A Performance Study of Deployment Factors in Wireless Mesh Networks

Joshua Robinson and Edward W. Knightly
Department of Electrical and Computer Engineering
Rice University, Houston, TX 77005
Email: jpr, knightly@rice.edu

Abstract—We present a measurement-parameterized performance study of deployment factors in wireless mesh networks using three performance metrics: client coverage area, backhaul tier connectivity, and fair mesh capacity. For each metric, we identify and study topology factors and architectural features which strongly influence mesh performance via an extensive set of Monte Carlo simulations capturing realistic physical layer behavior. Our findings include: (i) A random topology is unsuitable for a large-scale mesh deployment due to doubled node density requirements, yet a moderate level of perturbations from ideal grid placement has a minor impact on performance. (ii) Multiple backhaul radios per mesh node is a cost-effective deployment strategy as it leads to mesh deployments costing 50% less than with a single-radio architecture. (iii) Dividing access and backhaul connections onto two separate radios does not use the second radio efficiently as it only improves fair mesh capacity 40% to 80% for most users. This is in contrast to using the second radio to move half the user population to a new network operated on the second radio. This work adds to the understanding of mesh deployment factors and their general impact on performance, providing further insight into practical mesh deployments.

I. INTRODUCTION

Wireless mesh networks are a cost-effective last-mile access network, providing wireless Internet service to a large coverage area with low infrastructure cost [1]. As a result, many cities and ISPs plan to deploy two-tier mesh networks for city-wide Internet access [2]. A two-tier mesh network consists of an access tier, which provides connectivity to client devices, and a backhaul tier, which forwards traffic among mesh nodes to a wired Internet gateway. Current strategies for mesh network planning include three approaches: exhaustive site surveys to find optimal node placements [3], unplanned deployments [4], and reliance on general rules-of-thumb [5].

In this paper, we study mesh deployment by identifying key mesh topology factors and architectural features. We quantify their impact on performance with three metrics: client coverage area (characterizing the access tier), backhaul tier connectivity (capturing the ability to connect to a wired gateway node via *any* route), and fair mesh capacity (capturing idealized fair rates at the gateway node). Our methodology is to design and analyze a set of fractional factorial [6] Monte Carlo experiments in order to isolate and study each factor's influence. Our approach contrasts with prior work in that we consider the general impact of mesh topology factors and architectural features. For example, in [7], [8], optimal locations of WLAN nodes are determined as a function of

path-loss measurements for all location pairs; likewise, in [9], optimal wired gateway placement is determined as a function of the traffic matrix.¹

Our study first examines the coverage area of mesh networks. We find that while moderate perturbations from ideal grid placement do not significantly degrade coverage area, a random deployment requires $2\times$ more mesh nodes to achieve the same coverage target. The loss in coverage area due to randomness in node placement results from an increase in coverage dead spots, i.e. locations with little or no probability of connection. Additionally, we show that the hexagonal grid topology results in more coverage dead spots than a square or triangular grid and therefore requires twice the node density to achieve worst-case coverage guarantees.

Second, we define the average mesh node connectivity to study the availability of backhaul tier routes. Here, we find that random node placement only slightly degrades connectivity, exhibiting a 10% reduction at high densities. Further, random networks of moderate node density feature a higher probability of secondary, non-overlapping routes due to a fraction of the links having higher signal strength than any links in a grid topology. We show that using multiple *backhaul* radios per mesh node is a cost-effective solution for achieving connectivity targets, reducing the total network cost by up to 50%. Multiple radios reduce network cost by allowing significantly fewer mesh nodes to be deployed with only fractionally greater cost per mesh node. Also, we find that the marginal gain of systems with three or more radios is minimal.

Third, we calculate the ideal fair mesh capacity, i.e. the aggregate throughput at a wired gateway under per-user fairness constraints. We show that separating the access and backhaul tiers with a second radio is *not* an efficient use of a second radio, as users in the network experience a fair capacity improvement of *less* than double. This configuration does not fully take advantage of the second radio in the case where spatial reuse already allows some access links to operate without interfering with the wired gateway nodes. Additionally, we find that a random network provides less than half the fair capacity of a regular grid topology due to increased contention for wireless airtime at wired gateway nodes and poor coverage area. Consequently, random networks are not suitable for a large-scale mesh deployment.

¹Additional related work is discussed in Section VI.

The remainder of this paper is organized as follows. In Section II, we describe two-tier mesh networks and our investigation methodology. In Sections III, IV, and V we introduce and study our metrics: coverage area, connectivity, and fair mesh capacity. Section VI contrasts our work with existing literature, and finally Section VII concludes.

II. BACKGROUND

Two-tier mesh networks (Fig. 1) consist of a backhaul tier for interconnection between mesh nodes and an access tier for connection between infrastructure mesh nodes and client devices. A fraction, w , of these infrastructure nodes feature a wired connection to the Internet, which we refer to as wired gateway nodes. Mesh nodes communicate using omni-directional antennas across multihop paths to the wired gateways.

Fig. 1. A two-tier mesh network consists of infrastructure mesh nodes, which forward traffic to a gateway, and access nodes, which source and sink traffic.

A. Physical Layer for Mesh Links

The physical layer model used to determine the measurement-parameterized link behavior is a well-known pathloss model [10], which relates link distance to an expected signal strength on that link. The link behavior equations rely on a *pathloss exponent* α and a *shadowing standard deviation* σ_ϵ to accurately model link behavior, as described next.

The average received signal strength P_{dBm} is given as a function of distance d , a reference power level at distance d_0 , the pathloss exponent α , and a random shadowing term ϵ .

$$P_{dBm}(d) = P_{dBm}(d_0) - 10\alpha \log_{10}\left(\frac{d}{d_0}\right) + \epsilon \quad (1)$$

The shadowing term ϵ reflects the scattering environment and is represented as a zero-mean Gaussian random variable with standard deviation of σ_ϵ . The link's signal strength is represented as a random variable X and the probability that the average signal strength exceeds a threshold value T_{min} for a link of distance d is

$$Pr_d[X > T_{min}] = Q\left(\frac{T_{min} - P_{dBm}(d)}{\sigma_\epsilon}\right) \quad (2)$$

where $P_{dBm}(d)$ is expressed in Equation (1), σ_ϵ is the shadowing standard deviation, and the Q function is the complement

of the CDF of a standard Gaussian. The threshold value T_{min} represents the minimum acceptable average signal strength. Finally, to determine the achievable rate on a given link with Equation (1), we employ a measurement-based mapping from signal strength to expected link-layer throughput.

For our Monte Carlo simulations, we use the measured values of $\alpha = 3.27$, $\sigma_\epsilon = 5.94$ dBm, and $T_{min} = -75$ dBm from our case-study network [11], an outdoor 802.11b residential deployment in a Houston neighborhood. Additionally, we utilize the measured mapping between signal strength and achievable throughput from the same case-study network. The parameterization of physical layer behavior enables this study to extend to different pathloss environments and hardware architectures.

B. Methodology

This paper presents an extensive set of performance results for each metric over a wide range of mesh topologies. Our fractional factorial experiment design [6] focuses on each deployment factor's primary effects on performance as well as secondary effects wherein we jointly vary two factors. Dominant deployment factors and their effects appear in this paper. We employ Monte Carlo techniques to simulate the necessary topology realizations, as well as the link pathloss realizations for the connectivity metric. For each realization, we then solve for the metric using equations presented in the following sections. For this study, we focus on the performance of interior mesh nodes as the majority of deployed nodes in a moderately sized mesh network will be interior. We factor out edge performance effects by allowing edge nodes to participate in routes, but we do not report their performance results.

Fig. 2. The three regular tessellations shown where the mesh nodes are represented as triangles and links between mesh nodes as dashed lines.

We consider three regular tessellations as our baseline grid topology : triangular, square, and hexagonal (see Fig. 2). The random mesh topologies are modeled as homogeneous Poisson point processes. We also consider randomness as perturbation to a grid topology, reflecting the reality that perfect grid placement cannot usually be achieved due to constraints of the physical environment. Perturbation is modeled by displacing each mesh node by a random distance and angle. Further, we evaluate mesh node architectural choices, wherein we capture the use of multiple radios per mesh node, each able to operate concurrently.

III. ACCESS TIER COVERAGE

The first performance metric, *coverage area*, is the probability that an arbitrary client device connects to the mesh

network. In this section, we define the coverage area metric and examine the coverage area of networks with regular and random topologies. We also compute lower bounds for coverage area and then study the impact of mesh node perturbations on coverage area.

A client device connects to a mesh node if the average signal strength received from the mesh node is above a threshold, T_{min} . Therefore, to find the coverage area of a topology, we calculate the equivalent probability that it is not the case that a client location is unable to connect to any nearby mesh nodes:

$$\text{Coverage} = 1 - \prod_{\forall i} (1 - Pr_{d_i}[X > T_{min}])$$

where i represents each mesh node in the network and Pr_{d_i} is found with Equation (2).

We compute the coverage area as the probability of coverage for an arbitrary client location by averaging the coverage probabilities of all client locations for each topology realization. The pathloss exponent, $\alpha = 3.7$, for our experiments is the measured access tier pathloss in our case-study deployment [11]. The connection threshold, $T_{min} = -75$ dBm, corresponds to an average rate of 2 Mbps, which is the peak level of service provided in the case-study network.

A. Topology Factors

We first show the coverage gain of regular topologies over randomly deployed topologies before deriving worst-case coverage probabilities. Fig. 3 plots the coverage area as a function of node density for a regular square grid topology² and an unplanned network. The regular topology provides up to 20% higher coverage on average. A network operator that wants to target a coverage area of 95% will need to deploy approximately 20 nodes per km² in a regular grid. An unplanned deployment, though, requires almost double (39) the number of mesh nodes and therefore nearly double the deployment costs.

We also consider the coverage area of the three regular tessellations in Fig. 3. All three tessellations provide approximately equal coverage, with the hexagonal performing minimally worse. In order to understand the difference in coverage area between topologies, we next examine the impact of dead spots in coverage.

1) *Coverage Holes*: Fig. 4 plots the ratio of coverage holes as a function of mesh node density, where a *coverage hole* is a client location with less than a 50% probability of being connected to a mesh node. In other words, the fraction of coverage holes is the percentage of client locations that are more likely to be disconnected than not.

From Fig. 4 it is clear that the differences in coverage area are due to the relative fraction of locations that are poorly covered. In other words, the reduction in coverage is not uniform across the whole network, but instead there exist problem areas with little or no client coverage. The small differences in coverage area for the three regular topologies

²Monte Carlo techniques are not required for deterministic grid topologies because a client's coverage probability can be computed directly.

Fig. 3. Coverage area as a function of mesh node density for three regular topologies and random networks.

Fig. 4. Fraction of client locations that have lower than 50% chance of being connected to a mesh node, plotted as a function of mesh node density.

directly correspond to the relative fraction of coverage holes in those networks.

2) *Worst-Case Coverage*: To further study these coverage holes, we next consider the worst-case coverage probabilities. For the regular grid scenario, the worst coverage is at the center of each regular polygon (square, triangle, or hexagon). We express the worst-case distance from client to the nearest mesh node in each tessellation as m_i . i represents the tessellation: triangular ($i = 3$), square ($i = 4$), or hexagonal ($i = 6$).

$$m_4 = \frac{\sqrt{2}}{2} \times \frac{1000}{\sqrt{D}} = \frac{707.1}{\sqrt{D}}$$

The first term is the ratio of the distance between mesh nodes to the distance to the midpoint of the polygon. The second term is the polygon edge lengths as a function of the mesh node density D . Likewise, $m_3 = \frac{620.4}{\sqrt{D}}$ and $m_6 = \frac{1059.1}{\sqrt{D}}$.

The worst-case connectivity of tessellation i is a function of the worst-case client distance m_i and the number of mesh nodes m_i meters from the worst-case client location. Specifically, coverage of configuration i is $1 - (1 - Pr_{m_i})^i$,

where Pr_{m_i} is the probability of connection for a link of distance m_i as given in Equation (2).

Fig. 5. Worst-case coverage probability for each of the three regular tessellations as a function of mesh node density.

Fig. 5 plots the worst-case coverage probability. The hexagonal tessellation exhibits worst-case coverage probability significantly lower than the other two configurations. For a client location, the hexagonal configuration presents the greatest distance, m_6 , to a mesh node, and this is not offset by the fact that there are 6 mesh nodes at that distance. The hexagonal topology therefore requires twice as many mesh nodes to provide worst-case coverage guarantees.

We also consider the worst-case coverage of the random network. If K is the random variable for the number of mesh nodes in a circle of radius r , then probability that k mesh nodes are within radius r is [12]

$$Pr[K = k|r] = \frac{(\lambda\pi r^2)^k}{k!} e^{-\lambda\pi r^2} \quad (3)$$

Therefore, the percentage of time a client device in a random network has worse coverage probability than the worst-case in a regular network is the probability that no mesh nodes are within a circle of the appropriate radius.

The appropriate radius is found by observing that 90% of client locations connect to the nearest mesh node and therefore we assume that the coverage probability depends only on the distance to the nearest mesh node. We then find the probability in a random network that no mesh nodes are within a circle of radius corresponding to the worst-case distance in a regular network. Using this radius in conjunction with Equation (3), gives that for all mesh node densities, 20% of the client devices in a random network are farther away from a mesh node than the worst-case distance in a regular square grid of the same density. We also solve the triangular and hexagonal cases and find probabilities of 30% and 3% respectively. The coverage area of an arbitrary client location in a random network is thus shown to be lower than the worst-case coverage area of a regular network up to 30% of the time. Therefore, worst-case coverage probability indicates that the triangular and square grid topologies are most suitable for providing mesh coverage.

B. Grid Topology Perturbations

In most mesh network deployment scenarios, the mesh operator does not have complete control over the placement of mesh nodes. We next examine the impact of this realistic deployment scenario by introducing random perturbations to mesh node topologies (see Fig. 6).

Fig. 6. Illustration of random mesh node perturbations from an ideal grid topology.

Fig. 7 plots coverage area as average perturbation distance increases for the three regular tessellations. The density is 19.75 mesh nodes per km^2 , which corresponds to a square grid inter-node spacing of 225 meters. The coverage area declines as the perturbations increase, though only slightly at lower mesh node densities. With an average perturbation of 45 meters, or approximately $\frac{1}{5}$ the inter-node distance in the square grid topology, the coverage area decreases by only 2%. Note that the three regular configurations have similar coverage area regardless of perturbations.

Fig. 7. Coverage area as a function of average perturbation distance for square, triangular, and hexagonal configurations. The average coverage area for random topologies is also provided as a reference.

Fig. 7 also includes the average coverage area of a random network of the same density.³ As the average perturbations approach the inter-node distance, 225 meters for the square grid case, the coverage area also approaches that of the random network. Further experiments show that average perturbations of twice the inter-node spacing result in coverage equivalent to a random topology.

³The line is constant because we did not perturb the random networks.

IV. BACKHAUL TIER CONNECTIVITY

In this section, we introduce the *backhaul tier connectivity* metric to study the reliability of multihop routes from mesh nodes to wired gateways. We examine the connectivity of grid and random topologies, focusing on the existence of backup routes and the link symmetry assumption. We then quantify the connectivity gain due to additional backhaul radios.

Backhaul tier connectivity is the average probability that a mesh node has at least one path to a wired node in which each link's average signal level satisfies the minimum threshold, T_{min} . To calculate the connectivity of a mesh node, we evaluate the probability that a route from mesh node to a wired gateway exists (i.e. is usable at a desired performance level) as the inverse of the probability that no routes exist from the mesh node to any wired gateway. In other words, we find the complement of the probability that all routes fail to meet minimum quality requirements.

A. Topology Factors

Fig. 8 compares the considered grid and random topologies. The connectivity of random networks is equivalent to the grid networks at low and moderate mesh node densities. Only for high density networks do structured topologies outperform the average random topology. The random topologies do exhibit an order of magnitude greater standard deviation for the connectivity of individual mesh nodes. At low densities the randomness translates to a few, well-connected mesh nodes increasing the average connectivity, whereas at high densities the randomness inversely results in regions of low connectivity.

Fig. 8. Connectivity as a function of mesh node density for grid topologies and random topologies.

1) *Redundant Routes*: For fault-tolerance and multipath routing, it is important for a mesh node to have a second, non-overlapping route to a wired gateway. Fig. 9 plots the probability that a connected mesh is *also* connected via a secondary route. At low mesh densities, connected nodes are usually well connected (with a second path) and are only those nodes close to a wired gateway. Interestingly, at moderate densities, the random topology results in a greater number of redundant routes. This occurs because the random

structure features a fraction of higher quality links than any link in a regular topology. At higher mesh densities, all four mesh configurations automatically provide redundant routes. A further important distinction is that the random topology cannot provide connectivity guarantees due to the high standard deviation.

Fig. 9. Probability of a connected mesh node also able to connect via a second, non-overlapping route.

2) *Asymmetry*: Due to the presence of heterogeneous scatterers in an environment, link propagation is not always symmetric [13], i.e. the link shadowing is different for both directions. In Fig. 10, we consider two routing mechanisms for the asymmetric case: a strict protocol requiring identical upload and download paths and a second protocol allowing disjoint paths. The connectivity is lowest with the strict asymmetric routing because a useful upstream route no longer guarantees a useful downstream route. To highlight the importance of this assumption, we find the needed mesh node density to achieve 90% connectivity. The strict asymmetric results lead to 30% greater mesh node density to achieve this goal than the symmetric model. But with a less restrictive routing protocol, the target connectivity requires only 11% greater mesh node density than with the purely symmetric link assumption.

B. Multiple Backhaul Radios

Previously, we considered mesh nodes with one radio each and therefore one potential link per neighboring mesh node. This section now investigates installing multiple backhaul radios per mesh node. To model multiple backhaul radios, we assume that each of the R radios is set to the corresponding channel 1 through R . Each node therefore has R links between a neighbor, each with separate antennas and therefore we assume independent channel fading. We choose this channel assignment to maximize connectivity, as opposed to increasing spatial reuse [14].

Fig. 11 presents the relative gain in connectivity over a single-radio architecture for each multi-radio configuration in a square grid topology. A four-radio architecture results in up to five times greater mesh connectivity; yet each additional radio above two results in diminishing marginal gains. The gains

Fig. 10. Average connectivity in square grid topology as a function of mesh node density for both the symmetric and asymmetric link model. Strict asymmetric routing requires the upload and download path to traverse the same sequence of nodes.

for all multi-radio configurations diminish with increasing mesh node density due to the higher-density networks having essentially 100% connectivity and therefore no ability to improve.

Fig. 11. Gain in connectivity as a function of backhaul radios and mesh node density. Gains are relative to the connectivity of a single-radio backhaul.

Each additional radio in a mesh node is an additional resource at that node. Revisiting the economic motivations for mesh networks, we now consider how best to allocate the resources in a network, given that a mesh node and each additional radio has a given cost. The cost metric is normalized to the hardware and installation cost of a single mesh node and therefore does not require absolute cost values. The tradeoff here is the choice between a high density deployment with one radio, and a low density deployment with multiple radios per node. This choice is dependent on the resulting connectivity and the relative costs of radios and mesh nodes.

Fig. 12 plots the relative costs of deploying a multi-radio mesh network in a grid topology with average connectivity of 90%. The costs in Fig. 12 are normalized to the total cost of a single-radio deployment which also achieves the

Fig. 12. Lowest cost of network deployment with target connectivity of 90% for each multi-radio configuration, normalized to the cost of the cheapest single-radio deployment. Wire ratio is $\frac{1}{16}$.

90% connectivity target. A four-radio deployment leads to a network that costs as little as half of the cost of a single-radio configuration. Further, single radio systems are only best when the cost of the additional radio is at least 44% of the cost of a mesh node. Also note that the additional cost-savings of more than two radios diminishes rapidly. As illustration, note that four radios result in no more than 5% cost-savings over a three radio architecture, whereas a two-radio architecture leads to up to 30% cost-savings over a single-radio architecture. The results in Fig. 12 demonstrate the importance of a second backhaul radio for deploying cost-effective networks.

V. FAIR CAPACITY AT WIRED GATEWAYS

The final mesh performance metric, the *fair mesh capacity*, is the rate at which data is exchanged between the wireless mesh network and the wired Internet, given a per-user fairness constraint. This section defines the fair mesh capacity and explores the impact of mesh node density and topology configuration. Further, we study the deployment factor wherein a second radio is installed in each mesh node to handle only access tier traffic.

We express the fair mesh capacity of a wired gateway as the rate at which data flows between the wireless mesh network and the wired Internet. For a single gateway, this corresponds to the fraction of the wireless bandwidth that is used for transmitting data on the wired network. Note that some portion of the wireless bandwidth will not add to mesh capacity because of the time the gateway node must spend deferring to nearby multihop transmissions. This fraction of time corresponds to the normalized fair rate, labeled as δ . The fair mesh capacity is then the sum of the rates through each wired gateway.

To calculate fair mesh capacity, we first assume that all users in the network receive equal time shares of the wired links. The rate at each wired gateway is then determined by the wireless channel bandwidth and the fraction of time the gateway spends in useful transmissions. The total fair mesh

capacity, T , is

$$T = \sum_{j \in W} e_j \delta_j \quad (4)$$

where W is the set of all wired gateways, e_j is the wireless capacity of gateway node j , and δ_j is the backhaul rate of gateway j normalized to the wireless capacity e_j . In other words, δ_j is the fraction of e_j that corresponds to useful data transmitted across the wired interface of node j .

We calculate the normalized backhaul tier rate δ_j as the average utilization of the wireless medium at wired gateway node j . This is proportional to the number of links which interfere with the gateway, as well as the number of active routes which utilize those links. Our computations utilize Monte Carlo techniques to find locally optimal routes in each topology and then calculate the average fraction of time each gateway node must spend deferring. We assume an acknowledgment based MAC protocol, hence the receiving node also transmits and potentially interferes. The fair rates calculated in this section are ideal, assuming a fair multihop MAC protocol, perfect routing protocol, and a scheduling policy able to achieve desired rates.

A. Topology Factors

We begin by examining the fair mesh capacity of different topology configurations. Fig. 13 presents the fair mesh capacity results for the three regular network configurations. The fair mesh capacity of the hexagonal configuration is lowest, but not because of a lower coverage area (see Fig. 5). In fact, the access tier causes small performance degradation only at low mesh densities. The hexagonal configuration has lower fair capacity because the shorter links result in more links interfering with each gateway node. In other words, the hexagonal grid exhibits the worst spatial reuse. For all configurations, the marginal gain of increasing mesh node density decreases due to the exponential increase in interference as gateway nodes begin to overlap.

Fig. 13. Fair mesh capacity as a function of mesh node density and topology configuration.

Wired gateway density is a more critical deployment factor than mesh node density, as illustrated in Fig. 14. If the wired

Fig. 14. Fair mesh capacity as a function of wired gateway density and mesh node density.

density is fixed, adding mesh nodes does not automatically increase fair mesh capacity. The marginal benefit of increasing wired gateway density is an order-of-magnitude greater than the marginal gain due to increasing mesh node density. This result confirms that the availability of wired gateway nodes is the strongest factor determining the fair mesh capacity. Note that in the remainder of this section, we do not present hexagonal and triangular results because they exhibit only small performance deviations from the square topology.

Next, we contrast the fair capacity in random topologies with that of ideal grid-based networks. For random topologies, a fraction of the mesh nodes serve more client nodes than the other mesh nodes. In fair capacity calculations, these nodes receive greater weighting, a technique which extends to uneven user distributions. In Fig. 15, the average fair mesh capacity in a random network is less than half the fair mesh capacity in a grid topology. The difference in access tier coverage does not account for this large gap (recall Fig. 3).

Fig. 15. Average fair mesh capacity for random and square grid topologies as a function of mesh node density.

The additional loss in performance of a random network is due to the gateways in a random network having more than double the wireless resource demands. Further investigation

finds the reason for the increase in wired gateway demand to be that routes in a random network are on average one hop longer than the routes in a square grid network. A second reason for this increase is simply that a gateway in a random network contends with a larger number of links. Note that a grid network leads to less clustering of mesh nodes, where all links mutually interfere.

B. Dual Radio Architecture

We now consider the fair mesh capacity for square grid networks with one and two-radio hardware architectures. First, we calculate the normalized backhaul tier rate, δ , assuming Manhattan routing and two-hop neighbors interfering.

The π_i terms denote the percentage of mesh nodes in the network that are i -hops from a wired gateway. We generally denote the gateway air-resources required by a mesh node i -hops away as a_i . Rewriting δ in Equation (4) gives that the normalized backhaul tier rate per gateway is

$$\delta_j = \frac{1}{\sum_{i=0}^{\infty} a_i \pi_i}$$

The value of a_i corresponds to the number of interfering links traversed by a mesh node i -hops from gateway node j . Intuitively, the above equation gives the average utilization of the wireless medium at gateway j .

First, we examine a mesh network with a single radio for both access and backhaul links. Let the wire ratio $w = \frac{1}{16}$ and use the values for π_i of $\pi_0 = \frac{1}{16}$, $\pi_1 = \frac{1}{4}$, $\pi_2 = \frac{3}{8}$, $\pi_3 = \frac{1}{4}$, and $\pi_4 = \frac{1}{16}$. The values of a_i in this configuration do not exceed 3 as paths longer than three hops exhibit spatial reuse. Therefore $a_0 = 1$, $a_1 = 2$, $a_2 = 3$, and all other $a_i = 3$. The δ_j values of a mesh network with unified access and backhaul tiers are therefore $\frac{8}{23}$. In other words, the backhaul tier in a single-radio configuration is able to transmit or receive useful data 35% of the time.

Next, consider the two-radio architecture with the access tier on a separate radio such that it does not interfere with the backhaul links. The demands a_i decrease because the access links no longer interfere with the backhaul radio at the gateway node. The new values of a_i , which we denote $\hat{a}_i$, are thus $\hat{a}_0 = 0$, $\hat{a}_1 = 1$, and $\hat{a}_2 = 2$. For longer links, spatial reuse maintains $\hat{a}_i$ at a value of 3, gaining no benefit from the second radio. Solving for the normalized backhaul tier rate in a two radio architecture now gives $\hat{\delta}_j = \frac{16}{33}$, which means that the backhaul tier now transmits useful data 48% of the time.

These calculations lead to the conclusion that splitting the access tier onto a separate radio results in a $\frac{46}{33}$ improvement in δ_j . The next results calculate the fair mesh capacity as a function of the normalized backhaul tier rate for more general topologies. Note that the interference threshold is -86 dBm (450 meters on average).

Fig. 16 plots for both access tier configurations the fair mesh capacities for all clients associated with a *wireless* mesh node. We do not include the capacities of the clients associated with wired gateway nodes as their traffic does not traverse the backhaul tier. As a consequence of the access tier being

Fig. 16. Fair mesh capacity with separate and unified access and backhaul tiers. Not considering client capacities associated with wired gateways.

underutilized relative to the backhaul tier, the gateway clients achieve order of magnitude greater fair capacity. We consider approaches for improving the fair capacity for all clients, though.

The results in Fig. 16 demonstrate a 40% to 80% improvement due to the splitting the access tier traffic onto a dedicated radio. Recall the values of a_i above $i = 3$ do not decrease with the addition of a second radio, indicating no gain for those mesh nodes three hops away from a wired gateway. As mesh density increases, nodes require shorter communication paths, therefore reducing opportunities for spatial reuse. And without spatial reuse, the demand on the wired gateway node is greater, which in turn affords a greater opportunity for improvement by allowing access tier traffic to transmit concurrent with backhaul traffic. Doubling the number of radios results in less than double the fair mesh capacity, indicating that splitting access and backhaul tiers onto separate radios is not the best way to utilize a second radio. We argue that the approach of using the second radio to setup a second, independent mesh network, allows each network to service half the number of users as in a single-radio configuration.

VI. RELATED WORK

Placement protocols and algorithms have been proposed [7], [8] to arrive at an optimal set of access point locations using detailed pathloss measurement data. The drawbacks of this approach are the expense of the measurement study and the lack of resulting insight into general relationships between topology and performance. Currently, city-wide mesh networks are being designed using similar measurement-based approaches [3].

Mesh equipment manufacturers publish guidelines for the necessary mesh density based on average pathloss environments. For example, Tropos claims to provide ubiquitous coverage at a density of 4 to 8 per km² [5]. Other mesh networks [15] have used similarly rough guidelines or not published information about how their density was chosen.

Also closely related to this work is research on the effectiveness of random mesh deployments. An experimental

study [4] found that organic (randomly deployed) networks can provide broadband equivalent rates. The Roofnet study, though, focuses on single-tier mesh networks with only one active user. Our work additionally considers the access tier and multiple active flows. Further, we contrast the performance of a random deployment with a regular grid-based deployment.

Chandra et al [9] formulate the mesh planning problem in terms of placing wired gateway nodes given specific traffic demands. In contrast, our work considers the general problem of wired gateway placement in conjunction with other deployment factors.

Per-user capacity studies for wireless networks [16], [17] have derived scaling behaviors in general ad hoc networks. Li et al [18] study ad hoc network capacity with extensive simulations. These capacity studies consider random topologies with full traffic matrices and derive scaling factors, as well as the impact of the number of wireless interfaces and available channels [19]. We consider only traffic to and from wired gateways, as well as imposing a per-user fairness requirement and study performance scaling in conjunction with deployment factors.

Another closely related work [11] performs a measurement study of a mesh network deployment and highlights the importance of measurements in accurately planning and provisioning mesh networks. A prior measurement study [20] of outdoor mesh networks focuses on single-link behaviors. Our study builds on the experimental data and realistic physical layer parameterization to examine network-wide issues.

VII. CONCLUSIONS

In this work, we explore topology and deployment factors in mesh network design and study their impact on the performance of a general wireless mesh network. For this, we propose three metrics to capture three different components of a mesh network: access-tier coverage area, backhaul tier connectivity, and fair mesh capacity. For coverage area, we find that triangular and square grid topologies result in significantly better coverage than hexagonal and random topologies. Also, we find that moderate grid perturbations have only minimal impact on coverage area. Backhaul tier connectivity depends strongly on mesh node density and less on specific topology configuration; we show that random topologies only have marginally lower connectivity at high densities. Finally, the fair mesh capacity is dependent on many topology factors, including wired gateway density and grid configuration. We find that a hexagonal grid has the lowest capacity and showed that wired gateway density has an order of magnitude greater impact than mesh node density. For both coverage area and fair mesh capacity, we find randomly deployed networks to perform significantly worse than regular grid topologies.

We also consider architectural features in mesh deployments. Backhaul tier connectivity is greatly enhanced by additional backhaul radios, though the gain of systems with more than two radios is marginal. Further, we find that the popular method of using a second radio to carry only access tier traffic is an inefficient use of the additional resource.

ACKNOWLEDGMENTS

This research was supported by NSF grants ANI-0331620 and ANI-0325971, Hewlett Packard Labs, and the CISCO ARTI Program.

REFERENCES

- [1] R. Karrer, A. Sabharwal, and E. Knightly, "Enabling large-scale wireless broadband: the case for TAPs," in *Proceedings of HotNets-II*, Cambridge, MA, Nov. 2003.
- [2] S. Cherry, "Wi-fi nodes to talk amongst themselves," *IEEE Spectrum*, July 2006.
- [3] D. Neff, "Wireless Philadelphia business plan: Wireless broadband as the foundation for a digital city," 9 February 2005.
- [4] J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *Proceedings of ACM MobiCom*, Cologne, Germany, August 2005.
- [5] "Five keys to successful metro-scale wi-fi deployment," December 2004, tropos Networks White Paper.
- [6] G. E. Box, W. G. Hunter, and J. S. Hunter, *Statistics for Experimenters*. Wiley-Interscience, 1978.
- [7] M. Kamenetsky and M. Unbehaun, "Coverage planning for outdoor wireless LAN systems," in *International Zurich Seminar on Broadband Communications*, Zurich, Switz., February 2002.
- [8] Y. Lee, K. Kim, and Y. Choi, "Optimization of AP placement and channel assignment in wireless LANs," in *IEEE Conference on Local Computer Networks*, November 2002.
- [9] R. Chandra, L. Qiu, K. Jain, and M. Mahdian, "Optimizing the placement of integration points in multi-hop wireless networks," in *Proceedings of ICNP*, Berlin, Germany, October 2004.
- [10] T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies Series, T. Rappaport, Ed. Upper Saddle River, New Jersey: Prentice Hall, 1996.
- [11] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proceedings of ACM MobiSys*, Uppsala, Sweden, June 2006.
- [12] T. X. Brown, "Cellular performance bounds via shotgun cellular systems," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 11, Nov. 2000.
- [13] D. Kotz, C. Newport, and C. Elliott, "The mistaken axioms of wireless-network research," Dartmouth College Computer Science, Hanover, NH, Tech. Rep. TR2003-467, July 2003.
- [14] A. Raniwala and T. cker Chiueh, "Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network," in *Proceedings of IEEE INFOCOM'05*, Miami, FL, March 2005.
- [15] M. Ritter, R. Friday, M. Cunningham, and T. Logan, "The architecture of Metricom's MicroCellular Data Network (MCDN) and details of its implementation as the second and third generation Ricochet wide-area mobile data service," in *Broadband Communications for the Internet Era Symposium digest, 2001 IEEE Emerging Technologies Symposium on*, Richardson, TX, USA, September 2001.
- [16] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, Mar. 2000.
- [17] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *Proceedings of IEEE INFOCOM'03*, San Fransisco, CA, April 2003.
- [18] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proceedings of ACM MobiCom*, Rome, Italy, July 2001.
- [19] P. Kyasanur and N. Vaidya, "Capacity of multi-channel wireless networks: Impact of number of channels and interfaces," in *Proceedings of ACM MobiCom*, Cologne, Germany, August 2005.
- [20] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *Proceedings of ACM SIGCOMM*, Portland, OR, 2004.