Towards a Hybrid Network Model for Wireless Packet Data Networks

Hung-Yun Hsieh and Raghupathy Sivakumar
The GNAN Research Group
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 ad-hoc peer-to-peer wireless networks. We compare through simulations the performance of the two network models in terms of throughput, delay, power consumption, per-flow fairness, impact of mobility, impact of traffic locality, and impact of node distribution on the network performance. The simulation results show that while peer-to-peer networks perform better in terms of throughput, delay, and power, they suffer from unfairness, and poor performance in the event of mobility and low traffic locality.

We discuss the trade-offs involved in the performance of the two network models and contend that the trade-offs preclude the adoption of either of the network models as a clear solution for future wireless packet data networks. Thus, we present a simple hybrid wireless network model that uses a peer-to-peer network model in tandem with a conventional cellular network model. It supports a dual mode of operation and has the combined advantages of cellular and peer-to-peer wireless networks without suffering from the disadvantages of either. We present simulation results showing that the hybrid network model outperforms the conventional cellular network model in terms of throughput, delay, and power consumption, and achieves better fairness and resilience to mobility than the peer-to-peer network model.

1. Introduction

In recent years, the proliferation of mobile devices like hand-held PCs and PDAs has resulted in the rapid evolution of wireless packet data networks. Most wireless packet data networks use a *cellular* network model consisting of a base station or an access point to which the mobile hosts talk to directly. The base station in turn is connected to the backbone Internet through a distribution network. If the source and the destination lie within the same base station's cell, the base station serves as a relay between the hosts. Otherwise, the base station serves as a gateway between the wireless network and the backbone wireline network.

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 [1, 2], 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 sustain the accelerated growth in the number of users has severely exposed the limitations imposed by such a network model. The inability of the cellular model to effectively leverage the spatial reuse possible in wireless environments, and its reliance on a central coordinating entity that becomes a communication bottleneck has prompted researchers to investigate alternate network models for future wireless packet data networks [3, 4, 5]. Such models typically attempt to increase the amount of spatial reuse and decrease the dependency on the base station.

In this context, a special class of networks called *ad-hoc* networks [6] has gained attention by virtue of their ability to operate using a *peer-to-peer* network model in which there is no need for a base station. The mobile hosts use short-range transmissions and communicate with each other over multi-hop paths consisting of other hosts in the network. Using transmission ranges that are just large enough to ensure network connectivity allows the peer-to-peer model to potentially maximize the spatial reuse in the network. Although these properties in turn make the peer-to-peer network model an attractive candidate to be considered as a solution for future wireless packet data networks, we believe that the performance trade-offs between the cellular and the peer-to-peer network models need to be well understood as they can provide valuable insight towards the development of alternate and better network models. Towards this end, we make two contributions in this paper:

1. We present a simulations-based comparison of the performance trade-offs between the cellular and peer-to-peer network models. We compare the models along different dimensions including: end-to-end throughput, end-to-end delay, power consumption, fairness, impact of mobility, impact of traffic locality, and impact of node distribution. We conclude that although peer-to-peer networks can perform better in terms of end-to-end

throughput, delay, and power consumption, they fare badly in terms of the fairness they offer to network flows, and are highly vulnerable to mobility and traffic locality.
2. We use the insight gained from the comparison results to propose a simple hybrid network model that offers the performance benefits of peer-to-peer networks while not exhibiting the associated problems. Essentially, *the proposed hybrid model's performance tracks the bounding envelope of the super-imposition of the performance curves of the cellular and peer-to-peer network models*.

The rest of this paper is organized as follows: In Section 2, we describe the simulation model including the simulation environment and the metrics used. In Section 3, we present the simulation results and interpretations. In Section 4, we describe the hybrid network model and simulation results. Finally, Section 5 concludes the paper.

2. Simulation Model

We use the *ns-2* [7] network simulator with CMU wireless extensions for all simulations presented in this paper. The rest of the section describes the simulation model and environment in details.

2.1. Network Models

We use the IEEE 802.11 MAC protocol in the *Distributed Coordination Function* (DCF) mode for the peer-to-peer network model. The channel propagation model consists of a combination of a free space propagation model and a two-ray ground reflection model. The crossover point of the two models is called the *reference distance*. When the receiver is within the reference distance of the transmitter, the free space model where the signal attenuates as $1/r^2$ is used. Outside of this distance, the ground reflection model where the signal falls off as $1/r^4$ is used. The physical layer follows IEEE 802.11's *Direct Sequence Spread Spectrum* (DSSS) specifications, and the data rate of the channel is set to 2Mbps. Unless otherwise specified, the transmission power is set to the minimum power required to keep the network connected. All the simulations for peer-to-peer networks use *Dynamic Source Routing* (DSR) [8] as the routing protocol.

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 simulations, using the PCF leads to better performance for the cellular network model (due to the contention/collision free functioning of the PCF based MAC protocol). In this mode, the access point (or the base station) polls the mobile hosts using the polling bit in the IEEE 802.11 frame header. The mobile host then immediately transmits its data while the other hosts remain silent. The access point does the polling in a round-robin fashion. The physical layer for the cellular model is the same as the one used for the peer-to-peer network model.

2.2. Topology and Traffic Generation

For all simulations in this paper, we use a grid of 1500m x 1500m with four different network sizes of 50, 100, 200, and 400 nodes respectively. The initial positions of the nodes are randomly chosen in the grid. Specifically, for the 100-node topologies used, we measure the average degree of a node in the network to be 7.3, the minimum and maximum degrees to be 1 and 14 respectively, and the standard deviation of the degree to be 3.0. The density of the network is thus non-uniform and hence does not cause bias towards one of the two network models.

The movement generator uses the waypoint mobility model [9] for creating mobility in the scenarios. The waypoint model consists of two parameters: *speed* and *pause time*. For each node, a random destination is picked in the grid, and the node moves 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 time* amount of time, after which the whole cycle repeats again. We use different speeds of 0, 1 (pedestrian), 10, and 20 (vehicular) meters/second respectively. While we have performed simulations with different values for pause times, we present only those for a pause time of 0 seconds as they represent the worst-case scenario in terms of mobility.

Each node in the network acts as a CBR traffic source to a randomly chosen destination. We use TCP as the transport layer protocol. The choice of using TCP is 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, thus not using up unused resources upstream of the bottleneck link. (ii) TCP accounts for about 95% of the current Internet traffic volume [10]. We use different data rates for CBR traffic in our simulations to study the behavior of the two network models under varying network loads. However, due to lack of space, we present the simulation results only for the moderately loaded scenarios (128Kbps, 64Kbps, 32Kbps and 16Kbps per-flow data rate for 50, 100, 200, and 400 nodes respectively). We use a packet size of 1KB for all simulations. All simulations are run for a period of 100 seconds. Each data point is an average over 10 different simulations run with different seeds for the random number generator.

2.3. Metrics

We compare the two network models using three sets of simulations as follow:

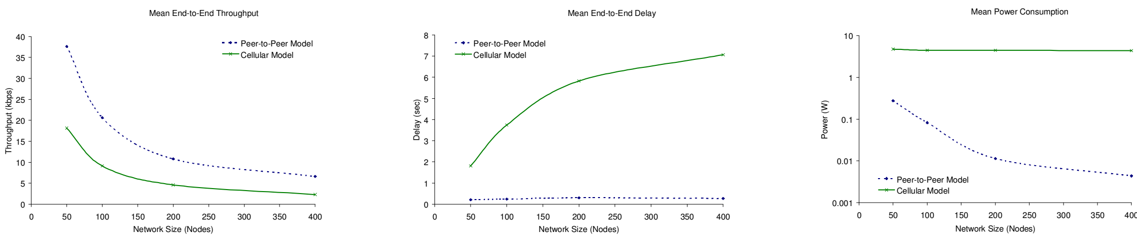

(a) Throughput (b) Delay (c) Power

Figure 1. Throughput, Delay, and Power (vs. Number of Nodes)

- *Throughput, Delay, and Power*: In the first set of simulations, we study the performance of the two network models in terms of the per-flow end-to-end throughput, end-to-end delay, and per-node power consumption. We study these metrics for different network sizes as described in Section 2.2. We use a constant load for each network size and use the minimum transmission power required to keep the network connected. All sources and destinations lie within the same cell and we do not consider mobility for this set of comparisons.
- *Impact of Transmission Power, Load, and Node Distribution*: In the second set of simulations we explore the impact of some of the parameters that we kept constant for the first set of simulations. Specifically, we study the performance of the two network models for different values of transmission power, different per-flow rates, and non-uniform node distributions. We present only throughput performance due to lack of space.
- *Fairness, Impact of Mobility, and Impact of Traffic Locality*: In the third set of simulations, we compare the two network models in terms of the fairness that they provide to network flows and the vulnerability they exhibit to mobility of nodes and locality of flows. We again use throughput as the metric to study the fairness, and impact of mobility and traffic locality. Note that the cellular network model can provide perfect fairness by virtue of the centralized coordination of the base station, and will not be affected by mobility since we consider only intra-cell mobility. However, this is not the case for the peer-to-peer model.

3. Simulation Results

3.1. Throughput, Delay, and Power

In this section we present simulation results comparing the cellular and peer-to-peer network models in terms of end-to-end throughput, end-to-end delay, and per-node power consumption. For each of the metrics, we show the performance comparisons for different network sizes of 50, 100, 200, and 400. Although most existing cellular networks serve only 30-50 users per cell, we expect future wireless networks to be forced to support more number of users per cell. Hence, we include much larger topologies in the comparisons to demonstrate that the trade-offs identified scale with the network size.

1. **Throughput**: Fig. 1(a) shows the average of the end-to-end throughputs of all flows for the two network models. The throughput is measured at the TCP sink after removing all control overheads introduced by the different protocol layers. Recall from the previous section that for each topology we make each of the nodes in the network a traffic source (with the destination picked randomly). The throughput for both network models goes down as the number of nodes (and hence flows) in the network increases. For the cellular model, the throughput will go down as $O(1/n)$ since the base station fairly arbitrates the available capacity among the n nodes in the network. On the other hand, for the peer-to-peer model, the spatial reuse and hence the raw capacity of the network will in fact increase with the network size. This is because as more nodes are distributed over the same grid area, the density of the network increases resulting in a smaller transmission range required to keep the network connected. However, this increase in network capacity is offset by the increase in the number of flows in the network. Since each flow traverses multiple hops, it can be shown that the end-to-end throughput decreases as $O(1/\sqrt{n})$ [11]. In the figure, although the throughput of both models decreases with the network size, the peer-to-peer model always exhibits a higher end-to-end throughput than the cellular model.
2. **Delay**: Fig. 1(b) shows the average end-to-end delay experienced by flows in the two network models as the network size increases. Since delay for a flow is directly proportional to the number of hops traversed by the flow, and inversely proportional to the flow's end-to-end throughput, this is an interesting metric to study in the

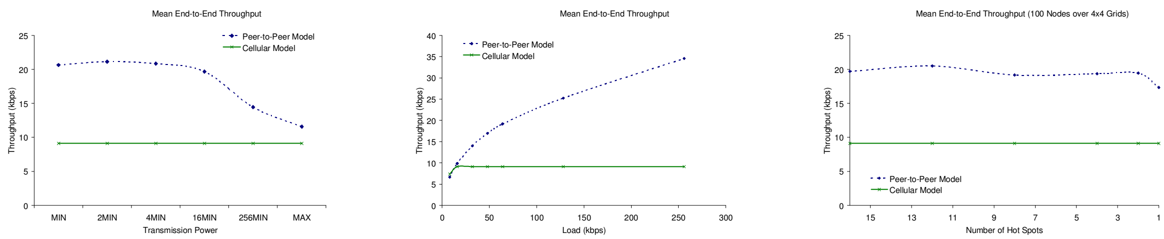

(a) Transmission Power (b) Load (c) Node Distribution

Figure 2. Impact of Transmission Power, Load, and Node Distribution

comparisons between the cellular and peer-to-peer network models. Although it can be shown that the difference in delays between the two network models is only of constant order[1], Fig. 1(b) shows a trend where the delay in the cellular model increases with the number of nodes while the delay in the peer-to-peer model remains the same. This is because of the fact that we measure the end-to-end delay from the time the transport protocol gives a packet to the routing layer to the time the packet reaches the destination. Since packets in the cellular model will remain in the link-layer queue until a transmission slot is given by the base station, the time packets spend in the link-layer queue is included in the delay shown. As the network size increases, this delay increases as nodes get access to the channel less often. In peer-to-peer networks, an increase in the network size has minimal effect on the delay since the communication is distributed and the delay solely depends on the end-to-end throughput.

3. **Power**: Fig. 1(c) shows the difference in the usage of power by mobile nodes in the two network models (not including the power used by the base station in the cellular model). We include all packet transmissions such as route discovery and maintenance packets in the peer-to-peer model. Note that in the power model we employ in the simulations, a packet transmission of distance r accounts for a usage of $O(r^4)$ of the battery power (for the 100-node topology the transmission power used in the peer-to-peer mode is approximately 200mW), while a packet reception accounts for a constant power consumption (for the 100-node topology it is about half the transmission power). Although flows traverse more number of hops of the order $O(\sqrt{n})$ in peer-to-peer networks, the transmission ranges decreases as $O(1/\sqrt{n})$. Hence, the amount of power consumed in the peer-to-peer model is less than that consumed in the cellular model by an order of $O(n\sqrt{n})$. This is substantiated by the results shown in Fig. 1(c).

3.2. Impact of Transmission Power, Load, and Node Distribution

In the last section we study the performance of the two network models for different values of the network size by using a constant transmission power and load for each network size, and a random node distribution. In this section we study the performance of the two network models for a 100-node network when the transmission power and load vary, and when the nodes are distributed in a non-uniform fashion onto the rectangular grid.

1. **Impact of Transmission Power**: Recall that in the last section we use the minimum transmission power required to keep the network connected as the transmission power. This results in maximizing the spatial reuse in the network. However, it might not always be possible to set the transmission power to the minimum possible value. In particular, using the minimum possible transmission power can make the network highly sensitive to node mobility and errors in the distributed transmission range estimation scheme. Hence, in Fig. 2(a) we study the impact of transmission power on the throughput performance of the two network models. As the figure illustrates, the performance gains of the peer-to-peer network model demonstrated in the last section holds good for small multiples of the minimum transmission power. Even at the largest possible transmission power (at which the network is fully connected), the peer-to-peer model provides better performance than the cellular model.
2. **Impact of Load**: While we have used a constant load for the performance comparisons in the last section, we show in Fig. 2(b) the impact of increasing load on the throughput experienced by flows for a specific network size (100 nodes). It can be seen that the throughput in the cellular model remains a constant after the initial rise since the

[1] An intuitive explanation for this observation is that in peer-to-peer networks the advantage in end-to-end throughput $O(\sqrt{n})$ is offset by the increase in the number of hops $O(\sqrt{n})$ a flow has to traverse.

maximum end-to-end throughput achievable in the cellular model is 10Kbps (capacity/flows/hops = 2Mbps/100/2). However, in the peer-to-peer model the average increases with the load for all the data points. In essence, the *knee* of the throughput vs. load curve for the cellular model is approximately 10Kbps, while the knee of the peer-to-peer model is significantly higher (for the 100-node topology, we observed the knee to be approximately 64Kbps) because of the increase in spatial reuse.

3. **Impact of Node Distribution**: While we have used randomly generated topologies for all the other simulations in this section, we show the effect of non-uniformly distributed topologies on the performance of the two network models. We take up throughput as the metric to demonstrate the impact. For the results shown in Fig. 2(c), we divide the 1500m x 1500m grid into a 4 x 4 array of same-sized smaller grids. We distribute 50 of the 100 nodes in the 1500m x 1500m grid randomly as before. However, for the other 50 nodes we randomly pick *x* of the 16 smaller grids and distribute the 50 nodes only among the *x* smaller grids. We vary the value of *x* from *1* (all 50 nodes within the same small grid) to *16* (same as the default case) and study the average throughput in the two network models (Fig. 3 shows the node distributions for *x=1* and *x=16*). As expected, the cellular model is not affected by the node distribution as we use a single sector cell and the distribution of nodes within the cell does not affect the performance experienced by flows. In the case of the peer-to-peer model, although the distribution impacts the throughput, the peer-to-peer model still has a higher throughput than the cellular model.

3.3. Fairness, Impact of Mobility, and Impact of Traffic Locality

We discuss in this section the remaining metrics for performance comparisons of the two network models.

1. **Fairness**: Fig. 4(a) shows the unfairness experienced by flows in the two models. We compute the unfairness index as the ratio of the standard deviation of the end-to-end throughput to the mean end-to-end throughput. As the results demonstrate, flows in the peer-to-peer network model can experience a high degree of unfairness. The primary reason for the unfairness is the underlying communication model that is based on a distributed and multi-hop environment. For example, if a flow A shares its path with four other flows while flow B does not share its path with any other flow, and given that all other parameters are the same, flow A will receive one fifth the throughput that flow B receives. Such unfairness is

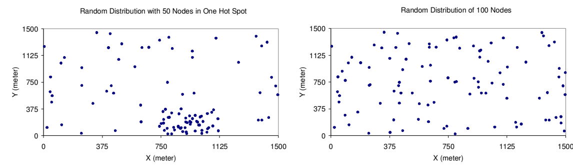

(a) One Hot Spot (b) No Hot Spot

Figure 3. Different Node Distributions

an inherent property of the underlying physical peer-to-peer network topology. While developing *fair* MAC and routing protocols can reduce the unfairness in peer-to-peer networks, we have shown in a related work that such protocols cannot entirely overcome the limitations of the underlying network model [12]. The cellular model on the other hand exhibits relatively much better fairness. While ideally the cellular model should not exhibit any unfairness at all, the unfairness exhibited in Fig. 4(a) is an artifact of the behavior of TCP, the transport protocol that we use. Specifically, since flows are served in a round-robin fashion, it is possible that flows that are served later in the cycle experience a larger round-trip time (including the delay experienced at the link-layer of the source) and hence lower throughput.

2. **Impact of Mobility**: Fig. 4(b) shows the average throughput for a network with 100 nodes as the speed at which nodes in the network move increases. Note that we do not study inter-cell mobility in this paper and hence the cellular network model is not affected by the mobility (thus the flat curve). However, the throughput of the peer-to-peer network model degrades with mobility because of two related but distinct phenomena: (i) *Network partitions*: The network can, for a period of time, remain partitioned into two or more components. If a flow has its source and destination in different components, all packets belonging to the flow are lost until the network reconfigures to a topology where the two components are connected. This can result in severe throughput degradation for the flow. (ii) *Route re-computations*: Even if the network remains connected, it is very likely that the route used by a flow will be broken because of the movement of one or more of the nodes along the path from the source to the destination. In such cases, all packets belonging to the flow are lost until an alternate route is computed, resulting in throughput degradation. It is interesting to note that in the presence of partitions the average throughput experienced by flows in the network might in fact be higher than the average without mobility. This is because flows within one component of the partitioned network manage to use up the capacity given up by the partitioned flows. In essence, the

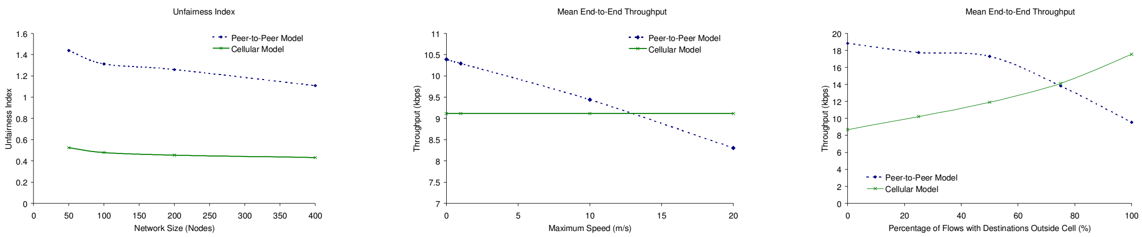

(a) Fairness (b) Mobility (c) Traffic Locality

Figure 4. Fairness, Impact of Mobility, and Impact of Traffic Locality

network is reduced to a state where only shorter flows thrive thus increasing the overall end-to-end throughput (recall that the end-to-end throughput is inversely proportional to the average hop count of flows in the network) at the expense of the partitioned flows (increasing unfairness). In our simulations the average throughput still decreases with increasing mobility due to the constant data rate (CBR traffic) source and the increasing overhead in route re-computations.

3. **Impact of Traffic Locality**: The peer-to-peer model is typically used in stand-alone networks where both the source and destination are within the same cell (or network). On the other hand, in a wireless network connected to the backbone Internet, there exist flows that have either the source or the destination outside the cell. We refer to such flows as *non-local* flows. All non-local flows will then use the base station as the destination even when functioning in the peer-to-peer mode. As the number of non-local flows increases, the contention in the local neighborhood of the base station will increase resulting in possible throughput degradation. Fig. 4(c) shows the throughput performance of the network models for different percentages of non-local flows in the network. As expected, the performance of the peer-to-peer model starts decreasing as the percentage of non-local flows increases, and falls down to much below the performance of the cellular network for a scenario where 100% of the flows are non-local. For the cellular model, non-local flows require only half the wireless bandwidth as local flows because they involve only one wireless transmission (source to base station or base station to destination) as opposed to the two wireless transmissions required by local flows. This explains the throughput increase experienced by flows in the cellular model as the percentage of non-local flows increases.

3.4. Summary of Results

Through our simulation results, we have demonstrated that while the peer-to-peer network model is better in terms of throughput, power, and delay (for varying network sizes, transmission ranges, loads and node distributions), it suffers from unfairness problems, and is highly vulnerable to network mobility and traffic locality.

4. Hybrid Network Model

In this section we present a simple hybrid network model that has the better throughput, delay, and power-consumption properties of the peer-to-peer network model, and at the same time does not exhibit the key drawbacks of peer-to-peer networks including unfairness, and vulnerability to mobility and traffic locality. We first present an overview of the hybrid network model and then show through simulation results the performance of the hybrid model. Due to lack of space, we do not present the details of the algorithms involved in the hybrid model.

4.1. Overview

The hybrid network model consists of a regular cellular infrastructure – with a base station – that supports a dual mode of operation. The channel is time-divisioned[2] into service periods with each service period consisting of n slots where n is the number of nodes (and hence flows) in the network. By default, the network operates in the peer-to-peer mode during all of the n slots.

Periodically, stations in the network convey to the base station the performance (throughput) observed by each of their flows during the last measurement period. If the throughput of a flow i is less than a threshold value (set to the expected throughput of the flow for the same period in a pure cellular network), the flow is assigned one of the n slots. If k flows are selected in this manner, during the next service period the network operates in the cellular mode for the first k of the n slots, with each of the selected flows served by the base station directly in the cellular mode. The network operates in the peer-to-peer mode for the remaining n-k slots. Flows served in the cellular mode do not receive any service in the peer-to-peer mode and

[2] Note that although we use time division to split the channel, the architecture does not stipulate a specific channel division scheme and other schemes like frequency division or code division can also be employed.

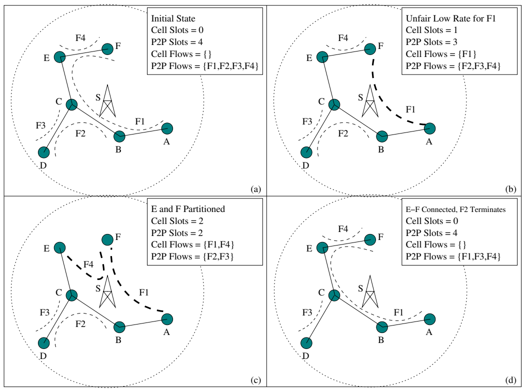

Figure 5. The Hybrid Network Model

vice-versa (to prevent out of order packet delivery from occurring and impacting TCP's performance). Irrespective of which flows are selected to be served in the cellular mode, all nodes in the network participate in packet forwarding during the peer-to-peer mode. Flows are periodically reverted back to the peer-to-peer mode, and their throughputs monitored as before.

Note that the performance of the cellular model is a lower bound for the performance of this hybrid network model. Irrespective of why flows experience throughput degradation (unfairness, mobility, or traffic locality) when served in the peer-to-peer mode, such flows will be switched to the cellular mode and thus would experience the same performance as they would in a cellular model. The worst-case performance of the hybrid model occurs when all the slots during a service period are allocated for the cellular model. In this case, the performance of the hybrid model will be the same as that of the cellular model.

Fig. 5 illustrates some of the aspects of the hybrid model. In Fig. 5(a), the system is in its initial state with four flows (F1-F4) in the peer-to-peer mode. The dotted circle represents the coverage area of the base station. In Fig. 5(b), flow F1 experiences low throughput because of the high contention along its path. Hence, F1 is switched over to the cellular mode and the channel is time-divisioned in the ratio of 3:1 between the peer-to-peer and cellular modes. In Fig. 5(c), nodes E and F are partitioned in the network. The throughput of flow F4 thus suffers, and consequently it is also switched to the cellular mode. Finally, in Fig. 5(d), nodes E and F are connected again, and flow F2 terminates, relieving the congestion on the multi-hop path between nodes A and F. Hence, both flows F1 and F4 are reverted back to the peer-to-peer mode.

4.2. Performance

We now show the performance of the hybrid model through simulations. The simulation model used is the same as the one described in Section 2. We use a network size of 100 with 50 TCP flows. We discuss the performance of the hybrid model in terms of throughput fairness, impact of mobility on throughput, and impact of traffic locality on throughput. Due to lack of space, we do not discuss the performance of the hybrid model in terms of the other metrics (throughput, delay, and power for static topologies, and impact of node distribution) since its performance will be the same as that of the peer-to-peer mode for the other metrics. (Recall that the peer-to-peer model performs significantly better than the cellular model for those metrics.)

- **Fairness**: In Fig. 6(a), we present the per-flow throughput distributions in the peer-to-peer model and the hybrid model for a data rate of 48Kbps. The figure also shows the average throughput of the cellular model. It can be seen that the hybrid model achieves a better distribution of throughput with none of the flows observing throughput much lower than the average throughput in the cellular model. This is because of the fact that all flows that experience significantly lower throughputs in the peer-to-peer mode will be switched to the cellular mode in the hybrid model. The marginal deviation experienced by some of the flow throughputs is because of the periodic switching of the flows back to the peer-to-peer mode (similar to the effect of bandwidth probing on TCP's performance). We are currently investigating compensation schemes to overcome this problem.
- **Impact of Mobility and Traffic Locality**: In Fig. 6(b), we compare the performance of the three network models in terms of the impact of mobility on the throughput. As discussed in Section 3.3, the throughput of the peer-to-peer model suffers for higher mobility rates while the throughput of the cellular model remains unaffected because the mobility is only intra-cell. The performance of the hybrid model tracks the performance of the peer-to-peer model during lower speeds. However, for the higher speeds, more flows experience throughput degradation and are switched to the cellular mode. Hence, the performance of the hybrid model tracks the performance of the cellular model for higher speeds and hence does not show the throughput degradation exhibited by the peer-to-peer model. In Fig. 6(c), we show the performance of the hybrid model for varying percentages of non-local flows in the network. As in the case of mobility, the performance of the hybrid model tracks that of the peer-to-peer model for lower percentages of non-local flows (when the peer-to-peer model is significantly better than the cellular model), and switches over to the cellular model for the extreme cases where most of the flows are non-local and hence the cellular model performs better than the peer-to-peer model.

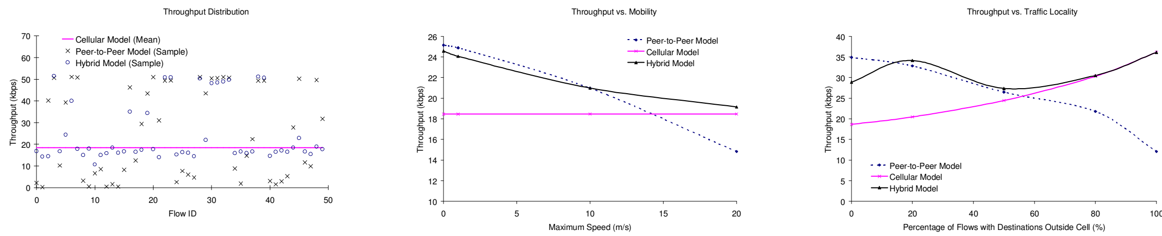

(a) Fairness (b) Mobility (c) Traffic Locality

Figure 6. Performance of the Hybrid Model

5. Issues and Conclusions

5.1. 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 stand-alone 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 three network models would exhibit in a multi-cells environment where inter-cell hand-offs are present. (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 might play a determining role in the choice of the network model for future wireless communication systems. These include the issues of security (e.g. 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 peer-to-peer networks) will enable these issues to be addressed more easily.

5.2. Conclusions

In this paper we evaluate the performance trade-offs between conventional cellular networks and peer-to-peer networks in terms of the throughput, delay, power consumption, fairness, and impact of mobility, traffic locality, and node distribution on network performance. We conclude that the two network models have significant trade-offs that preclude either of them from being adopted as the solution for future wireless packet data networks. We then propose a simple hybrid wireless network model that is based on a cellular infrastructure but operates in either the peer-to-peer mode or the cellular mode depending on which model provides better performance for the state of the network at any given time. We show through simulation results that the performance of the hybrid model closely follows the envelope of the graph obtained by superimposing the performance curves of the peer-to-peer and cellular network models, thus combining the advantages of the two network models.

References

[1] CTIA, "World of Wireless Communications," http://www.wirelessdata.org.

[2] L.-R. Hu and S. Rappaport, "Personal Communication Systems Using Multiple Hierarchical Cellular Overlays," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 2, pp. 406-415, Feb. 1995.

[3] 3G TR 25.924 V1.0.0, "*3GPP TSG-RAN; Opportunity Driven Multiple Access*," Dec. 1999.

[4] Y.-D. Lin and Y.-C. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications," in *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, Mar. 2000, pp. 1273-1282.

[5] C. Qiao and H. Wu, "iCAR: An Integrated Cellular and Ad-hoc Relay System," in *Proceedings of IEEE IC3N*, Las Vegas, NV USA, Oct. 2000, pp. 154-161.

[6] IETF MANET Working Group, "Mobile Ad-hoc Networks," http://www.ietf.org/html.charters/manet-charter.html.

[7] The Network Simulator *ns-2*, http://www.isi.edu/nsnam/ns.

[8] J. Broch, D. Johnson, and D. Maltz, "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks," Internet Draft *draft-ietf-manet-dsr-07.txt*, Feb. 2002.

[9] J. Broch, D. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," in *Proceedings of ACM MOBICOM*, Dallas, TX USA, Oct. 1998, pp. 85-97.

[10] CAIDA, "Analysis of Internet Traffic," http://www.caida.org/analysis/performance.

[11] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388-404, Mar. 2000.

[12] H.-Y. Hsieh and R. Sivakumar, "Improving Fairness and Throughput in Multi-hop Wireless Networks," in *Proceedings of IEEE ICN*, Colmar, France, July 2001, pp. 569-578.