University of Bombay, Bombay, India, and the M.S.E.E. degree from Purdue University, West Lafayette, IN. She is currently working towards the Ph.D. degree at the School of Electrical and Computer Engineering, Purdue University.

Her research interests include Internet routing and overlay networks, peer-to-peer systems, and mobile computing.

Ms. Pucha received the Violet Haas Award from Purdue University in 2005 and the Google Anita Borg Memorial Scholarship in 2006.



**Dimitrios Koutsonikolas** (S'04) received the B.S. degree from the School of Electrical and Computer Engineering, National Technical University of Athens (NTUA), Athens, Greece, in 2004. He is currently working towards the Ph.D. degree at the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN.

His research interests include multi-hop wireless ad hoc, mesh, and sensor networks.

Mr. Koutsonikolas has received a Tellabs Fellowship and a Ross Fellowship from Purdue University.



**Y. Charlie Hu** (S'90–M'03) received the M.S. and M.Phil. degrees from Yale University, New Haven, CT, in 1992 and the Ph.D. degree in computer science from Harvard University, Cambridge, MA, in 1997.

He is an Assistant Professor of Electrical and Computer Engineering and Computer Science at Purdue University, West Lafayette, IN. From 1997 to 2001, he was a Research Scientist at Rice University. His research interests include operating systems, distributed systems, networking, and parallel computing. He has published over 70 papers in these areas.

Dr. Hu is a member of USENIX and the Association for Computing Machinery (ACM). He received the Honda Initiation Grant Award in 2002 and the NSF CAREER Award in 2003. He served as a TPC Vice Chair for the 2004 International Conference on Parallel Processing, and a co-founder and TPC Co-Chair for the International Workshop on Mobile Peer-to-Peer Computing.



**Dimitrios Peroulis** (S'99–M'04) received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1998, and the M.S.E. and Ph.D. degrees in electrical engineering from The University of Michigan, Ann Arbor, in 1999 and 2003, respectively.

Since August 2003, he has been an Assistant Professor at the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN. His current research work is focused on microelectromechanical systems (MEMS) for multifunctional communications systems and sensors.

Dr. Peroulis has received several teaching awards from Purdue University. He has also been the recipient of three Student Paper Awards at IEEE MTT-S (2001 and 2002) and IEEE AP-S (2001).