

Downlink Capacity of Hybrid Cellular Ad Hoc Networks

Lap Kong Law, Konstantinos Pelechrinis, *Student Member, IEEE*, Srikanth V. Krishnamurthy, *Senior Member, IEEE*, and Michalis Faloutsos, *Member, IEEE*

Abstract—Augmenting cellular networks with shorter multihop wireless links that carry traffic to/from a base station can be expected to facilitate higher rates and improved spatial reuse, therefore potentially yielding increased wireless capacity. The resulting network is referred to as a *hybrid network*. However, while this approach can result in shorter range higher rate links and improved spatial reuse, which together favor a capacity increase, it relies on multihop forwarding, which is detrimental to the overall capacity. In this paper, our objective is to evaluate the impact of these conflicting factors on the overall capacity of the hybrid network. We formally define the capacity of the network as the maximum possible downlink throughput under the constraint of *max-min fairness*. We analytically compute the capacity of both one- and two-dimensional hybrid networks with regular placement of base stations and users. While almost no capacity benefits are possible with linear networks due to poor spatial reuse, significant capacity improvements with two-dimensional networks are possible in certain parametric regimes. Our simulations also demonstrate that in both cases, if the users are placed randomly, the behavioral results are similar to those with regular placement of users.

Index Terms—Capacity, hybrid networks.

I. INTRODUCTION

ONE can increase the spatial reuse in a cellular data network by reducing the range of direct cellular coverage and creating a larger number of cells, each of which houses an expensive base station (BS) that is directly connected to the Internet. In comparison with this approach, augmenting cellular networks with *ad hoc* wireless connectivity to improve spatial reuse is attractive in terms of the incurred cost [1], [2]. We call this latter type of network a hybrid cellular ad hoc network, or simply hybrid network. A natural question that arises is: do such hybrid networks indeed offer a higher capacity than the original cellular network?

While there has been a surge in interest in modelling various kinds of hybrid networks (to be elaborated in Section VI) [1]–[12], the above question on whether using multihop *ad hoc* wireless relays is beneficial (or even detrimental) in terms of the *capacity* of the original cellular network has not been answered to date. While the use of shorter range and, hence, higher rate wireless *ad hoc* links may improve spatial reuse (more simultaneous transmissions can occur) [13], the use of multihop relaying increases the number of wireless hops traversed, and

Manuscript received June 10, 2008; revised February 03, 2009; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor B. Levine.

The authors are with the Department of Computer Science and Engineering, University of California, Riverside, CA 92502 USA (e-mail: lklaw@cs.ucr.edu; kpele@cs.ucr.edu; krish@cs.ucr.edu; michalis@cs.ucr.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2009.2023651




Fig. 1. A pictorial representation of a hybrid network.

this drives down the achievable throughput. Given the two conflicting factors, it is unclear whether or not the capacity of the network will in fact increase relative to the original cellular network. In this paper, we determine under *what conditions* and by *how much* the downlink capacity of a hybrid cellular ad hoc network is higher/lower than that of the original pure cellular network. Both linear and two-dimensional hexagonal networks are considered.

A. The Hybrid Network in a Nutshell

To describe the hybrid network in brief, consider a two-dimensional hexagonal hybrid network as shown in Fig. 1. In this network, only users that are within the *reduced cellular coverage*, depicted with a dark hexagon, receive downlink traffic directly from the BS; this direct link between the BS and such users is called the *infrastructure component* of the hybrid network. The reduced cellular coverage enhances the transmission efficiency of the BS since the directly connected users, being close to the BS, usually have good channels to the BS; thus, the BS can now support higher rates to these users. On the other hand, users that are outside the reduced cellular coverage require that the *directly connected* users act as *proxies* and forward traffic from the BS. In fact, only a subset of these outside users may directly receive traffic from the proxies. These users will then have to act as *relays* and forward traffic to other users that are further away from the BS. The part of the network that delivers the data from the proxies to the outside users is called the *ad hoc component* of the hybrid network. Note that in the *ad hoc* component, transmissions are carried via multiple hops to the end destination; we use the term “*ad hoc*” to reflect this.

B. The Capacity of the Hybrid Network

We define the *capacity* of the hybrid network to be the maximum possible network throughput under the conditions

of *max-min* fairness.¹[14] We construct analytical models to compute the capacity of the one- and two-dimensional hybrid network under the aforementioned conditions; the capacity varies with system parameters, as we articulate later. Note that for a defined set of parameter values, the computed capacity may or may not be achievable in practice; however, the achievable throughput will always be less than or equal to the capacity. The computed expressions can also be trivially used to compute the capacity of the original pure cellular network.

The *distinguishing aspects of our work* can be summarized as follows.

- In computing the downlink capacity, *we consider interference from surrounding cells*. We account for such interference when modelling both direct BS-to-user communications and *ad hoc* user-to-user communications.
- Unlike in previous efforts (such as [8] and [12]), our model reflects the impact of the *distance spanned by the link on the achieved rate*. In other words, we examine the impact of rate versus range tradeoffs on the capacity of the hybrid network.
- Unlike most previous efforts, our model accounts for the fact that *the BS is always at one end of every communication*. Given that cellular users typically download content from the Internet [15] (as opposed to uploading), our work focuses on downlink communications.² Note that the considered traffic pattern has a big impact on the computed capacity since it affects the extent of possible spatial reuse in the *ad hoc* part of the network. With the considered traffic pattern, in the hybrid network, nodes that are closer to the BS carry more traffic and, thus, are likely to be active more often. We take this into account in our analysis.
- *We consider a fixed spectral band that is shared between the direct cellular communications and the ad hoc communications*. Most previous efforts assume the use of an additional spectral band for supporting multihop relay communications [2], [4]. Our approach facilitates a fair comparison between the capacities of the hybrid network and its original pure cellular counterpart.
- In addition to validating our analysis, our simulations show that *the capacity achieved by placing users randomly within the network is similar to that with the regular placements of users* (the analysis assumes regular placements).

C. Our Work in Perspective

Our analysis computes the downlink capacity in one- and two-dimensional hybrid networks. We compare this capacity with that of the original cellular network, which is a special case of the hybrid network when no *ad hoc* component exists. Hence, the capacity of the pure cellular network is trivially computed from the same analysis and with the same set of assumptions, making the comparison meaningful. Our analytical assessments suggest that due to poor spatial reuse in the *linear network*, almost no capacity benefits are seen. Capacity improve-

¹We seek to avoid the starvation of users with poor channel conditions, i.e., consider a scenario in which a fair share of the network capacity is provided to each user.

²The traffic patterns for the uplink case significantly differ from those of the downlink. All uplink traffic is destined from the users to the center of the cell, which makes the interference patterns and, as a consequence, the analysis of the two cases significantly different.

ments are possible in the case of the *hexagonal hybrid network* when three key parameters discussed below are appropriately tuned. With appropriately selected parameters, one can achieve the right tradeoffs, and the gains with the hybrid network can be as high as 70% more than those of the corresponding pure cellular network. In addition, we validate our analytical results with extensive simulations.

D. Fundamental Tradeoffs

We find that the capacity of the hybrid network depends on the following parameters: 1) the size of the *BS footprint*, 2) the *spectrum allocation* between the BS-to-user links and the user-to-user links, and 3) the *transmission range* of the multihop *ad hoc* wireless links. The first parameter b , the size of the BS footprint, determines the region of direct cellular coverage (as described earlier). The smaller the footprint is, the higher the rate at which the BS can transmit to its directly connected users will be; however, this will lead to a higher multihop forwarding overhead. Secondly, the fixed spectral bandwidth has to be appropriately apportioned between the infrastructure and the *ad hoc* components in order to ensure that one component does not become a bottleneck. The capacity depends on what fraction of this spectral bandwidth is assigned to the two components. Lastly, we can vary the transmission range r of the *ad hoc* links. Choosing a longer range would result in fewer hops; however, the rate on each hop will be low. One could instead choose to use a large number of high-rate short-range links at the cost of traversing more hops.

The rest of this paper is organized as follows. In the next section, we formulate the problem. In Section III, we compute the downlink capacity of the linear hybrid network. We extend our analysis to a hexagonal hybrid network in Section IV. A description of our simulation experiments and comparisons of analytical and simulation results form Section V. Related work is discussed in Section VI. We conclude in Section VII.

II. PROBLEM FORMULATION

In this section, we describe the network being considered and the models that we use; we list and justify the assumptions that we make. We repostulate some of the definitions from Section I for completeness.

A. The Hybrid Network

The hybrid network consists of two components: 1) the *infrastructure component* and 2) the *ad hoc component*. The infrastructure component refers to the part of network within which a user can communicate with its serving BS directly. The *ad hoc* component refers to the part of the network from which a user cannot communicate (effectively) with its serving BS directly; it requires its neighbors to relay traffic from the BS (across multiple hops). For clarity and the purposes of analysis, we use the concept of a *base-station footprint*. We formally define it to be the maximum distance from the BS within which a given user can communicate directly with the BS given a data-rate requirement.

In a linear hybrid network [see Fig. 2(a)], BSs are evenly placed on a straight line. A hexagonal hybrid network [see Fig. 3(a)] contains hexagonal cells, and each cell contains a single BS located at the center of the cell. The distance between any two BSs is the same and equal to a value $2D$, as alluded to

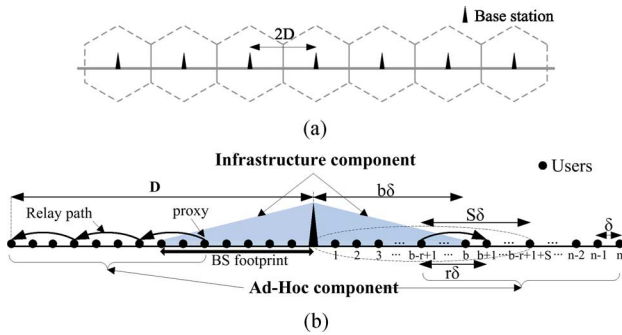


Fig. 2. Linear hybrid network. (a) Cell arrangement in the linear hybrid network. (b) Users arrangement within a cell.

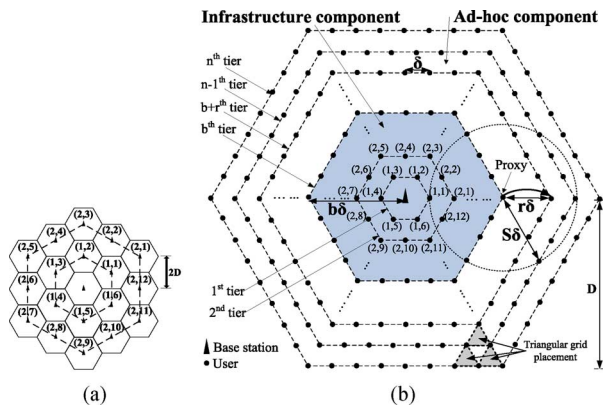


Fig. 3. Hexagonal hybrid network. (a) Cellular arrangement (cells are arranged in levels). (b) Users' arrangement within a cell (users are arranged in tiers).

by Figs. 2(a) and 3(a). In a hexagonal hybrid network, for each BS, the surrounding BSs are arranged in *levels* and are indexed as shown in Fig. 3(a) (dotted concentric hexagons). Note that even though we depict a finite number of BSs in Figs. 2(a) and 3(a), we consider hybrid networks with an infinite number of BSs in our analysis. Thus, any receiving user is subject to interference from an infinite number of simultaneous transmitters from surrounding cells. For the purposes of this paper, we limit ourselves to downlink transmissions, i.e., transmissions from the BS to the users. We expect that this is likely to represent the biggest portion of traffic flow given that Internet traffic largely consists of Web downloads [15].

In order to ensure the tractability of our analysis, we assume that users are *uniformly* and *regularly* placed in the network. We relax this assumption in our simulations and consider the more complex case where the users are randomly placed in the network. In the linear hybrid network, users are evenly spaced along a straight line on both sides of the serving BS [see Fig. 2(b)]. In the two-dimensional case, we consider the triangular grid placement of users [see Fig. 3(b)] since with this, users can be placed within a hexagonal cell perfectly with the highest density, for any given separation between adjacent users. In our formulation, the users are arranged and indexed in *tiers*; the tiers are represented by the dotted concentric hexagons in Fig. 3(b).

B. Configuration and Spectral Allocation

We assume that all users are equipped with two communication interfaces: 1) the *cellular interface* and 2) the *ad hoc inter-*

face. A user within the BS's footprint uses its cellular interface to communicate with the BS; it uses its other interface to communicate with the users outside the BS's footprint. Note here that the two interfaces operate on nonoverlapping channels. The channels are half-duplex. A *fixed spectral bandwidth* is shared between the two components. Let us denote by B , the total available bandwidth (in hertz); let γB (where $0 < \gamma \leq 1$) be the fraction of the total bandwidth that is allocated for infrastructure communications. Then, the bandwidth allocated for multihop *ad hoc* communications would be $(1-\gamma)B$.

C. Channel Properties and Models

We consider a simple path-loss channel propagation model. In particular, $P_{RX} = \frac{P_{TX}}{d^\alpha}$ where P_{TX} is the average transmission power, P_{RX} is the average received signal power, d is the distance between the transmitter and the receiver, and α is the path loss exponent, usually $\alpha \geq 2$. With a goal of capturing long-term channel effects, multipath fading is not considered. We also do not consider shadow fading.

To compute the capacity of a wireless channel, we model the cochannel interference as Gaussian,³ as in [13] and [16]; then the *channel capacity* is given by *Shannon's capacity formula* [17]

$$\text{Capacity} = W \log_2(1 + \text{SINR}) \quad (1)$$

where W is the bandwidth (in hertz) and SINR is the signal-to-(interference plus noise) ratio. Note that the *channel capacity* is the maximum achievable throughput on a link and will not be achieved in practice due to protocol semantics and overheads. The actual achievable link capacity depends on the link quality but also on protocol specifics and the interference. It also depends on the type of fading (frequency-selective versus nonselective) and the coherence time of the channel [18]. These are extremely difficult to analytically model. Thus we use a popularly considered upper bound given by the Shannon capacity.

D. Transferring Traffic From a BS to the Users Within its Footprint

As mentioned earlier, a BS transmits directly only to the users within its footprint. BSs perform transmissions in time-division multiplexed mode, i.e., at any given point in time, a BS transmits to only one user, and at full power P_{BS} . The transmission power P_{BS} is the same for all BSs in the network. We consider an aggressive approach in which all the BSs are transmitting all the time. Therefore, a receiving user experiences interference from all other BSs.⁴ This model has been proven to maximize the overall throughput of the system [19] and has been employed as the transmission strategy for 1xEV-DO [20].

An infrastructure link can be formed between a user and its closest BS. We define the *infrastructure link capacity* (ILC) as the maximum transmission rate that can be achieved if the user is the *only* user that is served by the BS. With our model, the ILC of a user is dependent on its location only. By computing the SINR

³In the presence of a large number of independent interferers, the central limit theorem supports this assumption.

⁴Such an aggressive approach would typically increase the collisions at users at the boundary of the cell in a pure cellular network. However, such users will not typically communicate with the BS directly in the hybrid network due to the smaller BS footprint. These users instead use multihop *ad hoc* links for communications, and thus, the above effects are less pronounced.

of a given node, we can use (1) to compute its ILC. However, the actual *servicing rate* of a user depends on the number of users, their ILCs, and the transmission scheduling strategy of the BS. To determine the transmission strategy of BSs, cellular network providers usually consider both throughput and fairness issues [21]. Here, we consider the transmission strategy that achieves maximum fairness.

E. Traffic Flow to the Ad Hoc Component via the Proxies and the Relays

Users that lie outside the BS footprint will have to rely on their neighbors to act as *relays* to deliver traffic from the BS. Note that these relays may in turn depend on other neighbor relays. Ultimately, the relays that are within the BS footprint relay traffic to those users that are outside the footprint (using the *ad hoc* interface). Thus, these users act as *proxies* to the BS.

As in typical ad hoc networks, carrier sensing is employed [16] within the *ad hoc* component. All users use the same transmission power P_{user} . A user can transmit (using its *ad hoc* interface) only if there is no other active transmitter that is located within its sensing range [13]. Note that the sensing range is only a function of the transmission power and the sensitivity of receivers [22].⁵

The capacity of an *ad hoc* link depends on the distance between the transmitting and receiving user. In order to sustain a desired transmission rate on an *ad hoc* link, the maximum distance that a transmission can span must be restricted. This distance is the *transmission range* (or simply range) of the link. Choosing a shorter range increases the achievable rate on each link; however, it also increases the number of hops needed to reach a destination user. We examine the tradeoffs between the two effects.⁶

We assume that a relay path from a proxy to a destination user is the minimum-hop path, given the transmission range. This assumption is commonly made in previous capacity analyses [23], [24] since it minimizes the relaying overhead incurred in the network.

F. Impact of Density

Note here that the imposed span of the links in the network, and therefore the achievable rate, is dictated by the density of users in the network. If the network is sparse, one would be forced to establish longer links, and therefore the achievable rates on the links will be limited. If, on the other hand, there is a dense population of *ad hoc* users, one could potentially form shorter and therefore higher rate links. We study the impact of user density in our performance evaluations. Note, however, that we do not consider extremely sparse cases where *ad hoc* links can only support extremely low data rates since these cases are likely to be of no interest.

G. Metrics of Interest

We wish to compute the capacity when the network achieves maximum fairness, i.e., all users receive approximately the same throughput from the BS. This avoids starvation for users with

⁵In IEEE802.11a, a typical transmission power is 15 dBm (32 mW) and the receiver sensitivity is around -90 dBm ($1e^{-9}$ mW). With this, the carrier sensing range $S \approx 421$ m.

⁶This distinguishes our work from other similar studies wherein the variability in the achievable rate with distance is not considered.




Fig. 4. An abstract representation of a hybrid network.

poor channel conditions. This capacity is given by the maximum *uniform throughput*, which is defined as the minimum of the achieved throughputs, considering all users in the network. We first determine the maximum uniform throughput of a) the infrastructure component and b) the *ad hoc* component, respectively. Depending upon the size of the BS footprint, the bandwidth allocation between the infrastructure and the *ad hoc* components, and the transmission range of an *ad hoc* link, the capacity bottleneck could appear either in the infrastructure part or in the *ad hoc* part. Therefore, the maximum uniform throughput of the hybrid network is the minimum of the maximum uniform throughputs of the two components.

Fig. 4 depicts an abstract representation of a hybrid network. The capacity between the BS and the proxies (as shown) is the uniform throughput of the infrastructure component, which we denote by λ_I . The proxies deliver information to the rest of the nodes in an *ad hoc* component at a rate we call the uniform throughput of the *ad hoc* component and denote by λ_A . Note that the interactions between the data transmissions from the proxies and the relay operations within the *ad hoc* component affect the maximum achievable λ_A (details to be discussed). We denote by λ_H the uniform throughput of the hybrid network. Then

$$\lambda_H = \min(\lambda_I, \lambda_A). \quad (2)$$

III. THE LINEAR HYBRID NETWORK

To begin with, we focus on a linear hybrid network [see Fig. 2(b)]. One may consider this topology to be representative of a vehicular network; a row of cars on the highway or a row of homes on a street, communicating via their wireless interfaces. We consider this topology for two reasons: a) its simplicity and b) it provides a means of capturing the impact of spatial reuse on capacity (as we discuss later when we compare the linear network with a two-dimensional hexagonal network). Let n denote the number of users on each side of the BS within a cell and let δ denote the distance between adjacent users. Then, the radius of each linear cell is $D = n\delta$. Let r denote the transmission range of an *ad hoc* link (for a given transmission rate requirement, as discussed in Section II), S the sensing range of each user, and b the size of the BS footprint. These distances are expressed in terms of units, each of which is of length δ . Note that $S \cdot \delta \geq r \cdot \delta \Rightarrow S \geq r$ and $n \cdot \delta \geq b \cdot \delta \geq r \cdot \delta \Rightarrow n \geq b \geq r$. We only consider the case where $b \geq r$ since the direct link from the more powerful BS is likely to be at least as long as an *ad hoc* link. For ease of discussion, users are indexed from one to n , as shown in the figure, and USER_i represents the user at a distance of i units from the BS. Note that the two users that are equidistant from the BS (on either side) have the same index. Given our minimum-hop routing strategy, each user, when establishing an *ad hoc* link, selects the neighbor that is *furthest* from itself from among its *reachable* neighbors.

With this, USER_i will always attempt to form an *ad hoc* link with USER_{i-r} and USER_{i+r} , if these users exist. Furthermore, USER_{b-r+1} to USER_b are the potential proxies since they can form links with both the BS and the users that are outside the BS footprint.

A. Uniform Throughput of the Infrastructure Component

As discussed, in order to obtain the maximum uniform throughput λ_I , the BS must ensure that all users in the infrastructure component receive the same throughput. Furthermore, the BS must select the best set of proxies and forward to these proxies the appropriate amounts of traffic destined to the users in the *ad hoc* component. We have to find the optimal BS transmission schedule that achieves this maximum fairness. We denote the ILC of USER_i by c_i and the fraction of time that the BS transmits to USER_i by τ_i . Given this, the above problem can be formulated as follows:

$$\begin{aligned} & \text{maximize} && \lambda_I \\ & \text{subject to} && \tau_i c_i = \lambda_I \quad \forall i = 1, \dots, (b-r) \\ & && \tau_j c_j = \left(1 + \left\lfloor \frac{n-j}{r} \right\rfloor\right) \lambda_I \\ & && \quad \forall j = (b-r+1), \dots, b \\ & && 2 \times \sum_{k=1}^b \tau_k = 1. \end{aligned} \quad (3)$$

The first constraint mandates that the throughput of the non-proxy users in the infrastructure component is the same and is equal to λ_I . The second constraint requires that each proxy receives a throughput of λ_I for itself and obtains an additional throughput of $\left\lfloor \frac{n-j}{r} \right\rfloor \lambda_I$ for the users in the *ad hoc* component, for which it acts as a relay. Note that $\left\lfloor \frac{n-j}{r} \right\rfloor$ is the number of users in the *ad hoc* component that are served by proxy USER_j . The third constraint ensures that the sum of the fractional shares assigned to the proxies does not exceed unity (the factor two accounts for the two sides of the BS).

To solve the maximization problem (3), we need to determine the infrastructure link capacity $c_i, \forall i = 1, \dots, b$. Given the regular placement of the BSs and the users, the SINR as perceived at any user can be determined. The distance between USER_i and the serving BS is $i\delta$, and the distance between USER_i and the interferers (the BSs from the surrounding cells) is $2D-i\delta, 4D-i\delta, 6D-i\delta, \dots$ and so on, and $2D+i\delta, 4D+i\delta, 6D+i\delta, \dots$. Therefore, the SINR at USER_i is

$$\begin{aligned} \text{SINR}_{\text{USER}_i} &= \frac{\frac{P_{\text{BS}}}{(i\delta)^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{P_{\text{BS}}}{(2kD+i\delta)^\alpha} + \frac{P_{\text{BS}}}{(2kD-i\delta)^\alpha} \right]} \\ &= \frac{\frac{1}{(i\delta)^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{1}{(2kD+i\delta)^\alpha} + \frac{1}{(2kD-i\delta)^\alpha} \right]} \\ &\leq \frac{\frac{1}{(i\delta)^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{2}{(2kD)^\alpha} \right]} = \frac{2^{\alpha-1} n^\alpha}{i^\alpha \zeta(\alpha)} \end{aligned} \quad (4)$$

where $n = \frac{D}{\delta}$ and $\zeta(\alpha)$ is the *Riemann zeta function* representing $\sum_{k=1}^{\infty} \frac{1}{k^\alpha}$. In the above derivation, we ignore the effect of noise since we consider the case of an interference-dominated

environment. The final inequality follows from the convexity of the function $\frac{1}{(2kD+i\delta)^\alpha} + \frac{1}{(2kD-i\delta)^\alpha}$ (Jensen's inequality).

From (1), the ILC is then computed to be

$$c_i \leq \gamma B \times \log_2 \left(1 + \frac{2^{\alpha-1} n^\alpha}{i^\alpha \zeta(\alpha)} \right). \quad (5)$$

Using the equality in the above expression, λ_I is found by solving the maximization problem defined in (3); standard linear programming approaches are applicable.

B. Uniform Throughput of the Ad Hoc Component

Next, we compute the maximum possible value for λ_A . Given that we only consider downlink traffic, the traffic from the BS must flow through the proxies to enter the *ad hoc* component. Therefore, the uniform throughput in the *ad hoc* component depends on how well the proxies transport traffic. To this end, we define and compute the following two quantities: a) **NORMALIZED RELAY LOAD**: The total relay load (due to proxy and relay transmissions) incurred for each unit of traffic delivered per user in the *ad hoc* component, and b) **ONE HOP THROUGHPUT**: The maximum total one-hop throughput (due to proxy and relay transmissions) in the *ad hoc* component. Since the uniform throughput of the *ad hoc* component is λ_A , the total traffic load that must be generated in the *ad hoc* component is at least **NORMALIZED RELAY LOAD** $\times \lambda_A$. This value cannot exceed what can be carried by the *ad hoc* component, i.e.,

$$\text{NORMALIZED RELAY LOAD} \times \lambda_A \leq \text{ONE HOP THROUGHPUT}. \quad (6)$$


Given the above inequality, if we are able to compute the **ONE HOP THROUGHPUT** and **NORMALIZED RELAY LOAD**, we can compute an upper bound on λ_A .

We first calculate **NORMALIZED RELAY LOAD**. As the distance between the destination user and the BS increases or the transmission range of an *ad hoc* link decreases, the number of relay operations (proxy and relay transmissions) needed and, thus, the relay load generated increases. As per the minimum-hop routing policy, the relay load that is incurred by relaying one unit of traffic to a user, say, USER_k , in the *ad hoc* component is $\lceil (k-b)/r \rceil$. Given this, the total relay load **NORMALIZED RELAY LOAD** (by considering users on both sides of the BS) is given by

$$\text{NORMALIZED RELAY LOAD} \geq 2 \times \sum_{k=b+1}^n \left\lceil \frac{k-b}{r} \right\rceil. \quad (7)$$

The use of the \geq sign is to account for imperfections in routing and BS scheduling. We take the minimum value here in order to obtain an upper bounded λ_A . Therefore, the total traffic load of the *ad hoc* component with a uniform throughput of λ_A is at least $2 \times \sum_{k=b+1}^n \lceil \frac{k-b}{r} \rceil \times \lambda_A$.

To determine the **ONE HOP THROUGHPUT**, we first calculate the transmission rate that can be sustained on an *ad hoc* link given its transmission range. Then, we determine an upper bound on the number of *ad hoc* transmissions that can occur simultaneously. The product of the two gives the maximum achievable one-hop throughput in the *ad hoc* component. Note that **ONE HOP THROUGHPUT** is different from the aggregate one-hop throughput that is defined in [23] and [24]. The definition used in [23] and [24] assumes that a node always has

Fig. 5. The infinite linear *ad hoc*; BSs are not shown.

data to transmit when given the opportunity to transmit and, thus, the maximum spatial reuse is exploited. However, due to the direction of traffic flow in the hybrid network, users that are in vicinity of the BS have to transmit more often than the others that are not directly connected to the BS. A user may not have data to transmit when given the opportunity to do so. We will take this heterogeneous load distribution into account when computing ONE HOP THROUGHPUT.

We consider an infinite linear *ad hoc* network⁷ and the case where the number of simultaneous transmissions in the *ad hoc* component is maximized. Therefore, the distance between adjacent *transmitting* users is to be minimized. Given the infinite *ad hoc* network, the average interference power and, hence, the SINR at any receiving user is equal. Without loss of generality, we focus on transmitter–receiver pair A, B, shown at the center of Fig. 5. Then the distances between the receiver B and the interferers from the left are $(S+r)\delta$, $(2S+r)\delta$, $(3S+r)\delta$, ... and from the right are $(S-r)\delta$, $(2S-r)\delta$, $(3S-r)\delta$, ... The SINR is then given by

$$\begin{aligned} \text{SINR} &= \frac{\frac{P_{\text{USER}}}{(r\delta)^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{P_{\text{USER}}}{((kS+r)\delta)^\alpha} + \frac{P_{\text{USER}}}{((kS-r)\delta)^\alpha} \right]} \\ &= \frac{\frac{1}{r^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{1}{(kS+r)^\alpha} + \frac{1}{(kS-r)^\alpha} \right]} \leq \frac{\frac{1}{r^\alpha}}{\sum_{k=1}^{\infty} \left[\frac{2}{(kS)^\alpha} \right]} = \frac{S^\alpha}{2r^\alpha \zeta(\alpha)}. \end{aligned} \quad (8)$$

The inequality in the last step follows from the convexity of the function $\frac{1}{(kS+r)^\alpha} + \frac{1}{(kS-r)^\alpha}$ (Jensen's inequality). We ignore the effects of noise since we consider an interference-dominated environment. We assume that each USER enjoys the highest achievable SINR (i.e., the above bound). Thus, the highest transmission rate that can be reliably sustained on an *ad hoc* link is given by

$$\text{Rate}_{\text{ad hoc}} \leq (1-\gamma)B \times \log_2 \left(1 + \frac{S^\alpha}{2r^\alpha \zeta(\alpha)} \right). \quad (9)$$

C. Upper Bound on the Number of Simultaneous *Ad Hoc* Transmissions

We denote by PROXY TX and RELAY TX the upper bounds on the number of simultaneous proxy and relay transmissions in the network, respectively (assuming maximum fairness). Then, the upper bound on the number of simultaneous *ad hoc* transmissions in the network will be (PROXY TX + RELAY TX). In fact, there is a correlation between PROXY TX and RELAY TX. One may envision that when a proxy transmits a packet, the packet is relayed a number of times depending on the distance

⁷Any large network with continued coverage can be thus approximated. The approximation helps us in eliminating boundary effects.

between the proxy and the destination until it reaches the destination. Due to the deterministic placement of users in our model, we can find an upper bound on RELAY TX in terms of PROXY TX and a multiplicative factor. This multiplicative factor is denoted by RP RATIO. In other words, RP RATIO represents the maximum possible ratio of the number of relay transmissions to the number of proxy transmissions in order for every user in the *ad hoc* component to receive a unit of data traffic. RP RATIO is given by

$$\text{RP RATIO} = \frac{2 \times \sum_{k=b+r+1}^n \left\lceil \frac{k-(b+r)}{r} \right\rceil}{2 \times (n-b)} \quad (10)$$

where the numerator reflects the number of relay transmissions in the *ad hoc* component (each USER_k, where $k = b+r+1, \dots, n$, incurs $\lceil \frac{k-(b+r)}{r} \rceil$ number of relay transmissions) and the denominator reflects the total number of proxy transmissions (which equals to the number of users in the *ad hoc* component). Then the maximum number of actual simultaneous *ad hoc* transmissions is upper bounded by PROXY TX \times (1 + RP RATIO). Our target is then to compute PROXY TX.

Due to carrier sensing, the fraction of time that a proxy can transmit is affected by the number of proxies and relays that are located within its sensing range and their *carried loads* (the carried load of a user is the amount of traffic that is to be transmitted by the user). We define the *normalized* carried load of USER_k, l_k , to be the carried load of USER_k in order for each user that relies on USER_k as a relay in the *ad hoc* component to receive a unit of data traffic. Then, l_k is given by

$$l_k = \begin{cases} \left\lfloor \frac{n-k}{r} \right\rfloor, & \text{if } b-r+1 \leq k \leq n-r \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

Essentially, l_k is equal to the number of users to which USER_k will relay traffic. For any given silent proxy (say, Proxy A), there could be at most *two* other active transmitters in its sensing range. In the best case, the two active transmitters simultaneously transmit, thereby restricting the transmission opportunity of Proxy A. This effect causes the fraction of time for which Proxy A can transmit to be higher than if the two active transmitters were to transmit sequentially. Thus, an upper bound on the fraction of transmission time of an arbitrary proxy is the ratio of the proxy's normalized carried load to the sum of its normalized carried load and *half* of the normalized carried loads of all the other users within its sensing range. Therefore, PROXY TX, the sum of the fraction of transmission times of all proxies, is

$$\text{PROXY TX} \leq 2 \times \sum_{k=b-r+1}^{\min(b,n-r)} \frac{l_k}{l_k + \frac{1}{2} \sum_{m=-S, m \neq 0}^S l_{|k+m|}} \quad (12)$$

where $|\cdot|$ is the absolute function; $\min(b, n-r)$ is the furthest that a proxy can be from the BS. Note that in the above formulation, we neglect the effects of *ad hoc* transmissions that occur in the neighboring cells. One might expect that proxy transmissions are not affected by adjacent cells since it is likely that $n \gg S$ and $n > b$ in typical scenarios. Furthermore, the effects of users transmissions at the boundary of an adjacent cell on the transmission possibilities of a proxy are likely to be insignificant, given that the carried load of these users is small.

From this, ONE HOP THROUGHPUT is computed to be

ONE HOP THROUGHPUT

$$\leq \text{PROXY TX} \times (1 + \text{RP RATIO}) \times \text{Rate}_{\text{ad hoc}}. \quad (13)$$

By using (6), (7), and (13), the uniform throughput of the *ad hoc* component λ_A is thus

$$\lambda_A \leq \frac{\text{PROXY TX} \times (1 + \text{RP RATIO}) \times \text{Rate}_{\text{ad hoc}}}{\text{NORMALIZED RELAY LOAD}}. \quad (14)$$

We defer a discussion on actual numerical results and the interpretation of these results to Section V.

IV. THE HEXAGONAL HYBRID NETWORK

In this section, we compute the capacity of a two-dimensional hybrid hexagonal network (Fig. 3). Most of the notations used and the assumptions made are similar to those with the linear hybrid network. As discussed in Section II, we arrange users in *tiers* (n in number) within a cell. A tier essentially consists of the users that lie on a hexagon of a certain major radius. A user in a cell is indexed as $\text{USER}_{i,j}$, where i ($i \geq 1$) represents the tier to which the user belongs and j ($1 \leq j \leq 6i$) identifies the user within that tier [see Fig. 3(b)]. The size of the BS footprint represents the number of directly reachable tiers from the BS. Thus, the region enclosed by the BS footprint has a hexagonal shape with the corresponding major radius. All users that are located between the $(b-r+1)$ th and the b th tiers are potential proxies. The hexagonal hybrid network is depicted in Fig. 3(b).

A. Uniform Throughput of the Infrastructure Component

We wish to maximize λ_I such that all users receive the same throughput (max-min fairness). We assume that the ILCs of users that are located within the same tier are the same and equal to the best ILC possible for any user within the tier⁸. Let c_i denote the ILC of a user in the i th tier and τ_i the fraction of time that the BS transmits to a user in the i th tier. The maximum value of λ_I is computed by solving the following:

$$\begin{aligned} & \text{maximize} && \lambda_I \\ & \text{subject to} && \tau_i c_i = \lambda_I \quad \forall i = 1, \dots, (b-r) \\ & && \tau_j c_j = \left(1 + \sum_{m=1}^{\lfloor \frac{n-j}{r} \rfloor} \frac{6(j+mr)}{6j}\right) \lambda_I \\ & && \quad \forall j = (b-r+1), \dots, b \\ & && \sum_{k=1}^b 6k\tau_k = 1. \end{aligned} \quad (15)$$

The above maximization problem differs from the one with the linear hybrid network in terms of the average number of users that is supported by each proxy (in the second constraint) and the number of users in each tier (in the third constraint). When c_i s, $\forall i = 1, \dots, b$ are determined, the maximization problem (15) can be solved as a linear program.

To determine c_i s, we compute the SINR of a user in the i th tier of users. We denote by $d_{i,j}$ the distance between $\text{USER}_{i,j}$ and its serving BS. $D_{i,j}^{u,v}$ is the distance between $\text{USER}_{i,j}$ and

the surrounding base station $\text{BS}_{u,v}$.⁹ Then, the SINR of $\text{USER}_{i,j}$ (and an upper bound on the SINR of the users in the i th tier) is given by

$$\begin{aligned} \text{SINR}_{\text{USER}_{i,j}} &= \frac{P_{\text{BS}}}{(d_{i,j})^\alpha} \\ &= \frac{\infty \sum_{u=1}^{\infty} \sum_{v=1}^{6u} \left[\frac{P_{\text{BS}}}{(D_{i,j}^{u,v})^\alpha} \right]}{\left(\frac{2}{\sqrt{3}i\delta} \right)^\alpha} \left(\text{since } d_{i,j} \geq \frac{\sqrt{3}i\delta}{2} \right) \\ &\leq \frac{\infty \sum_{u=1}^{\infty} \sum_{v=1}^{6u} \left[\frac{1}{(D_{i,j}^{u,v})^\alpha} \right]}{\left(\frac{2}{\sqrt{3}i\delta} \right)^\alpha} \\ &< \frac{\infty \sum_{u=1}^{\infty} \frac{6u}{(2uD)^\alpha}}{\sqrt{3}^{(\alpha+2)} i^\alpha \zeta(\alpha-1)} = \frac{2^{2\alpha-1} n^\alpha}{\sqrt{3}^{(\alpha+2)} i^\alpha \zeta(\alpha-1)} \end{aligned} \quad (16)$$

where $n = \frac{D}{\delta}$. Again, noise is ignored since we consider an interference dominated environment.

In the last step of the computation, we take advantage of the fact that the interference experienced by $\text{USER}_{i,j}$ from any particular level of surrounding BSs must be greater than or equal to the interference experienced if the $\text{USER}_{i,j}$ was located at the center of the serving cell; we also take advantage of the fact that the distance between the center of the serving BS and any BS in the u th level of surrounding BSs is at most $2uD$. Thus, $\sum_{v=1}^{6u} \left[\frac{1}{(D_{i,j}^{u,v})^\alpha} \right] \geq \frac{6u}{(2uD)^\alpha}$.

With the above, the ILC c_i is computed using (1) to be

$$c_i \leq \gamma B \times \log_2 \left(1 + \frac{2^{2\alpha-1} n^\alpha}{\sqrt{3}^{(\alpha+2)} i^\alpha \zeta(\alpha-1)} \right). \quad (17)$$

Then, λ_I is found by solving the maximization problem (15) by using standard linear programming techniques.

B. Uniform Throughput of the Ad Hoc Component

We use an approach similar to the one in Section III to determine the uniform throughput of the *ad hoc* component λ_A . We determine NORMALIZED RELAY LOAD, ONE HOP THROUGHPUT, $\text{Rate}_{\text{ad hoc}}$, PROXY TX, and RP RATIO as before.

We first determine NORMALIZED RELAY LOAD. As before, all traffic flows follow their minimum-hop paths. Thus, the relay load incurred by relaying one unit of traffic to a user in the k th tier of users (where $k > b$) is $\lceil (k-b)/r \rceil$. Since there are $6k$ users in the k th tier of users, the total relay load NORMALIZED RELAY LOAD is

$$\text{NORMALIZED RELAY LOAD} \geq \sum_{k=b+1}^n 6k \times \left\lceil \frac{k-b}{r} \right\rceil. \quad (18)$$

Therefore, the total relay load with a uniform throughput λ_A is at least $\sum_{k=b+1}^n 6k \times \left\lceil \frac{k-b}{r} \right\rceil \times \lambda_A$.

Then we compute the components that constitute ONE HOP THROUGHPUT [see (13)].

We first compute $\text{Rate}_{\text{ad hoc}}$. If one were to ignore the BSs in the hybrid network, due to the regular placement of users, one would observe an infinite hexagonal grid of users. Let us focus

⁹ $\text{BS}_{u,v}$ is the BS that is located in the u th level of surrounding BSs and with index v in that level. Note that u and v are defined with respect to the current cell where $\text{USER}_{i,j}$ is located. Refer to Fig. 3 for the naming convention of BSs and users.

⁸This provides an upper bound on the achievable throughput.

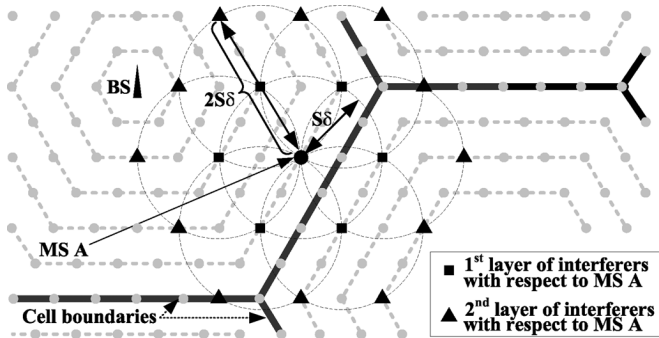


Fig. 6. Layers of interferers with respect to USER A.

on a specific user (say, USER A) as shown in Fig. 6. If USER A is an active transmitter, due to the carrier sensing constraint, other simultaneously active transmitters must be outside USER A's sensing range, i.e., at a distance $\geq S\delta$ from USER A. From geometric considerations, it is easy to verify that there can be at most six other active interferers that are at a distance of $S\delta$ from USER A [13], [16], [25]; these form the *first layer* of interferers with respect to USER A (Fig. 6). If these six "interferers" are transmitting, it is easy to see that there can be at most 12 other "interferers" that are at a distance of at most $2S\delta$ from USER A, forming the *second layer* of interferers (Fig. 6), and so on. Therefore, the t th layer of interferers has $6t$ interferers that are at a distance of at most $tS\delta$ from USER A. Thus, with the maximum number of simultaneously active *ad hoc* transmissions, the SINR at any receiver is computed to be

$$\text{SINR} \leq \frac{\frac{P_{\text{USER}}}{(r\delta)^\alpha}}{\sum_{t=1}^{\infty} 6t \frac{P_{\text{USER}}}{(tS\delta)^\alpha}} = \frac{S^\alpha}{6r^\alpha \zeta(\alpha - 1)}. \quad (19)$$

With this, the maximum transmission rate that can be reliably sustained on an *ad hoc* link is

$$\text{Rate}_{\text{ad hoc}} \leq (1 - \gamma)B \times \log_2 \left(1 + \frac{S^\alpha}{6r^\alpha \zeta(\alpha - 1)} \right). \quad (20)$$

The second step towards computing ONE HOP THROUGHPUT is to compute RP RATIO and PROXY TX. Recall that RP RATIO is the maximum possible ratio of the number of relay transmissions per proxy transmission, such that every USER in the *ad hoc* component receives a unit of data traffic. Therefore, RP RATIO is given by

$$\text{RP RATIO} = \frac{\sum_{k=b+r+1}^n 6k \left\lceil \frac{k-(b+r)}{r} \right\rceil}{\sum_{m=b+1}^n 6m} \quad (21)$$

where the numerator reflects the number of relay transmissions and the denominator the number of proxy transmissions. Recall that PROXY TX is the sum of the fractions of the transmission




Fig. 7. Determining the number of interfering users.

times of all the proxies. Let l_k be the normalized carried load of a user in the k th tier of users in order for each user in the *ad hoc* component to receive a unit of data traffic; l_k is given by

$$l_k = \begin{cases} \left\lfloor \frac{n-k}{r} \right\rfloor \sum_{m=1}^{\lfloor \frac{n-k}{r} \rfloor} \frac{6(k+mr)}{6k}, & \text{if } b-r+1 \leq k \leq n-r \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

Essentially, l_k is equal to the number of more distant users to which a user in the k th tier of users will relay traffic. Note that when a proxy is forbidden to transmit due to the carrier sensing constraint, there could be at most *six* other active transmitters in its sensing range (in the case of the linear network, there were two). Therefore, an upper bound on the fraction of time that a proxy can transmit is the ratio of the proxy's normalized carried load to the sum of its normalized carried load and *one-sixth* of the normalized carried loads of all other users within its sensing range. However, the determination of the number of interferers in each tier of users within a proxy's sensing range is difficult since it depends on the position of proxy, the size of the sensing range, and the orientation of the interfering tiers. We seek to solve this problem by geometry. Fig. 7(a) depicts a cell with hexagons representing the tiers of users. The sensing region of a proxy is shown by the shaded circle (sensing circle). The highlighted line segments inside the circle capture the locations of the interferers. Due to the regular placement of users, if we are able to find the length of each of these line segments, the number of interferers in each tier of interest can be computed. Let us now focus on one particular tier of interferers (say, the t th tier) and a proxy in the p th tier. We depict this scenario in Fig. 7(b). In the figure, the hexagon in bold represents the t th tier of interferers and the circle in bold represents the sensing region of a proxy located in the p th tier. Then, the interferers of the proxy are located on the highlighted hexagonal line segments that are covered by the proxy's sensing region. Our target is to compute a lower bound on the length of the line segments. Note that t could be smaller than, larger than, or equal to p . We draw two dotted circles around the hexagon representing the t th tier, as shown in Fig. 7(b). The inner circle is the largest circle that is contained in the hexagon. The outer circle is the smallest circle that contains the hexagon. We consider two angles θ_{inner} and θ_{outer} . θ_{inner} is the angle between the lines joining the BS and the intersection points of the inner circle and the sensing circle. Similarly, θ_{outer} is the angle between the lines joining the BS

and the intersection points of the outer circle and the sensing circle. The angles are computed using the law of cosines

$$\theta_{\text{inner}} = 2 \times \cos^{-1} \left(\frac{\left(\frac{t\sqrt{3}}{2}\right)^2 + p^2 - S^2}{2 \times \frac{t\sqrt{3}}{2} \times p} \right) \quad (23)$$

$$\theta_{\text{outer}} = 2 \times \cos^{-1} \left(\frac{t^2 + p^2 - S^2}{2 \times t \times p} \right). \quad (24)$$

Since the line segments of interest must lie between the inner and the outer circles, the sum of their lengths must be greater than or equal to the length of the arc A (refer to Fig. 7); this is of length $\frac{t\sqrt{3}}{2} \times \min(\theta_{\text{inner}}, \theta_{\text{outer}})$. Recall that users are regularly placed; adjacent users are of unit distance (unit in δ) apart from each other. Therefore, given a proxy in the p th tier, the lower bound of the number of users from the t th tier that interferes with its transmission I_p^t is simply computed to be

$$I_p^t = \left\lfloor \frac{t\sqrt{3}}{2} \times \min(\theta_{\text{inner}}, \theta_{\text{outer}}) \right\rfloor. \quad (25)$$

Note that the sensing circle might not intersect with both the inner and the outer circle for some tiers, and either one or both angles could be undefined. Since, in this case, the lengths of the line segments that are covered by the sensing circle are very small in general, we can simply ignore the interferers; this is a conservative lower bound on the number of interferers. With this, PROXY TX is computed as follows:

$$\text{PROXY TX} \leq \sum_{p=b-r+1}^{\min(b,n-r)} 6p \times \frac{l_p}{l_p + \frac{1}{6} \left[\left(\sum_{t=1}^n I_p^t \times l_t \right) - l_p \right]}. \quad (26)$$

Then, ONE HOP THROUGHPUT is bounded by using (13), and an upper bound on λ_A is computed using Inequality (14). We defer a discussion on numerical estimates of the capacity and the interpretations thereof to the next section.

V. EVALUATIONS

In this section, we present numerical results derived from our analytical models. In addition, we perform simulations to validate our analysis. We use the ns-2 simulator. We first show that the simulation results with regular placements of users fit the analytical results very well. We also perform simulations with random placements of users and show that the obtained results are similar to those with the regular placements; this in turn suggests that our analytical results are relevant even when the users are placed randomly.

The close match between the simulations and the analytical results is due to the following reasons. First, in both the simulations and analysis, the Shannon formula is used to compute the capacity of each link. More importantly, the interference diminishes drastically with distance. In other words, the more distant BSs project little interference (the approximations consider this, but this is ignored in simulations). Finally, we assume ideal scheduling that provides the maximum spatial reuse (as is assumed in the analysis) in the simulations.

¹⁰<http://www.isi.edu/nsbam/ns/>.

A. Simulation Settings

In the simulations, an ideal “reliable” media access protocol is assumed. A fair schedule (global knowledge is assumed) that provides the max-min fairness is implemented. Minimum-hop routing is enforced in the *ad hoc* component (in line with the analysis). The results with random placements of users are averages from 20 simulations, each of which uses a randomly generated scenario and runs for 15 s.

With the linear hybrid network, the diameter of a cell is 1 km (i.e., $D = 500$ m) and there are 100 users in each cell. For the regular placement of users, the 100 users are placed on a straight line with a BS at the center of the cell; thus, there are 50 users in each side of the BS (i.e., $n = 50$). The distance between adjacent users is fixed at 10 m (i.e., $\delta = 10$ m). For the random placement of users, each of the 100 users selects a random position within the cell. The size of the BS footprint b' ¹¹ varies from 100 to 500 m. The transmission range of an *ad hoc* link r' varies from 100 to 450 m, and the sensing range S' is fixed at 450 m.

In the hexagonal hybrid network, the diameter of a cell is 2 km (i.e., $D = 1000$ m). In the first set of experiments, there are ten tiers of users (i.e., $n = 10$), and the distance between adjacent users is 100 m (i.e., $\delta = 100$ m). With these settings, the total number of users in a cell is 330. For random placements of users, each of the 330 users is randomly positioned within the hexagonal cell. Note here that this number is reasonable, since the BS's footprint in today's cellular networks covers roughly 2 km, and one might expect the number of users to be fairly large. The size of the BS footprint b' varies from 100 to 1000 m. The transmission range of an *ad hoc* link r' varies from 100 to 400 m, and the sensing range S' is fixed at 400 m.

In both the linear and hexagonal hybrid networks, the total spectral bandwidth B is 1.25 MHz and the spectral allocations that we consider are $\gamma = 0.1, 0.3, 0.5, 0.7$, and 0.9 . The path-loss exponent $\alpha = 4$, unless otherwise stated.¹² Without loss of generality, we use the same value of α for both the direct BS-to-user link and the intrauser links. It is easy to extend the results if the value is different in the two cases.

B. Optimizing the Transmission Range of an Ad Hoc Link

An increase in the transmission range of an *ad hoc* link reduces the total relay load in the *ad hoc* component [see (7) and (18)]; however, it also reduces the transmission rate that can be sustained on each link [see (9) and (20)]. Striking the balance between the two is necessary for the highest capacity. To understand this effect, we compute the capacity of the hybrid network (using our analysis) by varying the transmission range of the *ad hoc* links in an exhaustive set of scenarios; our goal is to determine the optimal transmission range in each of these scenarios. Figs. 8 and 9 depict the transmission range that maximizes the capacity of the linear and hexagonal network, respectively. We observe that the optimal transmission range varies with the size of the BS footprint and the fraction of the spectrum allocated to the infrastructure component (γ). First, note that *the optimal transmission range decreases when γ increases*. When γ is large, the portion of the spectrum that is allocated for *ad*

¹¹Recall that in our framework, b , r , and S are defined in terms of units of length δ ; we will use b' , r' , and S' to represent the actual values (in meters) of these variables. The values of b , r , and S can easily be computed.

¹²We also examine the impact of varying the path-loss exponent later.




Fig. 8. Optimal transmission range of an *ad hoc* link changes with BS footprint and spectrum allocation. Note that $\delta = 100$ m and $S' = 400$ m (linear network).




Fig. 9. Optimal transmission range of an *ad hoc* link changes with BS footprint and spectrum allocation. Note that $\delta = 100$ m and $S' = 400$ m (hexagonal network).

hoc transmissions is small. In this regime, in order to achieve a higher λ_A , a smaller transmission range should be used in order to increase the link rate of *ad hoc* transmissions. Secondly, we observe that the optimal transmission range increases with the size of the BS footprint. When the size of the BS footprint increases, the bottleneck on the capacity of the hybrid network is the capacity of the infrastructure component since the BS now supports a large number of directly connected users, to which it has links with poor channel qualities. In this regime, choosing a larger transmission range for *ad hoc* link transmissions will lead to a higher λ_I . This is because, with a larger transmission range, users at the boundary of the infrastructure and the *ad hoc* components can reach proxies that are closer to the BS. These proxies have higher ILCs, and this, in turn, translates to a higher network capacity. This can be easily seen from the maximization problem defined in (3) and (15).

C. Capacity of the Linear Hybrid Network

The capacities with the linear hybrid network in various scenarios are shown in Fig. 10. Each of the points in the figure is obtained by using the optimal transmission range of an *ad hoc* link for the corresponding scenario. Therefore, the figure depicts the maximum achievable uniform throughput of the linear hybrid network while varying the size of the BS footprint and the spectrum allocation factor γ . Simulation results are shown along with the analytical computations. We see that the simulation results with both regular and random placements of users exhibit very similar behaviors and conform with the analytical




Fig. 10. The results of analysis and simulations with a linear hybrid network. The results are obtained by using the optimal transmission range of an *ad hoc* link at the corresponding scenario.

results. The horizontal line in the figure corresponds to the uniform throughput of the pure cellular network.¹³ The results show that the linear hybrid network, in most of the scenarios, does not provide higher capacities than the pure cellular network; this is especially true when γ is small (less than 0.7). This is because the capacity bottleneck appears in the infrastructure component, i.e., the BS cannot deliver traffic at the rates possible with the pure cellular network. For $\gamma = 0.5$ to 0.9, *sweet spots* (where the uniform throughput reaches a maximum) appear. A sweet spot is a result of two contradicting factors. Increasing the size of the BS footprint reduces λ_I , since now, a larger number of users with poor channel qualities will be directly reached by the BS. In contrast, it increases λ_A , since the total relay load is greatly reduced. Clearly, the sweet spot represents the operational point that provides the best tradeoff between the conflicting effects. We see that the linear hybrid network outperforms the pure cellular network only in a few limited scenarios. Moreover, the improvement is less than 10%. The poor capacity with the linear hybrid network is due to the poor spatial reuse in the *ad hoc* component.

D. Capacity of the Hexagonal Hybrid Network

For the hybrid network, the analytical results and the simulation results with both the regular and random placements of users are depicted in Fig. 11. In the figure, we see that the simulation results with regular placements of users exhibit similar behaviors and conform with the analytical results.¹⁴ The slight discrepancies between the results are due to the approximation in the analysis that the ILCs of users that are located within the same tier are the same and equal to the best ILC possible for any user within the tier (see Section IV). Note that the results depicted in Fig. 11 are obtained by selecting the *ad hoc* transmission range that maximizes the uniform throughput of the hybrid network. The uniform throughput of the corresponding pure cellular network is represented by the bold solid lines. The hybrid network yields capacity improvements over the corresponding pure cellular network in certain parametric regimes. However,

¹³When the pure cellular network is considered, the footprint corresponds to the complete cell; in other words, the footprint of the BS is the cell itself. Note here that we assume that the total number of BSs in the network is fixed. We believe that this is a reasonable assumption since BSs are generally expensive and increasing their number would drastically increase deployment cost.

¹⁴When $r' = 100$ m, some of the users are disconnected in the *random* scenarios due to the short transmission range of an *ad hoc* link. In this case, the uniform throughput of the hybrid network is thus equal to zero, as per our definition.

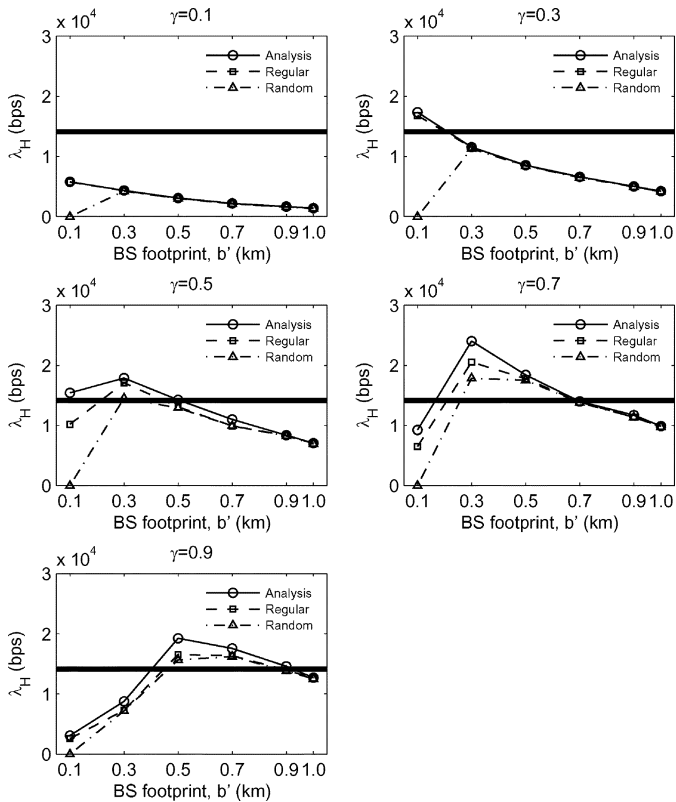


Fig. 11. The uniform throughput of the hybrid network with different sizes of BS footprint (b') and spectrum allocations (γ). An optimal transmission range of an *ad hoc* link is used in each of the scenarios. The solid line represents the capacity of the corresponding cellular network.

due to the capacity bottleneck of the infrastructure component when γ is small ($\gamma \leq 0.3$), the uniform throughput of the hybrid network is poor in general. However, when γ increases, the hybrid network begins to outperform the pure cellular network, as long as the BS footprint is not too large. When γ increases, the size of the BS footprint for achieving the maximum uniform throughput also increases. This is because a higher γ means a lower *ad hoc* component capacity. In order for the *ad hoc* component to achieve a high uniform throughput, the size of the BS footprint must be increased so that the relay load in the *ad hoc* component decreases. Again, due to the two contradicting factors that are in play (discussed previously), *sweet spots are observed with relatively large γ values*. Generally, the improvements with the hybrid network over the corresponding original pure cellular network can be as great as 70% when γ , b , and r are appropriately selected. The capacity gains with the hybrid network (as compared with the linear hybrid network) are due to the better exploitation of spatial reuse in the *ad hoc* component (more simultaneous *ad hoc* transmissions can happen) and the increased bit rates due to the use of shorter range links.

E. Comparison of Spatial Reuse

In order to validate the claim that the benefits in the hexagonal network (as compared to the case of the linear network) are due to a better exploitation of spatial reuse, we compare the spatial reuse in the two networks. We consider different sizes of the BS footprint. In our framework, once packets are injected into the *ad hoc* component by the proxies, the packets are always delivered at a destination; we do not impose any bound

TABLE I
SPATIAL REUSE OF PROXIES (HEXAGONAL, $r' = 250$ m)

BS footprint (b')	250m	300m	350m	400m	450m
Analysis	6.258	6.291	6.294	6.307	7.088
Simulation (Regular)	5.852	5.862	5.818	5.728	5.939
Simulation (Random)	5.664	5.828	5.811	5.753	6.006

on delays. With this, the amount of traffic that is delivered by the proxies dictates the maximum uniform throughput that is achieved on the *ad hoc* component (see Fig. 4). Therefore, computing the spatial reuse of the proxies (i.e., the number of concurrent proxy transmissions) is sufficient to capture the spatial reuse of the *ad hoc* component. We only show the number of concurrent proxy transmissions in the hexagonal hybrid network with different sizes of the BS footprint (Table I). One can easily see that there can be at most two concurrent proxy transmissions in the linear network (see Section III). We see that the results of simulations match the results from the analysis quite well. The slight discrepancy between the simulation results with the analytical results is mainly due to an approximation that was made in determining the number of interfering users in (25). Specifically, the approximation considers a lower number of interferers than what may be expected if one were to perform simulations. Also, the discrepancy between the simulation results with random placements of users and that with random placements of users is due to the increased number of hops on the minimum-hop routes, on average, in the random scenarios. When the average number of hops increases, the normalized carried load of each user will increase (since a packet will have to traverse more hops) and, in turn, reduce the fraction of time that the proxies can transmit. We see that the spatial reuse of the hexagonal hybrid network (as reflected by the number of concurrent proxy transmissions) is always greater than that of the linear hybrid network (which has at most two).

F. Regular Versus Random Deployment

We have shown that the simulation results with random placements of users conform with the analytical results as well as the simulation results with deterministic placements of users. Due to the uniform distribution of users, the topological properties of the network with random placements are similar to those of the network with regular placements of users. As indicated in [26], the connectivity of the *ad hoc* stations to the proxies with random placements of users is almost the same as that with deterministic placements of users. In the case of random deployments, however, if the transmission range of a user is too short relative to the interuser distance, the connectivity of the users to the proxies will degrade. In other words, one must ensure that the transmission range of a user is long enough to ensure that all users can reach the proxies with high probability. In our framework, the transmission range of a user is $r\delta$ and the interuser distance is δ . In the deterministic setting, it suffices to have $r = 1$ in order to ensure complete connectivity. We observe that with the random deployment, a value of $r = 2$ is enough to provide complete connectivity with a probability ≈ 1 . We observe that this results in an increase, on average, in the total relay load from 0.6% to 7.1% in the linear hybrid network and from 0.6% to 16% in the hexagonal hybrid network. This results in a slight degradation in capacity, as seen in Figs. 10 and 11.




Fig. 12. The maximum capacity improvements provided by the hybrid network over the corresponding pure cellular network with (a) different density of users (cell radius = 1000 m) and (b) different cell sizes [user density is fixed ($\delta = 100$ m)].

G. Impact of Density of Users in the Network

Next, we investigate the effects of varying the user density on the capacity of the hexagonal hybrid network. Given the poor performance of the linear hybrid network, we do not study it further; however, we expect the behavioral results to be similar in that case. In Fig. 12(a), we depict the *maximum* capacity improvements of the hexagonal hybrid network over the corresponding pure cellular network with different user densities (varied by adjusting the number of tiers of users in a cell). When the user density is low, the capacity improvements are also low. In this regime, since adjacent users are separated by large distances, a large transmission range has to be used on the *ad hoc* links in order to maintain connectivity. This in turn limits the achievable transmission rates on the *ad hoc* links and, thus, reduces the capacity. On the other hand, as the user density increases, the capacity improvements become better. As the user density increases, the distance between adjacent users decreases, and this leads to two advantages: first, one can now establish shorter links; thereby, higher transmission rates can be achieved. Secondly, it improves the flexibility/granularity of the transmission range selection. While the benefit provided by the second advantage is not immediately apparent, it is vital for maximizing the uniform throughput in the *ad hoc* component (note as discussed earlier, the capacity varies with r). Note that after a certain density, the capacity improvements converge to a stable saturated value ($\approx 70\%$). This is because the rates cannot be increased beyond a certain level (as dictated by the Shannon capacity and the available raw bandwidth). Furthermore, since the transmission power is fixed, the maximum spatial reuse cannot be increased beyond a certain point. As we increase node density, the reuse increases initially and reaches this point; if the node density is increased further, the spatial reuse will not further increase. Thus, the capacity gain will not increase either.

H. Impact of Cell Size

We evaluate the impact of the cell size on the maximum capacity improvements that can be achieved with the hexagonal hybrid network. In Fig. 12(b), we depict the results by varying the cell radius. Note that the user density is fixed ($\delta = 100$ m). Therefore, when the size of the cell increases, the number of tiers of users in a cell increases. From the results, we see that the capacity improvements generally increase with an increase in the cell size. This seems counterintuitive at first; however, note that the capacity of the corresponding cellular network decreases when the cell size increases since there are more

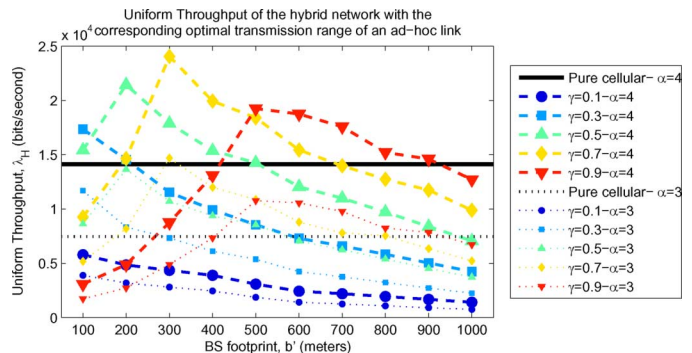


Fig. 13. The effect of path-loss exponent at the hexagonal hybrid network.

users with poor channel qualities. With the hybrid network, these users benefit to a greater extent from the high-rate *ad hoc* transmissions. It may seem that the penalty due to multihop transmissions will increase with cell size and could hurt capacity; however, since the sensing range is fixed, increasing the cell size can actually increase the number of simultaneous *ad hoc* transmissions in the network. The combined effects result in a slight increase in the capacity of the hybrid network when the cell size increases.

I. Impact of Path-Loss Exponent

In Fig. 13, we present the results for the case of the hexagonal hybrid network for two different values of the path-loss exponent ($\alpha = 3$ and $\alpha = 4$). Our results indicate that a smaller path-loss exponent results in reduced absolute capacity due to an increased impact of interference due to decreased attenuation with distance. As a consequence, the spatial reuse is decreased. Note, however, that this effect impacts the pure cellular network as well. The gains from the usage of a hybrid network are thus still viable.

VI. RELATED WORK

Most previous efforts on computing the capacity of wireless networks consider either a pure *ad hoc* or a cellular network [23], [27]. There are a limited number of efforts on various kinds of hybrid networks. The distinguishing aspects of our work compared to these efforts were highlighted in Section I.

We discuss the previous efforts in brief below and elaborate on how our work differs from these efforts in specific aspects.

A. Capacity of Hybrid Wireless Networks

In [1], a hybrid network with a sparse layout of base stations connected by high-bandwidth links is studied. It is shown that with n *ad hoc* nodes and m base stations, the benefits provided due to the base stations are insignificant when m grows asymptotically slower than \sqrt{n} . In [9], in a similar network setting, it is shown that $\Theta(1)$ per node throughput can be achieved when the nodes can optimally control their transmission powers. In [10], it is shown that a hybrid network with n wireless nodes and n^d access points, interconnected by wires, can achieve throughput gains only when $(1/2) < d < 1$. Unlike our work, the above efforts assume that the traffic flows are established between randomly chosen *ad hoc* nodes and the base stations or access points are for relaying purposes only.

In [8] and [12], the authors consider a hybrid network that is similar to what we consider, but with a small number of relays and a small number of hops. However, they *do not consider rate variability with range*; furthermore, the transmission range of relays is always equal to the distance between the relays. Furthermore, they *do not consider load heterogeneity of relays* due to the downlink traffic flow pattern. Reference [8] provides a simple informal analysis on the capacity gain and evaluates the throughput gain of the network with four relays. The work in [12] evaluates the throughput gain of the network by considering geographical routing. In [11], the authors evaluate the tradeoffs between channel reuse efficiency and increased interference in simple and small hybrid network topologies. In [28], the authors consider the capacity improvements of a CDMA-based cellular network when it is augmented with *ad hoc* relays. However, unlike in this paper, they allow only for one "*ad hoc*" hop and they assume out-of-band relays. Liu *et al.* [29] investigate the impact of network dimensionality and geometry on the capacity of hybrid wireless networks. They show that different dimensions can lead to significantly different scaling laws with regards to capacity. However, they only consider one- and two-dimensional grid and strip networks. Furthermore, they do not explicitly examine the impact of rate versus range tradeoffs on the capacity of the hybrid network as we do here.

B. Hybrid Network Architectures

In [3], a multihop cellular network (MCN) is proposed. It is shown that the throughput with the MCN exceeds that of the pure cellular network. However, it is assumed that the *transmission range of the BSs is the same as that of the mobile stations (MSs)*. BSs are for *relaying purpose only*. In [4], an integrated cellular and *ad hoc* relaying system is proposed; *ad hoc* relaying stations are strategically deployed to improve call service quality; capacity gain is not the main focus of this work. The work described in this paper differs from the above efforts in that the transmission range of a BS (the size of the BS footprint) can be varied and is independent to the transmission range of the MSs. Furthermore, these efforts assume peer-to-peer traffic patterns and do not consider rate variability with range.

In [5], the authors consider the *ad hoc* network model for wireless packet data networks and show that this can support better spatial reuse and result in better throughput per unit power spent. However, again, the BSs are for relaying purposes only. In [2], the authors propose a unified cellular and *ad hoc* network (UCAN) architecture to improve individual user throughput and the aggregate downlink throughput. However, UCAN assumes the use of an additional spectral band for *ad hoc* communications.

VII. CONCLUSION

In this paper, we consider an approach of augmenting cellular networks with *ad hoc* wireless connectivity to improve spatial reuse. In particular, we reduce the coverage area of the base station and require that users outside the area rely on other users within the area for connectivity. Our objective is to evaluate if this approach can yield capacity benefits over that of the original cellular network. While the use of the relays to form a hybrid network provides shorter higher rate communication links, multihop forwarding contributes to a reduction in capacity. We analytically compute the capacity of a regular linear and hexagonal

hybrid network under the conditions of max-min fairness. We also perform simulations with both deterministic and random placements of users. We observe that the linear hybrid network does not provide significant capacity improvements over the pure cellular network. However, the hexagonal hybrid network provides higher capacities than a corresponding pure cellular network by better exploiting the spatial reuse in certain parametric regimes of operation. With careful parametric choices, the capacity of the hexagonal hybrid network can exceed that of the corresponding pure cellular network by as much as 70%.

ACKNOWLEDGMENT

The authors wish to thank M. Grob, R. Pankaj, D. Schultz, V. Gupta, and other QUALCOMM staff for suggestions and comments.

REFERENCES

- [1] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *Proc. IEEE INFOCOM*, San Francisco, CA, 2003, vol. 2, pp. 1543–1552.
- [2] H. Luo, R. Ramjee, P. Sinha, L. Li, and S. Lu, "UCAN: A unified cellular and ad-hoc network architecture," in *Proc. ACM MobiCom*, San Diego, CA, 2003, pp. 353–367.
- [3] Y.-D. Lin and Y.-C. Hsu, "Multihop cellular: A new architecture for wireless communications," in *Proc. IEEE INFOCOM*, Tel Aviv, Israel, 2000, vol. 3, pp. 1273–1282.
- [4] H. Wu, C. Qiao, S. De, and O. Tonguz, "Integrated cellular and ad hoc relaying systems: iCAR," *IEEE J. Sel. Areas Commun.*, vol. 19, no. 10, pp. 2105–2115, Oct. 2001.
- [5] H.-Y. Hsieh and R. Sivakumar, "On using the ad-hoc network model in cellular packet data networks," in *Proc. ACM MobiHoc*, Lausanne, Switzerland, 2002, pp. 36–47.
- [6] S. Lee, S. Banerjee, and B. Bhattacharjee, "The case for a multi-hop wireless local area network," in *Proc. IEEE INFOCOM*, Hong Kong, China, 2004, vol. 2, pp. 894–905.
- [7] R. Chakravorty, S. Agarwal, S. Banerjee, and I. Pratt, "MoB: A mobile bazaar for wide-area wireless services," in *Proc. ACM MobiCom*, Cologne, Germany, 2005, pp. 228–242.
- [8] H. Viswanathan and S. Mukherjee, "Performance of cellular networks with relays and centralized scheduling," in *IEEE Veh. Technol. Conf.*, Orlando, FL, 2003, vol. 3, pp. 1923–1928.
- [9] A. Agarwal and P. R. Kumar, "Capacity bounds for ad hoc and hybrid wireless networks," *Comput. Commun. Rev.*, vol. 34, no. 3, pp. 71–81, Jul. 2004.
- [10] S. Toumpis, "Capacity bounds for three classes of wireless networks: Asymmetric, cluster, and hybrid," in *Proc. ACM MobiHoc*, Tokyo, Japan, 2004, pp. 133–144.
- [11] J. Cho and Z. J. Haas, "On the throughput enhancement of the down-stream channel in cellular radio networks through multihop relaying," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 7, pp. 1206–1219, Jul. 2004.
- [12] S. Mukherjee and H. Viswanathan, "Analysis of throughput gains from relays in cellular networks," in *Proc. IEEE GLOBECOM*, San Francisco, CA, 2003, pp. 3471–3476.
- [13] H. Zhai and Y. Fang, "Physical carrier sensing and spatial reuse in multirate and multihop wireless ad hoc networks," in *Proc. IEEE INFOCOM*, Barcelona, Spain, 2006, pp. 1–12.
- [14] S. Sarkar and K. N. Sivarajan, "Fairness in cellular mobile networks," *IEEE Trans. Inf. Theory*, vol. 48, no. 8, pp. 2418–2426, Aug. 2002.
- [15] A. Balachandran, G. Voelker, P. Bahl, and P. Venkat Rangan, "Characterizing user behavior and network performance in a public wireless LAN," in *Proc. ACM SIGMETRICS*, Marina del Rey, CA, 2002, pp. 195–205.
- [16] X. Yang and N. H. Vaidya, "On physical carrier sensing in wireless ad hoc networks," in *Proc. IEEE INFOCOM*, Miami, FL, 2005, pp. 2525–2535.
- [17] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [18] J. Camp and E. W. Knightly, "Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation," in *Proc. ACM MobiCom*, San Francisco, CA, 2008, pp. 315–326.

- [19] A. Bedekar, S. Borst, K. Ramanan, and P. Whiting, "Downlink scheduling in CDMA DATA networks," in *Proc. IEEE GLOBECOM*, Rio de Janeiro, Brazil, 1999, pp. 2653–2657.
- [20] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, "CDMA/HDR: A bandwidth-efficient high-speed wireless data service for nomadic users," *IEEE Commun. Mag.*, vol. 38, no. 7, pp. 70–77, Jul. 2000.
- [21] S. Lu, V. Bharghavan, and R. Srikant, "Fair scheduling in wireless packet networks," *IEEE/ACM Trans. Netw.*, vol. 7, no. 4, pp. 473–489, Aug. 1999.
- [22] T.-S. Kim, H. Lim, and J. C. Hou, "Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks," in *Proc. ACM MobiCom*, Los Angeles, CA, 2006, pp. 366–377.
- [23] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.
- [24] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proc. ACM MobiCom*, Rome, Italy, 2001, pp. 61–69.
- [25] P. Santi, *Topology Control in Wireless Ad Hoc and Sensor Networks*. New York: Wiley, 2005.
- [26] J. Robinson and E. Knightly, "A performance study of deployment factors in wireless mesh networks," in *Proc. IEEE INFOCOM*, Anchorage, AK, 2007, pp. 2054–2062.
- [27] T. Bonald and A. Proutiere, "Wireless downlink data channels: User performance and cell dimensioning," in *Proc. ACM MobiCom*, San Diego, CA, 2003, pp. 339–352.
- [28] D. Zhao and T. D. Todd, "Cellular CDMA capacity with out-of-band multihop relaying," *IEEE Trans. Mobile Comput.*, vol. 5, no. 2, pp. 170–178, Feb. 2005, in , vol. 5.
- [29] B. Liu, P. Thiran, and D. Towsley, "Capacity of a wireless ad hoc network with infrastructure," in *Proc. ACM MobiHoc*, Montreal, PQ, Canada, 2007, pp. 239–246.



hybrid networks.

Lap Kong Law received the B.S. degree in computer science from The Chinese University of Hong Kong in 2002, and the Ph.D. degree in computer science and engineering from the University of California, Riverside, in 2007.

He is currently a Software Engineer with the Wireless Control Team, Trapeze Networks. His research interests include routing protocols for unicast, broadcast, and multicast communications in ad hoc networks; TCP performance over wireless networks; and capacity analysis of cellular/ad hoc



ysis. He is also interested in mathematical foundations of communication networks.

Konstantinos Pelechrinis (S'07) received the diploma in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2006, and the M.Sc. degree from the Computer Science Department, University of California, Riverside, in 2008, where he currently is pursuing the Ph.D. degree.

His research interests include wireless networking, especially security-related issues that span the full protocol stack. He is involved in protocol design, real-world experimentation, and performance anal-



Hoc Networks: Technologies and Protocols (Berlin, Germany: Springer-Verlag, 2005). He is the Associate Editor-in-Chief for *ACM MC2R*.

Dr. Krishnamurthy received a National Science Foundation CAREER Award from ANI in 2003.

Srikanth V. Krishnamurthy (SM'07) received the Ph.D. degree in electrical and computer engineering from the University of California at San Diego in 1997.

From 1998 to 2000, he was a Research Staff Scientist with the Information Sciences Laboratory, HRL Laboratories, LLC, Malibu, CA. Currently, he is an Associate Professor of computer science at the University of California, Riverside. His research interests are primarily in wireless networks, network security, and Internet technologies. He is Co-Editor of *Ad*

Michalis Faloutsos (M'09) received the Bachelor's degree from the National Technical University of Athens, Athens, Greece, in 1993, and the M.Sc. and Ph.D. degrees from the University of Toronto, Toronto, ON, Canada, in 1995 and 1999, respectively.

He is a Faculty Member with the Computer Science Department, University of California, Riverside. He is actively involved in the community as a Reviewer and a Technical Program Committee Member for many conferences and journals. With

his two brothers, he co-authored a paper on powerlaws of the Internet topology (SIGCOMM'99), which is one of the top 10 most cited papers of 1999. His most recent work on peer-to-peer measurements has been widely cited in popular printed and electronic press such as slashdot, *ACM Electronic News*, *USA Today*, and *Wired*. His interests include Internet protocols and measurements, peer-to-peer networks, network security, BGP routing, and ad hoc networks. Most recently, he has focused on the classification of traffic and identification of abnormal network behavior. He also works in the area of Internet routing, ad hoc networks routing, and network security, with emphasis on routing.