

Performance Comparison of Cellular and Multi-hop Wireless Networks: A Quantitative Study*

Hung-Yun Hsieh and Raghupathy Sivakumar

School of Electrical and Computer Engineering
Georgia Institute of Technology

{hyhsieh, siva}@ece.gatech.edu

ABSTRACT

In this paper we study the performance trade-offs between conventional cellular and multi-hop ad-hoc wireless networks. We compare through simulations the performance of the two network models in terms of raw network capacity, end-to-end throughput, end-to-end delay, power consumption, per-node fairness (for throughput, delay, and power), and impact of mobility on the network performance. The simulation results show that while ad-hoc networks perform better in terms of throughput, delay, and power, they suffer from unfairness and poor network performance in the event of mobility.

We discuss the trade-offs involved in the performance of the two network models, identify the specific reasons behind them, and argue that the trade-offs preclude the adoption of either network model as a clear solution for future wireless communication systems. Finally, we present a simple hybrid wireless network model that has the combined advantages of cellular and ad-hoc wireless networks but does not suffer from the disadvantages of either.

1. INTRODUCTION

In recent years, the proliferation of mobile devices like handheld PCs and PDAs has resulted in the rapid evolution of wireless packet data networks. Such networks have typically assumed a cellular network model to provide wireless access to the mobile devices. The model consists of a base station or an access point covering a (circular) service area, with all subscribed mobile devices directly communicating with the base station. The base station in turn might be connected to a backbone network. When both source and destination hosts reside in the same cell, the base station acts as a relay between the source and destination. However, when the destination is not within the same cell as the

source, the base station forwards packets from the source to a distribution network in the backbone¹. Wireless data networks are typically characterized as either local-area or wide-area based on the size of their coverage cells.

Some of the key reasons behind the adoption of the cellular model for wireless data networks include the ability to simply reuse existing voice network infrastructures for packet data [18], and the simplicity of the model due to the presence of a central coordinating entity – the base station. However, the inability of wireless data networks to scale to high data rates and thus to sustain the accelerated growth in the number of users² has prompted researchers to explore alternate network models to improve network performance [3], [9], [13], [15]. Most, if not all, of the alternate models are based on reducing cell sizes to increase the amount of spatial reuse³ in the network, thereby enhancing network capacity.

In this context, a special class of networks called ad-hoc or multi-hop wireless networks [1], [16] have gained attention by virtue of their ability to operate with a peer-to-peer network model, where the mobile devices potentially use a minimal transmission range just large enough for the network to be connected. In the absence of any established infrastructure (like the base station), the hosts in an ad-hoc network communicate with each other over multi-hop paths consisting of other hosts in the network. Using cells sizes just enough to maintain network connectivity allows ad-hoc networks to potentially maximize spatial reuse.

While we present a detailed discussion of related work later in this paper, to the best of our knowledge, there has been no work that comprehensively compares the performance trade-offs between conventional cellular networks and multi-hop ad-hoc networks. We believe that such a comparison is warranted and can provide valuable insight towards the development of alternate and better wireless network models. Towards this end, we make two contributions in this paper:

* This work was partly supported by funds from the Yamacraw Organization (<http://www.yamacraw.org>).

¹ In WWANs, the base stations are typically connected to a private network of the wireless network provider that in turn is connected to the backbone network.

² Next generation wireless systems, popularly known as 3G wireless [17], have projected data rates of merely 350 Kbps outdoors and 2 Mbps indoors.

³ Spatial reuse is the ability to have multiple simultaneous transmissions in non-overlapping sections of the network.

- (i) *We present a comprehensive, simulations-based comparison of the performance trade-offs involved in the cellular and ad-hoc network models.* We compare the models along different dimensions including: raw network capacity, end-to-end throughput, end-to-end delay, power consumption, fairness, and impact of mobility on network performance. We first show that, as expected, ad-hoc networks can maximize spatial reuse and hence increase raw network capacity. More interestingly, we demonstrate that the multi-fold increase in raw network capacity does not necessarily translate into an equivalent increase in the end-to-end throughput. Furthermore, we show that ad-hoc networks do perform better than cellular networks in terms of end-to-end throughput, delay, and power consumption. On the other hand, ad-hoc networks fare badly in terms of the fairness they offer to the network nodes, and are vulnerable to frequent network partitions and route failures that essentially prohibit them from being used in critical (and commercial) applications where connectivity is a requirement.
- (ii) While the focus of this paper is to compare the two network models, we use the insight gained from the comparison results as the basis to *propose a simple hybrid wireless network model* that offers the performance benefits of ad-hoc networks without exhibiting the associated connectivity problems. Briefly, the hybrid network model retains the physical infrastructure of a conventional cellular network. However, the network hosts by default operate in a multi-hop ad-hoc mode where the base station merely acts as another host in the network. The mobile hosts use a special dedicated control channel to keep the base station updated about their local topology information. The base station periodically checks the status of the network and directs the mobile hosts to increase or decrease their transmission power levels depending on whether or not the network is partitioned. When the transmission power reaches a level at which the throughput of the network goes down below that of a corresponding network operating in cellular mode, the base station switches the network to operate in cellular mode. If at a later point, the topology reverts to an ad-hoc friendly configuration, the base station switches the network back to ad-hoc mode. Ideally, *the hybrid model's performance will track the bounding envelope of the superposition of the performance curves of the cellular and ad-hoc network models.* We present simulation results that demonstrate the performance achieved by the hybrid network model.

The scope of this paper is, however, limited to comparing the performance of the two network models within a single cell. Hence, the results presented in the paper are suited for local-area and wide-area wireless networks in which the sender and the receiver are present within the same cell⁴.

The rest of this paper is organized as follows: In Section 2, we present a discussion on related work in this area. In Section 3, we describe the network models, algorithms, and metrics used in the simulations. In Section 4, we present the simulation results and

interpretations. In Section 5, we describe the hybrid network model and simulation results that demonstrate its performance. Finally, in Section 6, we conclude the paper.

2. RELATED WORK

In this section, we discuss some related works that address the issue of increasing the capacity of wireless networks by exploiting spatial reuse and other mechanisms.

The hierarchical cellular network model (HCN) [15] is the outcome of efforts in the traditional voice networks community to increase spatial reuse to provide high data rates, especially in regions with dense populations of mobile devices. While the HCN model is conceived with the purpose of leveraging spatial reuse (at the lower levels), it is different from the approach presented in this paper in the following ways: (i) The model is still based on the conventional cellular model except for reduced cell sizes. However, since the cell sizes are constant, the spatial reuse of a hierarchical cellular network is fixed irrespective of the population or the topology of the network. Therefore, such networks cannot realize the maximum spatial reuse possible. (ii) Hierarchical cellular networks increase spatial reuse at the expense of a more elaborate backbone infrastructure comprising of a large number of base stations. (iii) The HCN model uses a multi-channel approach that can potentially lead to under-utilization of the network capacity (in the event that one of the levels in the cell hierarchy is not fully utilized).

In [9], the authors propose a new wireless network model called the multi-hop cellular network (MCN). Briefly, the model involves mobile devices farther away from the base station communicating with the base station using a multi-hop path consisting of other mobile devices in the cell. The MCN model, however, does not address the key drawback of ad-hoc networks, namely the vulnerability to frequent network partitions and route failures due to mobility, and it remains susceptible to network partitions. Further, the authors do not explicitly address the issue of minimizing transmission ranges to maximize spatial reuse. Hence the network model may not achieve the maximum spatial reuse (and thus network capacity) possible. Finally, the authors do not compare the two network models based on other metrics like delay and power consumption, and do not evaluate the performance of the hybrid model in terms of the fairness offered to network nodes.

In [13], the authors propose a network model called iCAR that involves special stations called ad-hoc relay stations (ARSs) to complement a conventional cellular network model. The ARSs are strategically placed between regular cells in the network and help relay data in a congested cell through a neighboring non-congested cell. However, this approach requires the underlying channel to be split between the regular base stations and the ARSs, thus reducing capacity in cells that do not require an ARS or cannot be served by an ARS. (This can be due to multiple reasons including unavailability of an ARS, heavily loaded neighboring cells, etc.)

Lastly, in [3], the authors analytically derive several properties pertaining to the capacity of wireless networks. We verify several of their analytical results in this paper through the interpretation of our simulation results.

⁴ In a later work we have studied the impact of traffic locality (with sources and destinations potentially not within the same cell) on the performance of the ad-hoc network model [6].

3. SIMULATION MODEL

We use the *ns2* (version 2.1b6) network simulator [12] developed at the Lawrence Berkeley National Laboratory in collaboration with the VINT project for all simulations. Wireless extensions from contributions of the CMU Monarch Project are widely used to simulate mobile wireless networks. Details about the simulation model and environment are presented in the rest of this section.

3.1 Network Models

The physical layer implementations of *ns2* include a radio interface that models the node as a Lucent 915 MHz WaveLAN device following IEEE 802.11's 2 Mbps *Direct Sequence Spread Spectrum* (DSSS) specifications. The channel model consists of a combination of a free space propagation model and a two-ray ground reflection model. The crossover point for the two models is called the *reference distance*. When a transmitter is within the reference distance, the free space model where the signal attenuates as $1/r^2$ is used. Outside of this distance, the two-ray ground reflection model where the signal falls off as $1/r^4$ is used. Unless otherwise specified, all the simulations for multi-hop ad-hoc networks use the IEEE 802.11 MAC protocol in the *Distributed Coordination Function* (DCF) mode and the *Dynamic Source Routing* (DSR)⁵ routing protocol [7].

For the cellular network model, we use the IEEE 802.11 MAC protocol in the *Point Coordination Function* (PCF) mode. Since we assume backlogged sources for all our simulations, use of the PCF leads to better performance for the cellular network model (due to the collision-free functioning of the PCF-based MAC protocol). Since *ns2* currently does not support the PCF mode, we have extended the *ns2* implementation of IEEE 802.11 to include the PCF mode of operation. In this mode, the access point or the base station polls the mobile hosts using the IEEE 802.11 polling frames in a round robin fashion. The polled mobile host then immediately transmits its data while all the hosts remain silent. Upon successful receipt of the data, the access point relays it to the destination and then continues polling. The physical layer implementations for the cellular model is the same as the one used for the ad-hoc network model.

3.2 Topology and Traffic Generation

For all of the results in this paper, we use a 1500m x 1500m grid in which four different network sizes of 50, 100, 200, and 400 uniformly distributed nodes are deployed respectively. The movement generator uses the waypoint mobility model [1] to create mobility in the scenarios. The waypoint model consists of two parameters: *speed* and *pause*. For each node, a random destination is picked in the grid, and the node is made to move towards the destination based on a uniformly distributed variable in the range $[0, speed]$. Once the node reaches the destination, it remains static at that position for *pause* amount of time, after which the whole cycle repeats. We use different speeds of 5, 10, 15, and 20 m/s respectively in the simulations. While we have performed simulations with different values for pause times, we

present only those for a pause time of 0 second as they represent the worst-case scenario in terms of mobility.

Every node acts as the traffic source to a randomly chosen destination in the network. We use *Constant Bit Rate* (CBR) traffic source with the packet size of 1 Kbytes for all flows in our simulations. TCP is used to transport traffic based on two reasons: (i) In all our comparisons, more focus is given to end-to-end performance rather than hop-by-hop performance. Since end-to-end performance is influenced by the bottleneck link, by using TCP the source adapts itself to the characteristics of the bottleneck link, thereby not using up available resources in the upstream of the bottleneck link. (ii) TCP accounts for about 95% of the current Internet traffic volume [2]. Depending on the network topologies, we choose various data rates for all flows to study the behavior of the two network models under varying network loads (light to heavy) and network sizes (sparse to dense). However, due to lack of space, we present only the results for the moderately loaded scenarios in network topologies of 100 and 400 nodes. For a complete set of results, please see [4].

3.3 Algorithms

- *Minimum Transmission Power*: We use a simple probing mechanism to determine the minimum power required for any given topology. Since all nodes in the given network are made to transmit at the same power, the algorithm starts off at the minimum possible power level and increases the power at increments of ϵ . We use a value of 5mW for ϵ in our simulations. Before each increment, the network is checked for partitions. On the detection of a partition in the network, the increment is added and the process continues. If the network is connected, the algorithm terminates and returns the transmission power level at that instant.
- *Medium Access Protocols*: We use three different scheduling protocols in our simulations: (i) *IEEE 802.11*: This is the default medium access control protocol used by the *ns2* network simulator for multi-hop wireless networks. Most commercial wireless networks make use of this standard MAC protocol. However, recent studies have shown that the IEEE 802.11 protocol can exhibit both short-term and long-term unfairness to flows [8]. Hence, to preclude the performance of the network from being affected by a sub-optimal MAC protocol, we have also implemented and used two ideal scheduling protocols that are described below. (ii) *Ideal Link Scheduling*: This ideal protocol provides perfect MAC fairness at the node level. Specifically, it tries to be as fair as possible when giving nodes access to the shared channel, and given a set of fair schedules, chooses the one that will maximize throughput [10]. (iii) *Ideal Flow Scheduling*: While the previous protocol provide node level fairness, it might not be the optimal algorithm to use because nodes serving more number of flows will still receive the same share of the underlying channel as those with fewer flows. Consequently all flows traversing such nodes would experience significantly lower bandwidth. Ideal flow scheduling solves this problem by optimally scheduling flow transmissions rather than node transmissions. The differences in the performance of the two ideal scheduling algorithms will be demonstrated in the next section.

⁵ In a related work we show that the end-to-end throughput is governed more by the underlying scheduling discipline than the routing protocol per se [5]. Hence we do not present results for any idealized routing protocol in this paper.

3.4 Metrics

We use the following metrics to compare the performance of the two network models. For each of the metrics, we briefly provide the motivation for using it to compare the two network models.


- *Raw Network Capacity*: This is the upper bound for the maximum instantaneous traffic in the network. It is essentially the product of the maximum number of simultaneous transmissions possible in the network and the capacity (data rate) of the channel. Note that this is a constant for a given network topology.
- *Fair Network Capacity*: While the raw network capacity is determined by using a MAC protocol that aims to only maximize network throughput, a more practical MAC protocol would be one that provides fair channel access to all nodes in the network. Fair network capacity is the throughput of the network when using such a MAC protocol. This can be thought of as the total network throughput when all sources have destinations in their immediate neighborhood. Both the raw and fair network capacities are measured to give a sense as to what the upper bounds for the performance of ad-hoc networks are. Note that in the case of the cellular network model, both the raw and fair network capacity would be the same and will be equal to the channel capacity.
- *End-to-End Throughput*: The end-to-end throughput is measured as the average of the end-to-end throughputs of all flows in the network. This is an important metric that reflects the real capacity of the network as far as end-to-end connections are concerned.
- *End-to-End Delay*: This is the average end-to-end delay experienced by flows in the network. Since flows traverse multiple hops in an ad-hoc network, this metric gains importance as the delay is directly proportional to the hop length of the path and inversely proportional to the end-to-end throughput.
- *Power Consumption*: This metric keeps track of the average total power consumed per node. In an environment where the battery power of mobile devices is a scarce resource, the significance of this metric is evident. Besides, this metric is important in our comparison because just like in the case of delay, it would be interesting to see the cumulative effect of longer path lengths and shorter transmission ranges in ad-hoc networks on the average power consumption. In our simulations, the power consumption per packet transmission and reception are determined from the physical layer specifications.
- *Fairness*: For each of the above metrics, we measure the per-node (per-flow) variance to monitor the fairness properties of the network model.
- *Impact of Mobility*: Under this measure, we monitor two metrics: the probability of network partitions and route failures, and the associated degradation in end-to-end throughput. Since the focus of this paper is only a single cell scenario, the cellular model will not be affected by mobility. However, because of short transmission ranges of ad-hoc networks and the random mobility patterns, the network will be vulnerable to network partitions and route failures.

4. SIMULATION RESULTS


In this section, we present the simulation results for the different metrics introduced in the previous section⁶. In order to study the upper bound for the throughput in ad-hoc networks, we first measure the network capacity using an unfair MAC and a fair MAC. We then show for two different sizes of network topologies: (i) mean and variance of the end-to-end throughput, (ii) mean and variance of the end-to-end delay, and (iii) mean and variance of power consumption. Finally, we study the impact of mobility on the performance of the two network models. We present the percentages of network partitions and path re-routing that occurred during the course of a simulation. We also show the degradation in end-to-end throughput for different mobility rates. Unless explicitly specified, all simulation results are obtained from averages of 10 random samples, each with a running period of 60 seconds.

4.1 Network Capacity

Figure 1(a) shows the maximum network throughput possible for




(a) Unfair MAC




(b) Fair MAC

Figure 1: Network Capacity

⁶ Due to lack of space we do not present an analytical treatment to the results presented in this section. However, a detailed analytical interpretation can be found in [4].



(a) Mean




(b) Variance


Figure 2: End-to-End Throughput (100 Nodes)

the specific topologies used. The throughput is measured as follows. Every node in the network chooses a destination in its immediate neighborhood so as to maximize the possible spatial reuse (the neighbor with the minimum two-hop degree is chosen as the destination). The MAC protocol then approximates⁷ the maximum independent set such that every node in the set can transmit without interfering with each other's transmissions. Once the set is computed, only nodes in the set are allowed to transmit for the entire duration of the simulation. The sum of the throughput of all the nodes is plotted. The throughput is computed for different power levels. The *MIN* power level represents the lowest power level possible such that the network remains connected. It is computed using the algorithm specified in the previous section. The absolute values for the *MIN* power level are different for the different topologies and are not shown in the figure. As the figure shows, when the transmission power levels increases, the amount of spatial reuse decreases, leading to lower network capacity. Figure 1(b) shows the network capacity when using a fair MAC. While the unfair MAC attempts to maximize the network throughput without consideration for fairness, the fair

⁷ Computing the maximum independent set is a NP-hard problem.



(a) Mean




(b) Variance

Figure 3: End-to-End Throughput (400 Nodes)


MAC attempts to provide fair service to the nodes in the network, and given a set of fair allocation, tries to maximize throughput. Since the fair MAC gives precedence to fairness over throughput, it can be observed that the absolute values for the throughput have decreased.

4.2 End-to-End Throughput


In this section, we present the mean and variance of the end-to-end throughput for the different network sizes. As mentioned earlier, we present results only for the moderately loaded scenarios. It is evident from the results that although the raw network capacity increased multi-fold, the average end-to-end throughput per flow does not witness an equivalently significant increase. In spite of this, the end-to-end throughput for the ad-hoc network is higher than that in the cellular network for all power levels. The different curves shown in Figure 2(a) are for the cellular network and the three scheduling algorithms for the ad-hoc network introduced in the previous section. As expected, the ideal flow scheduling algorithm outperforms both the ideal link scheduling algorithm and IEEE 802.11. Moreover, since the scheduling discipline focuses on flow level scheduling, the variance of its end-to-end throughput remains the lowest among the three. However, its variance is still higher than that of the




(a) Mean



(a) Mean



(b) Variance



(b) Variance

Figure 4: End-to-End Delay (100 Nodes)

Figure 5: End-to-End Delay (400 Nodes)


cellular model because of the inherent limitations imposed by the network topology. (Consider two sets of flows, one in which all flows have sources and destinations on either side of a cut vertex cv of the underlying graph, and the other in which flows do not have overlapping paths. The flows in the first set have to contend with each other at cv and hence will receive significantly lower throughput than the flows in the second set.)

4.3 End-to-End Delay


In this section, we present and study the end-to-end delay performance of the two network models. As before, we present the simulation results for different network sizes, moderate loads, and for the different scheduling algorithms. Although the average path length of ad-hoc networks goes up with the increasing network size and decreasing transmission power, the corresponding increase in throughput due to the increased spatial reuse overcompensates for the increase in path length, thus leading to better end-to-end delay performance in ad-hoc networks. In Figure 4(a), IEEE 802.11 appears to perform better than the ideal scheduling algorithm in terms of delay because the load supported by IEEE 802.11 (as shown in Figure 2(a)) is much less than the load supported by the ideal link scheduling algorithm.

4.4 Power

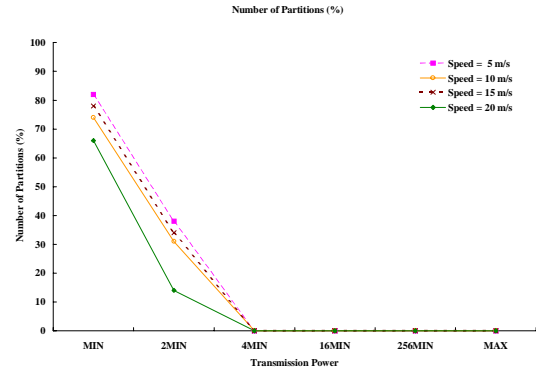
In this section, we present the mean and variance of the per-node power consumption in the two network models. As shown in Figure 6(a), the average power consumption per node decreases as the number of nodes in the network increases. This is notwithstanding the fact that the average path length also increases with increasing number of nodes resulting in more number of transmissions per packet. The reasoning behind this observation is tied to the power attenuation function used in our simulations. For example, consider the topology with 100 nodes. In the cellular mode, the average path length is 2 hops. However, the average per-hop distance is about 650m. Hence the transmission power consumption for an end-to-end transmission of a packet is proportional to $1 * 650^d$ (We do not take into account the base station's power consumption). In the ad-hoc mode, the average path length for the topology increases to about 6 hops. However, the per-hop distance reduces to around 210m. Hence the total transmission power used for an end-to-end transmission is proportional to $6 * 210^d$. Simulation results show the ad-hoc mode consumes power that is reduced by a factor of around 18 compared to the cellular mode.




(a) Mean




(b) Variance



(a) Network Partitions



(b) Path Re-routes



(c) Total Failures

Figure 6: Power Consumption

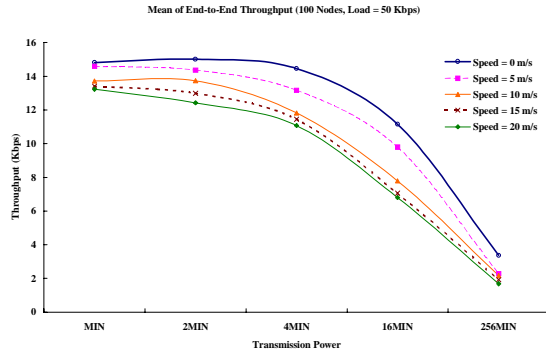
4.5 Impact of Mobility

In this section, we present results to demonstrate the impact of mobility on the performance of ad-hoc networks. The cellular network model is not affected by mobility since we do not consider nodes roaming from one cell to another, and all mobility is restricted to within one cell. We use two metrics to convey the impact of mobility. The first metric used is a measure of the frequency of network partitions with varying mobility rates. Specifically, we run the simulation for a period of 100 seconds, and check the network status every 1 second. Figure 7(a) shows the percentage of the samples in which the network is partitioned into more than one component. Although simulations are run for different network sizes, here we present only the results for the 100-node topology for want of space. As evident, with smaller power levels, the network is more vulnerable to partitions. On the other hand, as seen in earlier sections, the performance of the network in terms of throughput, delay and power improves with decreasing power levels. Hence, the challenge is to leverage the performance improvement exhibited by the ad-hoc network model while addressing the issue of connectivity.

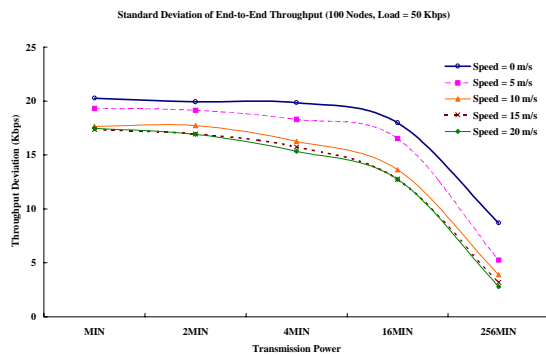
In Figure 7(b), we present the fraction of the samples for which newer paths have to be computed for flows as a result of mobility.

Figure 7: Impact of Mobility

While re-routing is not as catastrophic a phenomenon as partitioning, it serves as a drain on the network resources as evident in Figure 8, which shows the mean and variance of end-to-end throughput with increasing mobility rates. As expected, the performance of the network suffers with increasing speeds as packets are dropped due to mobility and due to the route re-discovery overhead. A detailed study of the degradation in



(a) Mean



(b) Variance

Figure 8: Mobility Induced Throughput Degradation

network performance in terms of throughput and delay with mobility is presented in [1]. The fraction of re-routes for the minimum power scenarios is less because of the network being predominantly partitioned. Figure 7(c) thus shows the sum of the failures due to partitions and path re-routes.

5. HYBRID NETWORK MODEL

In this section, we present a simple hybrid wireless network model that combines the advantages of ad-hoc networks (better throughput, delay, and power characteristics) with that of cellular networks (guaranteed connectivity)⁸. While the key focus of this paper is to identify the performance trade-offs between cellular and ad-hoc networks, we use the insights gained by studying the performance trade-offs to construct the hybrid network model. In the rest of the section, we first describe the new network model, and then provide simulation results comparing its performance with that of the conventional network models.

⁸ However, the hybrid model still does not address the issues of unfairness and impact of locality. In a later work, we have addressed these issues [6].

5.1 Architecture

The hybrid wireless network model consists of the same physical infrastructure as a conventional cellular network with an access point serving a certain coverage area. However, the mobile hosts operate by default in a multi-hop mode. The access point communicates with the mobile hosts through a separate control channel⁹. In the downstream, the control channel is used for the base station to convey two pieces of information to the mobile hosts: (i) the medium access mode to operate in, and (ii) the transmission power level. In the upstream, the mobile hosts use the control channel to periodically send to the base station: (i) the location of the mobile host (e.g., GPS location), and (ii) the observed throughput to offered load ratio. Periodically, the base station checks if there is a partition in the network by using the physical locations of the mobile hosts and their current transmission levels. In the event that the network has partitioned, the base station computes a new and increased power level using the *CONNECT_POWER* algorithm specified in Figure 9. The base station then conveys to the mobile hosts the newly computed power, and the mobile hosts adjust their transmission power levels accordingly.

Input:

Multi-hop wireless network $M = (N, L)$

Output:

Minimum power level p for network connectivity

Algorithm:

```

sort node pairs in non-decreasing order of mutual distance
initialize  $|N|$  components, one per node
for each  $(u, v)$  in sorted order, do
    if component( $u$ )  $\neq$  component( $v$ )
         $p = \text{power\_for\_distance}(\text{distance}(u, v))$ 
        merge component( $u$ ) and component( $v$ )
    if number of components is 1, then END

```

Figure 9: Algorithm *CONNECT_POWER*

The *CONNECT_POWER* algorithm is an adaptation of the *CONNECT* algorithm presented in [14], where the authors propose both centralized and distributed algorithms to control topologies of ad-hoc networks. The algorithm presented here is different since it assumes all mobile hosts to use the same power level. This is a conscious design decision rather than a simplification, because using different power levels at different nodes requires very sophisticated MAC protocols [11], and such protocols are still in the early stages of research. Furthermore, the work presented in [14] does not guarantee connectivity when operating in the distributed mode and hence does not address the problem that the hybrid model solves. On the other hand, the centralized algorithm presented in [14] cannot be used as is in a purely ad-hoc network because of the distributed nature of the

⁹ Note that the additional channel used in the hybrid model is entirely for control information unlike in iCAR, and hence uses a very small fraction of the underlying channel capacity.

network. In the hybrid model, this is not an issue since the base station runs the centralized algorithm.

If the network is not partitioned, then the base station checks to see if the total end-to-end throughput of the network is greater than that of the same network if it were to operate in a cellular mode. The throughput in the cellular model is computed as $C/2n$, where C is the channel capacity and n is the number of nodes in the network. If so, the base station directs the mobile hosts through the downstream control channel to change to the cellular mode of operation. When at a later point in time, the base station finds that the network topology is such that reverting back to an ad-hoc mode would increase total end-to-end throughput, it directs the mobile hosts to switch back to a multi-hop mode.

5.2 IEEE 802.11 and Hybrid Model

Switching from the cellular mode to an ad-hoc mode and vice versa when using IEEE 802.11 is simple since the MAC standard supports the coexistence of the DCF and PCF coordination functions. When the base station wants the network to operate in a cellular mode, it sends a downstream frame with the polling bit set and the mobile hosts immediately begin to operate in a cellular mode (PCF) until the base station explicitly sends a CF_END (end contention free period) frame, when they start contending again for the channel (DCF). However, the responsibility of switching from the MAC protocol's infrastructure mode (cellular) to the independent mode (ad-hoc) and vice versa still needs to be handled at the mobile host.

5.3 Simulation Results

In this section, we present simulation results that demonstrate the performance of the hybrid network model. The topology consists of a 1500m x 1500m grid with 100 nodes moving around randomly using the waypoint mobility model. Maximum speed for the mobility is set to 20 m/s, and different loads ranging from 20 Kbps to 1 Mbps are used for this simulation. Figure 10 shows the mean end-to-end throughput experienced by flows in the hybrid model against that in a non-adaptive ad-hoc network. It can be observed that as the load increases, the separation between the two curves widens. This is because of the fact that the same mobility scenario being used across all loads. Thus for the same periods of time when the network is partitioned, the non-adaptive model would lose more packets in a heavily loaded scenario. The absence of any partitions in the hybrid model results in all flows enjoying end-to-end connectivity throughout the simulation period. Therefore, although instantaneous throughput of the hybrid model goes down during partitions (due to increase of transmission power level), it stays much higher than that of the non-adaptive ad-hoc model, resulting in higher cumulative throughput.

6. SUMMARY

6.1 Issues

The following are some of the issues not addressed in the paper: (i) The scope of the paper is limited to the performance of the two network models within a single cell. Hence, the results presented are applicable only to standalone wireless networks consisting of a single cell, or local-area and wide-area wireless networks in the Internet where most of the traffic is within a single cell. Ongoing work is looking into what kind of performance the two network




Figure 10: Performance of Hybrid Wireless Network Model

models would exhibit when a significant number of the mobile hosts within a cell communicate with hosts outside the cell, and hence use the base station as the destination within the cell. We plan to present the results of this work in the near future. (ii) The metrics used to compare the two network models in this paper are chosen purely from the perspective of network performance. However, there are several other issues that may play a determining role in the choice of the network model for future wireless communications systems. These include security (authentication, integrity, and confidentiality), pricing, and billing. While we acknowledge that the above-mentioned issues are important when moving from a cellular model to a hybrid network model, we believe that having connectivity with the backbone at all times (unlike in pure ad-hoc networks) will enable these issues to be addressed more easily. (iii) The simple hybrid network model proposed in this paper does not address the issue of unfairness in ad-hoc networks. In a later work, we have proposed an updated hybrid network model that addresses this issue [6].

6.2 Conclusions

In this paper we evaluate the performance trade-offs between conventional cellular networks and multi-hop ad-hoc networks in terms of the raw network capacity, end-to-end throughput, end-to-end delay, power consumption, fairness, and impact of mobility on network performance. We conclude that for a single cell network, ad-hoc networks perform better than cellular networks in terms of throughput, delay, and power. However, they prove to be unfair to network hosts, show sharp degradation of performance in the event of mobility and can suffer from potential network partitions. We argue that the latter two characteristics prevent them from being adopted as the solution for future wireless communications systems.

Additionally, we propose a simple hybrid wireless network model that is based on a cellular infrastructure but operates in either the ad-hoc mode or the cellular mode depending on which model provides better performance for the state of the network at any given time. The base station of the cellular infrastructure plays the controlling role in deciding the specific mode that the network should operate in. We show through simulation results that the performance of the hybrid model combines the advantages of the

two network models. We argue that given the state of current MAC standards like IEEE 802.11, the overhead involved in switching from one mode to another is negligible and the switching itself requires no changes at the medium access control layer.

7. REFERENCES

- [1] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu and J. Jetcheva. A Performance Comparison of Multi-hop Wireless Ad Hoc Network Routing Protocols. In *Proceedings of ACM/IEEE MOBICOM*, Dallas, TX USA, October 1998.
- [2] CAIDA. <http://www.caida.org/analysis/performance/>
- [3] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388-404, March 2000.
- [4] H.-Y. Hsieh and R. Sivakumar. Performance Comparison of Cellular and Multi-hop Wireless Networks. GNAN Research Group Technical Report 00-01, November 2000.
- [5] H.-Y. Hsieh and R. Sivakumar. Improving Fairness and Throughput in Multi-hop Wireless Networks. In *Proceedings of IEE/IEEE ICN*, Colmar, France, July 2001.
- [6] H.-Y. Hsieh and R. Sivakumar. Sphinx: A Hybrid Network Model for Wireless Packet Data Networks. GNAN Research Group Technical Report 01-01, January 2001.
- [7] D. B. Johnson and D. A. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. *Mobile Computing*, Chapter 5, pp. 153-181, Kluwer Academic Publishers, 1996.
- [8] C. E. Koksal, H. I. Kassab and H. Balakrishnan. An Analysis of Short-Term Fairness in Wireless Media Access Protocols. In *Proceedings of ACM SIGMETRICS*, Santa Clara, CA USA, June 2000.
- [9] Y.-D. Lin and Y.-C. Hsu. Multihop Cellular: A New Architecture for Wireless Communications. In *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.
- [10] H. Luo, S. Lu and V. Bharghavan. A New Model for Packet Scheduling in Multihop Wireless Networks. In *Proceedings of ACM/IEEE MOBICOM*, Boston, MA USA, August 2000.
- [11] J. P. Monks, V. Bharghavan, and W. W. Hwu. Transmission Power Control for Multiple Access Wireless Packet Networks. In *Proceedings of IEEE LCN*, Tampa, FL USA, November 2000.
- [12] The Network Simulator ns-2. <http://www.isi.edu/nsnam/ns/>
- [13] C. Qiao and H. Wu. iCAR: an Intelligent Cellular and Ad-hoc Relay System. In *Proceedings of IC3N*, Las Vegas, NV USA, October 2000.
- [14] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment. In *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.
- [15] S. S. Rappaport and L.-R. Hu. Microcellular Communication Systems with Hierarchical Macrocell Overlays: Traffic Performance Model and Analysis. *Proceedings of the IEEE*, vol. 82, no. 9, pp. 1383-1397, September 1994.
- [16] Nokia Rooftop. <http://www.rooftop.com>
- [17] L. Swasey. 3G Wireless: Global Data & Voice Networking Strategies. Market Research Report, Global Information, Inc.
- [18] Wireless Data Forum. <http://www.wirelessdata.org>