

On Using the Ad-hoc Network Model in Cellular Packet Data Networks

Hung-Yun Hsieh and Raghupathy Sivakumar
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA
{hyhsieh, siva}@ece.gatech.edu

ABSTRACT

While several approaches have been proposed in literature for improving the performance of wireless packet data networks, a recent class of approaches has focused on improving the underlying wireless network model itself. Several of such approaches have shown that using peer-to-peer communication, a mode of communication used typically in ad-hoc wireless networks, can result in performance improvement in terms of both throughput and energy consumption. However, the true impact of using the ad-hoc network model in wireless packet data networks has neither been comprehensively studied, nor characterized. In this paper, we investigate the benefits of using an ad-hoc network model in cellular wireless packet data networks. We find that while the ad-hoc network model has significantly better spatial reuse characteristics, the improved spatial reuse does not translate into better throughput performance. Furthermore, although considerable improvement is seen in energy consumption performance, we observe that using the ad-hoc network model as-is might actually degrade the throughput performance of the network. We identify and discuss the reasons behind these observations. Finally, using the insights gained through our performance evaluations, we discuss strawman versions of three techniques which when used in tandem with the ad-hoc network model result in better throughput, energy consumption, fairness, and mobility-resilience characteristics. Through our simulation results, we motivate that using the ad-hoc network model in conventional wireless packet data networks is a promising approach when the network model is complemented with appropriate mechanisms.

1. INTRODUCTION

The mobile Internet user population has undergone a tremendous growth in the past few years. The growth has severely exposed the limitations of current wireless packet data networks in terms of the data rates that they can support. For example, existing 2.5G networks support bandwidths to the tune of a few tens of Kbps per user, while next generation 3G wireless networks support bandwidth of a mere 384Kbps per channel outdoors and 2Mbps per

channel indoors. Moreover, the above rates are valid for a limited number of users per cell (typically five users per cell [1]) and will decrease with increasing number of users. The limitations have in turn inspired a considerable body of research toward improving performance of wireless data networks. Achievements of such research include smarter radio transmission (e.g. adaptive array antennas) [17], better channel access schemes [30], more efficient scheduling schemes [19], faster and intelligent hand-offs [25], and transport protocols that are wireless-aware [4, 26].

Although the aforementioned approaches do improve the performance of wireless data networks, they are inherently limited by the network model that they operate on, namely the *cellular network model*. In the cellular network model, mobile-stations communicate directly with the base-station and do not interact in any manner with the other mobile-stations inside the same cell. Due to the *centralized* model used in cellular networks, for a given number of users n , the throughput per user can be shown to be of the order of $O(\frac{1}{n})$ [11, 14]. In addition, since mobile-stations communicate directly with the base-station, the average power consumption per transmission is of the order of $O(R^k)$ where R is the radius of the base-station cell, and k is the attenuation factor that usually varies between 2 and 6 depending on the signal propagation model [31]. Therefore, the only avenue for improving throughput or power consumption is to decrease the coverage area per base-station, thus reducing the number of users served and reducing the cell radius. While this approach has been adopted in schemes like hierarchical cellular networks [24], the drawback is the high infrastructure cost involved in deploying a large number of base-stations and the associated distribution networks.

Over the last few years, several approaches for an alternate network model have been proposed to improve the performance of cellular data networks [2, 18, 23, 27, 33]. An interesting and important commonality between such approaches has been the similarity between the proposed network model and the model used in a special class of wireless networks called *ad-hoc networks*. While ad-hoc networks have typically not been considered as an integral part of the mainstream Internet, a considerable amount of research has gone into protocol development for such networks due to the unique applications they have in other areas such as military environments, disaster relief operations, and sensor networks. Ad-hoc networks were conceived for environments that lack the services of an established backbone infrastructure, and hence the mobile-stations in an ad-hoc network act as routers or forwarders and communication is enabled through multi-hop routes.

Such a peer-to-peer mode of communication has distinct performance benefits as well over the conventional cellular model including better spatial reuse characteristics, lower energy consumption, extended coverage areas, etc. In conventional cellular networks,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBIHOC'02, June 9-11, 2002, EPFL Lausanne, Switzerland.
Copyright 2002 ACM 1-58113-501-7/02/0006 ...\$5.00.

base-stations communicate with mobile-stations directly, and the mobile-stations operate in a purely *peer agnostic* fashion. The fundamental element of the proposed alternative network models is thus the adoption of the peer-to-peer mode of communication in cellular packet data networks to improve performance.

In [2], the authors propose an approach called ad-hoc GSM (A-GSM), wherein mobile-stations participate in relaying to improve coverage and robustness against radio link failures. The approach uses received signal strength as the parameter in the decision process to switch from direct base-station communications to mobile-station relays. In opportunity driven multiple access (ODMA) [27], a scheme considered under the 3G endeavor, the high data rate coverage of the cell is increased at the boundaries by allowing mobile-stations inside the original high data rate coverage area to act as relays for mobile-stations outside. In [23], the authors propose an integrated cellular and ad-hoc relay (iCAR) approach wherein special mobile relays are placed between cells to relay traffic from an overloaded cell to a relatively under-loaded cell and therein achieve load balancing. In the mobile-assisted data forwarding (MADF) approach [33], a similar relaying scheme is proposed to reduce delay. In [18], the proposed model called multi-hop cellular networks (MCN) uses multi-hop relays for the mobile-stations to reach the base-station while reducing the transmission power of the mobile-stations and the base-station. MCN is shown to improve throughput performance when sources and destinations co-exist in a single wireless cell without mobility. In [14], a simple hybrid model is proposed wherein mobile-stations are used for multi-hop relays, and the base-station coordinates the network topology by directing the power adaptation at the mobile-stations. Although the purpose is to maximize spatial reuse and reduce network partitions, the model is evaluated only for the scenario in which all sources and destinations are co-located within the same cell.

Despite the fact that peer-to-peer communication is used in these proposed models, the true impact of using such a communication model is yet to be comprehensively understood. In this paper, we investigate the benefits of using peer-to-peer communication¹ in cellular packet data networks. We evaluate the use of peer-to-peer communication along a variety of performance metrics, and conclude that while the better spatial reuse characteristics of peer-to-peer communication results in better throughput per unit power, it does not translate into better performance in terms of throughput. *On the contrary, we demonstrate that using peer-to-peer communication on an as-is basis can result in degraded throughput.* We identify and discuss the reasons behind these observations. We then discuss three approaches that in tandem translate the spatial reuse benefits of peer-to-peer communication into better performance in terms of throughput, throughput per unit power, fair service allocation and mobility-resilience. Using strawman realizations of the three approaches, we show the performance enhancements that can be achieved and hence motivate further investigation along these lines.

Thus, the contributions of this paper are twofold:

- We show that using peer-to-peer communication on an as-is basis has its benefits in terms of throughput per unit power, but results in degradation of throughput. Moreover, we demonstrate that peer-to-peer communication exhibits unfair service allocation and is more vulnerable to mobility induced performance degradation.

¹In the rest of the paper we interchangeably use the phrases *ad-hoc network model* and *peer-to-peer communication model*, although in reality the ad-hoc network model encompasses more than just peer-to-peer communication.

- We discuss a set of three approaches that in tandem address the drawbacks that we identify for the peer-to-peer communication, and consequently leverage the spatial reuse advantages to provide better throughput, throughput per unit power, fairness, and resilience to mobility.

The rest of the paper is organized as follows: In Section 2 we discuss the evaluation model used for simulations presented in this paper. In Section 3 we evaluate the peer-to-peer communication model and characterize its benefits and drawbacks in the cellular packet data environment. In Section 4 we present a set of three approaches and show how these approaches improve the performance. Finally, in Section 5 we discuss some important issues with the proposed approaches and conclude the paper.

2. EVALUATION MODEL

In this section, we describe the evaluation model used and the assumptions made in the rest of the paper. We use the *ns2* network simulator [22] - along with the multi-hop wireless extensions from the CMU Monarch group - for all our simulations. Unless otherwise specified, the simulations are based on a 100 node topology. We focus only on a single base-station cell as shown in Figure 1. The implication of a multiple-cell environment is discussed in Section 5. We consider a one square mile area as the coverage area of the base-station. Each simulation is run for a 60 second period and each data-point in the presented results is averaged over 40 simulation runs with different random seeds. In the rest of the section we describe the other aspects of the evaluation model:




Figure 1: Communication Model

- *Physical Layer:* We use a single channel with a data rate of 2 Mbps. The propagation model used is a combination of the free space propagation model and the two-ray ground reflection propagation model. The signal degrades as $1/r^2$ within the cross-over distance (about 90m) and as $1/r^4$ above it.
- *Medium Access:* For the cellular network model, we use the *Point Coordination Function* (PCF) mode of the IEEE 802.11 MAC standard as the medium access control protocol. For the peer-to-peer model, we use the *Distributed Coordination Function* (DCF) mode of the IEEE 802.11 standard as the medium access control protocol.
- *Routing:* We use the dynamic source routing (DSR) protocol as the routing protocol for the peer-to-peer model [5].
- *Transport:* We use TCP as the transport protocol for all our simulations. Although performance evaluation efforts in the pure peer-to-peer (ad-hoc) networking research community have typically used UDP as the transport protocol [6], we use TCP because of the target environment that we consider in this paper. Specifically, 95% of the traffic in the Internet

consists of TCP flows [7], and it is reasonable to assume that traffic due to mobile Internet users will not deviate from this behavior.

- **Traffic Model:** Since the goal is to see how well the peer-to-peer model does in comparison to the cellular network model, we assume that the only mobile-stations in the cell are the ones that have data to send or receive. Specifically, unlike in other peer-to-peer evaluation efforts [6] that focus on stand-alone networks, we do not assume that the number of mobile-stations in the network is more than the number of flows in the network. Due to the scarce resources available at mobile-stations, we believe it is reasonable to assume that mobile-stations that do not want to remain connected to the backbone (having data to send or receive) will be unlikely to participate in the network operation. Therefore, in our simulations, *all the flows in the network have destinations outside the cell and hence traverse through the base-station*. Again, this reflects a realistic scenario since a dominant portion of the Internet traffic (>85%) consists of accesses to servers (http, ftp, smtp, etc.) [7]. Unless otherwise mentioned, we use a CBR application with 32Kbps data rate for each of the 100 flows.
- **Energy Model:** We use the default energy model provided by the *ns2* simulator to measure the power consumption. Specifically, for a 250m transmission range the transmit power used by *ns2* is 0.282W. In general the transmit power used for a transmission range of d is proportional to d^4 due to the signal propagation model described earlier. We use only the transmit power for our energy computations. In the peer-to-peer model, all packet transmissions including route computation and maintenance packets are incorporated. The transmit power consumption in the cellular model is based on the distance from the transmitting mobile-station to the base-station.
- **Topology:** We randomly distribute the mobile-stations on a $1500m \times 1500m$ grid. Although we have performed simulations with skewed distributions, such distributions do not have any significant impact on the nature of results that we show in this paper. Hence, due to lack of space we do not present results for other classes of node distributions.
- **Mobility Model:** We use the way-point mobility model for the scenarios with mobility. Nodes randomly pick a destination within the grid, and move toward the destination with a speed that is uniformly distributed between 0 and $speed_{max}$ meters per second. The $speed_{max}$ for a mobility scenario used is the value referred to in the x-axis of the mobility graphs. The pause time in the way-point model is set to 0. Since we consider only a single-cell network, the mobility of mobile-stations is restricted to within a cell. Consequently, the mobility model has no impact on the cellular network model.

3. USING THE AD-HOC NETWORK MODEL IN CELLULAR DATA NETWORKS

In this section, we evaluate the performance when using peer-to-peer communication in a cellular packet data network, and compare the performance against that of the conventional cellular network model. The metrics used for the comparison are: (i) Spatial reuse, (ii) Throughput, (iii) Power, (iv) Throughput per Unit

Power, (v) Fairness, and (vi) Impact of Mobility. While the per-flow throughput is an obvious metric to use in any network performance evaluation, we briefly provide the motivation for choosing the other parameters: (i) Spatial reuse: The key advantage of the peer-to-peer network model is the increase in spatial reuse that stems from its short-range transmissions [16]. By studying the number of simultaneous transmissions possible in the network during a single transmission slot, we characterize the degree of spatial reuse possible in the peer-to-peer model. Note that this metric will always be *one* for the cellular model. (ii) Power and Throughput per Unit Power: Since mobile-stations typically rely on a limited power source such as a battery, energy conservation has come to be realized as a critical factor to be considered in mobile networking. Hence, we study both the average power consumption, and throughput per unit power as metrics indicative of the energy conservation properties of the two network models. (iii) Fairness: While the cellular network model is an almost perfectly fair system by virtue of the centralized scheduling or channel-provisioning performed at the base-station, the same cannot be said for the peer-to-peer network model [29]. Due to the distributed nature of the network protocols, it is very possible that certain flows receive more service than others in the network. The goal of monitoring the fairness in terms of the normalized standard deviation of throughput is to characterize the fairness property of the peer-to-peer network model. (iv) Impact of Mobility: Since we consider only a single-cell network in this paper, the mobility of mobile-stations (within the cell) has no impact on the performance of the cellular network model. However, the peer-to-peer network model uses a shorter transmission range and hence will be more susceptible to failures due to mobility. We characterize the impact of mobility on the performance of the peer-to-peer network model by studying the throughput degradation due to mobility.

3.1 Spatial Reuse

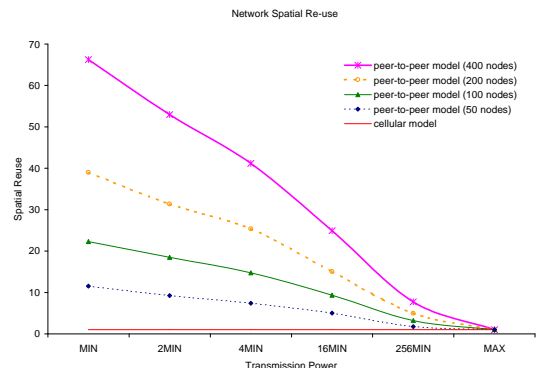


Figure 2: Spatial Reuse

In Figure 2, we present the degree of spatial reuse observed in the peer-to-peer network model as the transmission power of the mobile-station varies. For reference, we also present the spatial reuse possible in the cellular network model. It is interesting to note that at the minimum transmission range (using the minimum power such that the network stays connected), as the number of nodes in the network increases, the spatial reuse in the peer-to-peer

network also increases. This is because of the fact that the minimum transmission range is inversely proportional to the density in the network, and spatial reuse is inversely proportional to the transmission range. Hence, the spatial reuse increases with the increase in density of nodes within the cell. *This is a key property of peer-to-peer networks that promises an inherently better scalable network model than the cellular model where the capacity remains constant with increasing density (thus resulting in decreasing per-user throughput).*

3.2 Throughput




Figure 3: Throughput

In Figure 3, we show the end-to-end throughput (measured at the TCP sink) averaged over all flows for the two network models as the load (per-flow data rate) is varied. The interesting observation is that not only has the spatial reuse improvement shown in Section 3.1 not translated into better per-flow throughput, but the throughput in the peer-to-peer network model is in fact lower than that observed in the cellular network model. We now attempt to explain this observation in three parts:

1. **Multi-hop Routes:** Although the spatial reuse is increased, since a flow traverses multiple hops in the peer-to-peer network model, the end-to-end throughput of a flow, while directly proportional to the spatial reuse, is also inversely proportional to the hop-length. Moreover, since the expected hop-length in a dense network is of the order of $O(\sqrt{n})$, a tighter bound on the expected per-flow throughput is $O(\frac{1}{\sqrt{n}})$ [11]. While this bound is still higher than that of the cellular network model ($O(\frac{1}{n})$), the following two reasons degrade the performance even more.
2. **Base-Station Bottleneck:** The degree of spatial reuse and expected per-flow throughput of the peer-to-peer network model discussed thus far is valid for a network where all flows have destinations within the same cell. However, the scenario for the results presented in this section is one wherein the base-station is the destination for all flows (in the wireless component of their end-to-end path). Hence, any increase in spatial reuse possible cannot be fully realized as the channel around the base-station becomes a bottleneck and has to be shared by all the flows in the network. Note that this is not an artifact of the single-channel model adopted in our simulations. As

long as the resources around the base-station are to be shared by all the flows in the network (irrespective of the number of channels), the performance of the flows will be limited to that of the cellular network model.

3. **Protocol Inefficiencies:** The protocols used in the cellular network model are both simple and centralized (with the base-station performing most of the coordination) and operate over a single hop leading to very minimal performance degradation because of protocol inefficiencies. However, in the peer-to-peer network model, the protocols used are distributed (IEEE 802.11 and DSR), and operate over multiple hops. The inefficiencies that arise because of the distributed operation of the medium access and routing layers, and the multi-hop operation at the transport layer (the multi-hop path results in more variation in latency, losses, and throughput for TCP [12]) translate into a further degraded performance.

The combination of the above factors results in the spatial reuse improvement achieved in the peer-to-peer network model not translating into a corresponding increase in per-flow throughput.

3.3 Throughput Per Unit Power

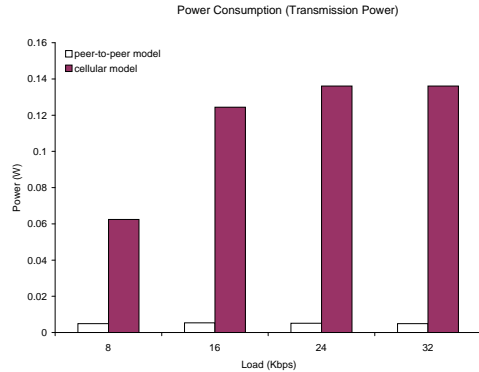


Figure 4: Power Consumption

Figures 4 and 5 present the average per-node power consumption and the average throughput per unit power for the two network models. Recall from Section 2 that the transmit power for a successful packet reception is $O(r^k)$, where r is the distance between the transmitter and the receiver and k is the attenuation factor used in the propagation model. In the case of the cellular model, the distance from a mobile-station to the base-station is always of the order of $O(R)$ irrespective of the number of mobile-stations within the cell, where R is the radius of the cell. In the case of the peer-to-peer network model, since a multi-hop route is used, the distance between any transmitter and receiver is significantly smaller (of the order of $O(\frac{R}{\sqrt{n}})$). However, the end-to-end power consumption is also directly proportional to the hop-length $O(\sqrt{n})$. Hence, while the average power consumption in the cellular network model is of the order of $O(R^k)$, in the peer-to-peer model it is of the order of $O(\sqrt{n} * (\frac{R}{\sqrt{n}})^k)$ resulting in a considerable improvement in power consumption by a factor of $O(\sqrt{n}^{k-1})$, where k is usually greater than 2. For the specific scenarios used for the presented results, k is equal to 4, and R is approximately 1000m (half of the diagonal of

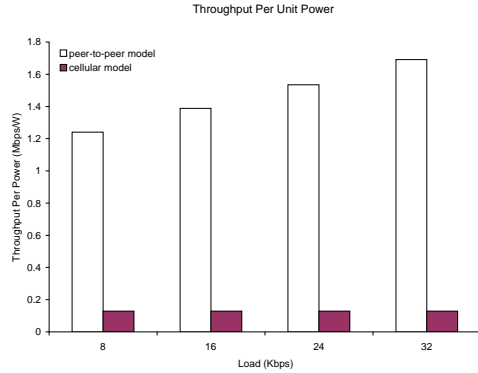


Figure 5: Throughput Per Unit Power

the square grid of size $1500m$), while the transmission range and hop-length in the case of the peer-to-peer model are approximately $250m$, and 3.5 respectively. This translates into roughly a twenty-fold decrease in the power consumption in the peer-to-peer model which is also seen in the results presented (although R is $1000m$, since the optimal transmission range is used in the cellular model, the expected transmission range is approximately $\frac{2 \times R}{3}$).

Although the throughput is smaller in the case of the peer-to-peer model, because of the significantly lowered power consumption, the throughput per unit power turns out to be considerably higher than that of the cellular model. We present the throughput per unit power results in Figure 5. For the results presented, the improvement in throughput per unit power increases by a factor of approximately 12 for the peer-to-peer model. This is an interesting result since it implies that, although the throughput is lowered, for applications that are insensitive to latency, the peer-to-peer model can be used to achieve significant savings in power consumption.

3.4 Throughput Fairness

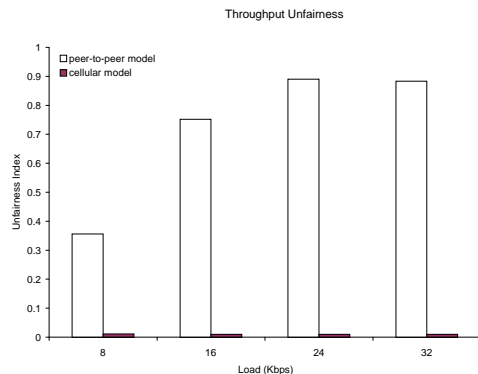


Figure 6: Fairness

In this section we present results demonstrating the fairness prop-

erties (or lack thereof) of the peer-to-peer network model. Note that because of the centralized round-robin like scheduling of the cellular network model, perfect fairness is achieved. However, in the peer-to-peer network model, because of the distributed protocols and because traffic can be distributed randomly, the throughput enjoyed by the different flows can be very different. We compute the unfairness index as the standard deviation normalized to the average throughput. Since the goal is to compare the performance of the peer-to-peer model against that of the cellular model, the standard deviation is computed using $\min(\frac{C}{f}, S_i)$ for each flow i , where C is the capacity of the channel, f is the number of flows, and S_i is the throughput of flow i . Equivalently, we consider that a flow is “unfairly” treated only if its throughput is lower than that it would have enjoyed in a cellular network model. Achieving throughput higher than that in a cellular network model is an incentive to participate in the peer-to-peer model, and hence we do not use the absolute throughput enjoyed by such flows in the fairness metric.

In Figure 6 we present the unfairness indices for the two network models. It can be seen that while the cellular model has a minimum unfairness index, the peer-to-peer network model suffers from considerably higher unfairness. The principal source of unfairness in the peer-to-peer network model is the protocol inefficiency discussed in Section 3.2. We substantiate this argument in the next section.

3.5 Impact of Mobility

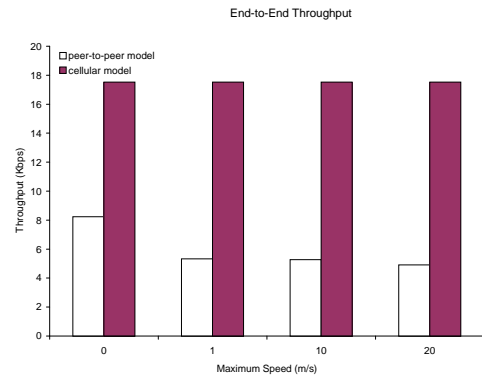


Figure 7: Impact of Mobility

While all the results presented thus far are for static scenarios, in this section we present results for scenarios in which the stations are mobile. The mobility model used is described in Section 2. Recall that because of the single-cell environment, the mobility of stations will not have any impact on the performance of the cellular model. While inter-cell mobility will have an impact on the cellular model’s performance, we discuss the implications of a multiple-cell environment in Section 5.

Due to lack of space, we present only the throughput results for the mobility scenarios. Figure 7 shows the throughput performance of the two models with increasing mobility. Note that the cellular model’s throughput is unaffected by the mobility. However, the performance of the peer-to-peer model degrades with increasing mobility due to the overheads caused by mobility induced route failures, route re-computations, associated losses, and the transport protocol’s reaction to such phenomena. In keeping with the focus

of the paper, we do not delve deeper into the above reasons and instead refer interested readers to related work that has investigated similar phenomena [3, 12].

It can be inferred from the results shown in Figure 7 that not only does the peer-to-peer network model exhibit lower throughput performance, it also shows degrading performance when the stations are mobile within the cell.

3.6 Summary of Results

In this section, we have shown that although the peer-to-peer network model enjoys better spatial reuse, it does not translate into better performance in terms of throughput. However, because of the short-range transmissions, the power consumption is greatly reduced and therefore the throughput per unit power is increased when compared to that of the cellular network model. Finally, we also show that the peer-to-peer network model exhibits poor fairness characteristics and exhibits further degradation in throughput with increasing mobility of the stations in the network. We summarize the key reasons for the drawbacks of the peer-to-peer model as follows:

- *Protocol Inefficiencies:* Since the peer-to-peer network model has traditionally been used only in stand-alone network environments, the network protocols used are typically distributed in nature with more focus on robustness and overall performance rather than on providing good per-flow service. As a result of the distributed operations, the protocols are inherently inefficient and in turn result in the peer-to-peer network model exhibiting poor throughput and unfairness in a cellular environment.
- *Base-Station Bottleneck:* Although the peer-to-peer network model uses short-range transmissions and hence increases the degree of spatial reuse in the network, since a majority of the flows in the cell are destined for destinations in the backbone Internet, the destination within the wireless cell for these flows happens to be the base-station. This results in the channel around the base-station becoming a bottleneck, limiting the throughput performance of the peer-to-peer network model to that of the cellular network model.
- *Impact of Mobility:* The peer-to-peer network model uses short-range transmissions and multi-hop routes. The impact of mobility is hence greater in the case of the peer-to-peer model than in the case of the cellular model. Specifically, when only intra-cell mobility is considered, the cellular model remains unaffected by the mobility of the stations in the network, whereas the performance of the peer-to-peer model decreases with increasing mobility.

4. ENHANCING PERFORMANCE OF THE AD-HOC NETWORK MODEL IN CELLULAR DATA NETWORKS

We now discuss a set of three approaches that address the drawbacks of the peer-to-peer network model identified in Section 3. We show that when the proposed approaches are used in tandem with a peer-to-peer network model, the spatial reuse benefits of the peer-to-peer model are translated into better performance in terms of throughput and power consumption. These approaches also provide fair service to flows and exhibit resilience to mobility.

4.1 Assisted Scheduling

While one option to solve the protocol inefficiency problem is to develop better distributed algorithms, another feasible option is to leverage the existence of the base-station and the availability of control channels that exist between the base-station and mobile-stations within the cell. Specifically, we consider an approach where the base-station plays an active role in the scheduling of flows in the network. Note that this is very similar to the role played by base-stations in a conventional cellular network [9]. The difference however lies in the fact that the scheduling is now done for multi-hop flows. The base-station periodically draws up a schedule for multi-hop transmissions within the network that maximizes throughput subject to fair service. The schedule is then broadcast to the mobile-stations through the control channel. The algorithm used by the base-station to perform the scheduling is shown in Figure 8.

```

Mobile-Station and Base-Station Communication
From each mobile-station  $n$ :
  Send neighbor list  $L(n)$  (or GPS information)
  If the source node of flow  $f$ , send flow backlog  $B(f)$ 

From base-station:
  Broadcast schedule of transmissions  $SCHEDULE$ 

At Base-Station
State Maintenance:
  Whenever  $L(n)$  is updated:
    Update the connection matrix  $CM$  of all mobile-stations based on  $L(n)$ 
    Obtain new route  $R(f)$  for every flow  $f$  using  $CM$ 

Scheduling:
  Initialize  $OVERFLOW$  to  $\emptyset$ 
  At time  $t$ , draw a new transmission schedule with  $slot$  time slots:
  Initialize  $SCHEDULE$  to  $OVERFLOW$ 
  Put each new backlogged flow  $f$  in  $LIST$ 
  While  $LIST$  is not empty:
    Find flow  $f$  with minimum service  $S(f)$ 
    status = Schedule-flow( $f, t$ )
    if status == SUCCESS
      decrement backlog counter  $B(f)$ 
      increment service counter  $S(f)$ 
      if  $B(f) == 0$ 
        remove  $f$  from  $LIST$ 
    else
      remove  $f$  from  $LIST$ 

Schedule-flow( $f, t$ )
   $sno \leftarrow t$ 
  Traverse each hop  $h$  of  $R(f)$  starting from the first hop
  Find the first slot  $s \geq sno$  such that transmission on hop  $h$ 
  will not interfere with transmissions already scheduled in slot  $s$ 
  if  $s < t + slot$ 
    schedule hop  $h$  in slot  $s$  by updating  $SCHEDULE$ 
  else if  $h$  is not the first hop
    schedule hop  $h$  in slot  $s$  by updating  $OVERFLOW$ 
  else
    return FAILURE
   $sno \leftarrow s + 1$ 
  return SUCCESS

```

Figure 8: Assisted Scheduling

Essentially, the base-station iterates through the list of flows with backlogged services, and for each flow schedules the hops along the path that the flow traverses. Once all the flows are accommodated within the schedule, the base-station iterates once again through the schedule and attempts to fill in more end-to-end transmissions for the flows within the schedule. The process is repeated until the schedule cannot be filled in with any flow. Unlike in related approaches [20], the base-station always tries to provide fair service before trying to enhance throughput. In other words, even when the “refilling” process is done, flows with less service are provided priority over flows with more service. Flows that have schedules overlapping the current schedule-period have slots reserved during the next schedule-period irrespective of the newly contending flows during the next scheduling operation (i.e. flows once scheduled are not preempted).

At each mobile-station, a single output queue is maintained for

all packets to be forwarded. When the MAC layer requests for a packet from a specific flow (according to the schedule drawn by the base-station), a *selective dequeue* mechanism is used to dequeue the first packet (from the head of the queue) that belongs to that flow. For results presented in this section, the base-station computes a shortest-path route for each flow. Note that more optimal routes (e.g. load-balanced routes) can be chosen by the base-station, although at the expense of more complexity.

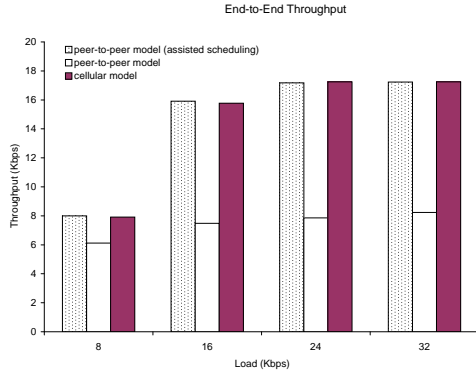


Figure 9: Throughput (Assisted Scheduling)

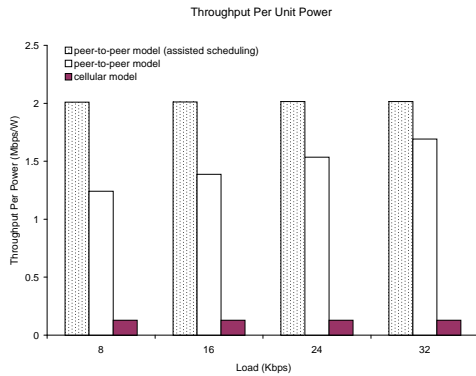


Figure 10: Throughput Per Unit Power (Assisted Scheduling)

In Figures 9, 10, and 11, we show the throughput, throughput per unit power, and fairness results for the assisted scheduling approach. It can be observed that when using assisted protocols the throughput now is the same as that of the cellular model (as opposed to worse in the case of the vanilla distributed protocols). The throughput per unit power results also show improved performance than in a cellular model, and the fairness achieved is almost the same as that in the cellular model.

The advantages of using assisted scheduling to improve the performance are threefold: (i) The base-station has global information and hence can draw up the optimal schedule for maximizing throughput and fairness. (ii) Since the schedule is constructed in a

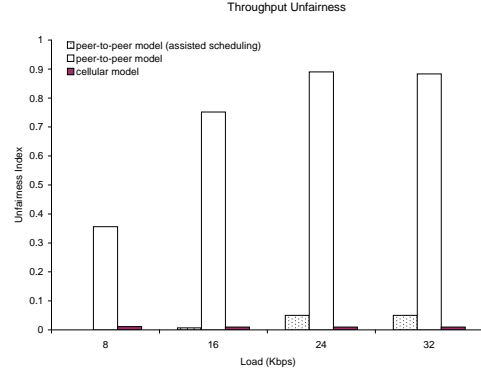


Figure 11: Fairness (Assisted Scheduling)

centralized fashion, the *band of contention* [13] of the medium access control protocol is reduced to one. This is in contrast to IEEE 802.11 where the band of contention is two (when a transmission occurs, stations in the vicinity of both the receiver and the transmitter cannot transmit or receive due to control packet exchanges between the transmitter and receiver - although in an ideal setting mobile-stations not in the vicinity of the receiver should be able to transmit, and mobile-stations not in the vicinity of the transmitter should be able to receive). (ii) The protocol inefficiencies because of the distributed operations are reduced.

However, there are several issues associated with the proposed approach: (i) *Communication Overheads*. The mobile-stations have to periodically inform the base-station about location information (either GPS or neighborhood information) in order for the base-station to construct the schedule. Also, the base-station has to periodically broadcast the schedule to the mobile-stations. However, in typical cellular networks, there already exist several control channels for exchanging information between the mobile-stations and the base-stations (including transmission schedules [9]). Hence, such control channels can be used for the information exchange. (ii) *Base-Station Complexity*. The base-station needs to compute the routes for all flows and maintain the transmission schedules. Since base-stations in the cellular environment typically perform complex operations, they can be provided with more computational power if necessary. (iii) *Mobile-Station Complexity*. The mobile-stations have to furnish the base-station with the location information in addition to information about new flows and terminating flows. While using GPS is an option, the location information can also be in the form of neighborhood information learned through receiving beacons or snooping on transmissions.

Note that the approach described above is primarily for demonstrating that protocol inefficiency can be reduced by using base-station assisted protocols, rather than to propose a specific mechanism for scheduling in such networks. In fact, in a different context, similar approaches have also been proposed. In multi-hop packet radio networks, (base-station) coordinated channel access schemes such as “spatial TDMA” (STDMA) [21] have been used to increase network capacity by drawing transmission schedules optimized to network topology. Scheduling algorithms that are adaptive to traffic dynamics [10] and immune to topology changes (due to mobility)

[8] have also been proposed². The results presented in this section motivate further consideration of such base-station assisted protocols when the peer-to-peer communication model is used in cellular packet data networks.

4.2 Hybrid Stations

With the large number of wired static hosts in the Internet geographically spread, we consider an approach wherein static hosts can optionally “subscribe” to wireless service. Such hosts will for all purposes act as any other mobile-station in the network. However, their primary role will be to serve as a gateway between the wireless cell and the wired Internet using the backbone connectivity that they possess. Such *hybrid stations* within the wireless cell can significantly alleviate the bottleneck at the channel around the base-station. *Note that the hybrid stations will not perform the role of a base-station. They will function as any other mobile-station, and send and receive packets accordingly.* The key difference in their functionality is that they act as gateways between the wireless and wired domains. Specifically, the hybrid station, upon receipt of a packet, routes the packet to the destination as if it were sending the packet³. On the reverse path, a packet from a static host destined for the mobile-station will reach the base-station which then re-routes the packet to the corresponding hybrid station that serves the mobile-station the packet is destined for. The overheads introduced by the additional routing (from the base-station to the hybrid station) can be overcome by adding appropriate intelligence to the hybrid station, or reserving a set of network addresses for each hybrid station that the station can then temporarily assign to mobile-stations it is serving. Hand-offs need not be performed from one hybrid station to another as long as the mobile-station is within the same cell. While relaxing this constraint might lead to better performance, we are in the process of studying the implications.

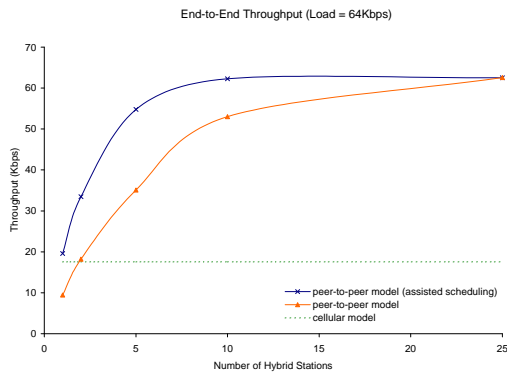


Figure 12: Throughput (Hybrid Stations)

The principal advantages gained because of the proposed approach is the removal of the bottleneck of the channel surrounding the base-station. Because of the distributed nature of the traf-

²In a related work [13] we have shown that in multi-hop wireless networks, the flow scheduling algorithm presented in this section achieves better performance than the node scheduling or link scheduling algorithm used in the STDMA approaches.

³This assumes the absence of *ingress filters* in the subnet of the hybrid station. If ingress filters are present, some form of address translation has to occur at the hybrid station.

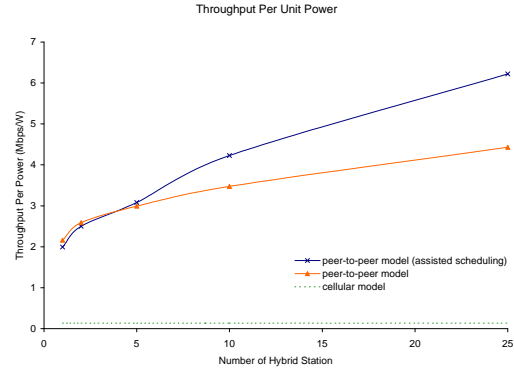


Figure 13: Throughput Per Unit Power (Hybrid Stations)

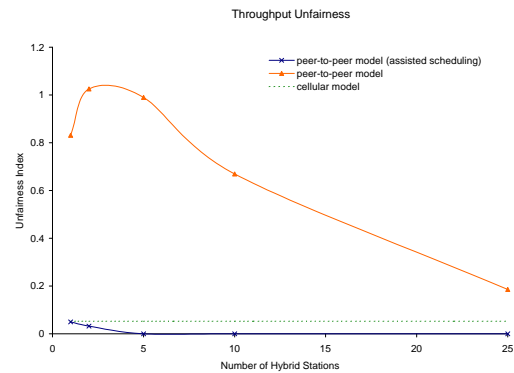


Figure 14: Fairness (Hybrid Stations)

fic flows in the cell, the performance of the peer-to-peer network model can be expected to significantly improve. In Figures 12, 13, and 14, we present results for the performance of the peer-to-peer model in the presence of hybrid stations. We increase the per-flow data rate from 32Kbps to 64Kbps to better show the improvement in network performance by using hybrid stations. Several observations can be made from the results: (i) The throughput, and throughput per unit power of the peer-to-peer network model with hybrid stations are significantly better than that of the cellular network model (for both the original and the assisted protocols), showing that network capacity is not limited by the bottleneck of the channel at the base-station. (ii) The number of hybrid stations required is not of the order of the number of mobile-stations in the network. Specifically, for the simulations shown, 10 hybrid stations are sufficient to serve 100 flows in the network. In the presence of 10 hybrid stations, the achieved throughput can reach the flow rate of 64Kbps when assisted scheduling is used. (iii) Network fairness is improved using the hybrid stations as shown in Figure 14. Using hybrid stations together with assisted scheduling, the fairness index is better than that in a pure cellular model because all flows now achieve higher throughputs (recall the definition of unfairness

index in Section 3.4). (iv) The results shown in Figure 12 are for a static scenario. In Figure 15, we can see that the performance degrades as before with increasing mobility, and can become less than that of the cellular model even in the presence of hybrid stations. The load and number of hybrid stations used for the mobile scenario are 32Kbps and 5 stations respectively.

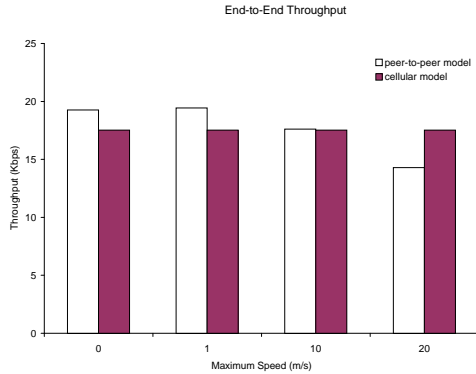


Figure 15: Impact of Mobility (Hybrid Stations)

The key issue in the hybrid stations approach lies in the assumption that static Internet hosts can be reconfigured as hybrid stations. However, such an approach for alleviating base-station bottlenecks is not an entirely new concept. In a loosely related context, the hybrid stations described above can be thought of to be IEEE 802.11 access points that provide additional service to mobile-stations subscribed to the cellular network. The BRAN HiperLAN/2 effort [28] includes a provision for the converged operation of 3G core networks with HiperLAN/2 access points with the primary goal of relieving cellular base-station bottlenecks. However, this effort is constrained to scenarios with a single hop between the mobile-stations and HiperLAN/2 access points. In a recent work, a converged architecture which supports multi-hop routes between mobile-stations and IEEE 802.11 access points has been proposed [15].

In another context, AT&T with the goal of exploiting its considerable cable infrastructure has plans to install access points at the houses of its customers with cable modems [32]. The goal is to provide local-area wireless access by installing such access points at as many houses as possible. Mobile users would then be able to access the Internet through these access points. While the example has been cited to substantiate our argument that the hybrid stations approach is not infeasible, notice that AT&T’s proposal is completely different from our solution because it does not use the existing cellular infrastructure, and mobile-users cannot use multi-hop paths to reach the access points. Therefore, we believe that if sufficient motivation is provided (e.g. free installation and a share of the revenue), wireline Internet users can be convinced to convert their static-stations to hybrid stations.

4.3 Dual Mode Service

Recall from Section 4.2 that even in the presence of the hybrid stations, mobility can have a negative influence on the performance of the network. Specifically, if stations in the network are mobile, the peer-to-peer model might be exposed to frequent route failures resulting in severe performance degradation. We use a dual mode

service to handle flows that thus suffer due to mobility. Specifically, flows are served in either the cellular or the peer-to-peer mode. The two modes are provided time-division access⁴ to the channel. In the initial state of the network, all flows are served in the peer-to-peer mode. Periodically, stations in the network convey to the base-station the performance (throughput) observed by each of their flows during the last observation period. If the throughput of a flow is less than a threshold value (set to the expected throughput of the flow for the same period in the cellular model) due to mobility, the flow is selected to be served in the cellular mode. The time division between the peer-to-peer and the cellular modes is done based on the number of flows selected to be served in the cellular mode. Flows served in the cellular mode do not receive any service in the peer-to-peer mode and vice-versa. Irrespective of which flows are selected to be served in the cellular mode, all mobile-stations in the network participate in packet forwarding during the peer-to-peer mode.

In the event of mobility, flows being served in the peer-to-peer mode can experience throughput degradation either because of network partitioning, or because of the overhead involved in the distributed route re-computation. In either case, by virtue of the flow’s throughput degradation, the flow will be selected by the base-station to be served in the cellular mode there-on. Flows are periodically reverted back to the peer-to-peer mode, and have their throughputs monitored.

| Variables | |
|-----------|--|
| n | → number of flows |
| $cTime$ | → time division allocation for cellular mode |
| $pTime$ | → time division allocation for peer-to-peer mode |
| ep | → time division cycle period ($cTime + pTime$) |
| mp | → mobility monitoring period |
| tp | → throughput monitoring period |
| rp | → cellular mode retention period |
| $Hmp(i)$ | → throughput over mp for flow i |
| $Htp(i)$ | → throughput over tp for flow i |
| $Hag(i)$ | → aggregate throughput for flow i |
| $Tag(i)$ | → time since start of flow i |
| $Trv(i)$ | → time for flow i to revert to peer-to-peer mode |
| $Loc(i)$ | → location or neighborhood information for node i |
| SF | → set of flows currently selected to be in cellular mode |
| $Cref$ | → reference cellular-mode throughput ($CThresh - b$) |

Figure 16: Variables Used in the Dual Mode Algorithm

Figures 17 and 18 present the pseudo-code for the algorithms, and Figure 16 lists the variables used in the algorithms. The channel is time-divided into periods of length ep , and each ep period is further divided between the peer-to-peer and cellular modes. At the end of every mp time units (set to multiples of ep), each mobile-station reports to the base-station its observed per-flow throughput, and location information (lines 4-6). We assume the use of a separate control channel for this information transfer. The base-station consolidates (lines 16-18) the information it receives from the mobile-stations.

An exception is made when the aggregate throughput $Hag(i)$ that flow i has received so far since its inception is higher than what flow i would have observed in the cellular network model. In this case, the base-station does not switch flow i into the cellular mode. The reasoning behind the exception is as follows: In order to maximize the network utilization the degree of spatial reuse has to be increased. Hence, it is desirable to have as few flows as possible in the cellular mode (since the channel allocation for the cellular mode is proportional to the number of flows served in that

⁴Note 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.

```

At Mobile Station  $i$ 
  After Every  $ep$  Time:
  1  $cTime \leftarrow ep - pTime$ 
  2 Participate in peer-to-peer mode for  $pTime$  period
  3 At end of  $pTime$  period, participate in cellular mode for  $cTime$  period

  After Every  $mp$  Time:
  4 Send  $(Hmp(i), Loc(i))$  to base-station, and reset  $Hmp(i)$ 
  5 Receive  $(SF, pTime)$ 
  6 Update link-layer state with  $SF$ 

  Peer-to-peer Mode:
  7 Forward packets in buffer (except those belonging to flows in  $SF$ )
  8 Use IEEE 802.11 for medium access and DSR for routing
  9 Update  $Hmp(i)$  based on packets received at destination

  Cellular Mode:
  10 If polled by base-station:
  11  $pkt \leftarrow SelectiveDequeue(SF)$ 
  12 send  $pkt$  to base-station

  SelectiveDequeue:
  13 Dequeue first packet  $pkt$  in link buffer such that
  14  $pkt$  belongs to one of the flows in  $SF$ 

```

Figure 17: Mobile-Station Operation

```

At Base-Station
  After Every  $mp$  Time:
  For each flow  $i$ :
  15 Receive  $(Hmp(i), Loc(i))$  from mobile-station  $i$ 
  16  $Htp(i) \leftarrow Htp(i) + Hmp(i)$ 
  17  $Hag(i) \leftarrow Hag(i) + Hmp(i)$ 
  18 if flow  $i$  is partitioned
  19  $SF \leftarrow SF + \{i\}$ 
  20 if  $Hag(i) > Cref * Tag(i)$  and  $i$  not partitioned
  21  $SF \leftarrow SF - \{i\}$ 
  22 if  $Trv(i) \leq currentTime$ 
  23  $SF \leftarrow SF - \{i\}$ 
  24  $pTime \leftarrow \frac{ep * (n - |SF|)}{n}$ 
  25 Send  $(SF, pTime)$  to mobile-stations

  After Every  $tp$  Time:
  For each flow  $i$ :
  27 if  $Htp(i) < Cref * tp$  and  $Hag(i) < Cref * Tag(i)$ 
  28  $SF \leftarrow SF + \{i\}$ 
  29  $Trv(i) \leftarrow currentTime + rp$ 
  30 reset  $Htp(i)$ 
  31

```

Figure 18: Base-Station Operation

mode). At the same time, the cellular mode is essential to ensure fairness, and to support flows that would have otherwise received very low throughputs in the peer-to-peer mode. Therefore, *in the fairness model supported by the dual mode service, the network utilization is maximized subject to the condition that every individual flow will receive a throughput that is no less than $CThresh - b$, where $CThresh$ is the expected throughput in the cellular network model, and b is a small constant*. Thus, flows that have aggregate rates $Hag(i)$ higher than what they would have received in a cellular network, are not switched to the cellular mode even if their instantaneous throughputs $Htp(i)$ are relatively lower.

When a flow is switched to the cellular mode, a timer $Trv(i)$ is associated with the flow (line 30). $Trv(i)$ is set to the amount of time the flow will stay in the cellular mode before it will be reverted back to the peer-to-peer mode (lines 23-24). Also, every mp time units, the base-station uses the location information to see if any of the source-destination pairs are partitioned. Flows corresponding to partitioned source-destination pairs are brought into the cellular mode (lines 19-20). Similarly, flows that are not partitioned anymore, but were switched to cellular because of a partition, are reverted back to the peer-to-peer mode (lines 21-22). The base-station also sends to the mobile-stations the updated set of selected flows SF and the corresponding peer-to-peer time to be used for the next set of ep periods (lines 25-26). Each mobile-station, upon




Figure 19: Throughput (Dual Mode Service)

receiving the updated information from the base-station, appropriately updates its link layer for selective deque (lines 13-14) and sets its $cTime$ and $pTime$ timers. It then functions in the peer-to-peer mode for $pTime$ time units (line 2). In the peer-to-peer mode, the mobile-station forwards packets in its link buffer that do not belong to any of the flows in SF .

Finally, after $pTime$ time units, the nodes stop functioning in the peer-to-peer mode and switch to the cellular mode. In the cellular mode, the base-station explicitly polls sources for packets, or forwards packets to destinations (based on a flow-based round-robin scheme). In this phase, polled mobile-stations forward only packets belonging to flows in SF to the base-station. The base-station immediately forwards packets to the destination if the destination happens to be in the same cell, or sends it out to the backbone distribution network. Flows coming into the cell from the distribution network are served under the same fairness scheme used for serving outgoing or intra-cell flows.

The advantage of the proposed dual mode of operation is that flows in the worst mobility scenario will receive throughput that is very close to that of the cellular model. Note that because of the overhead involved in periodically switching flows to the peer-to-peer mode, the effective performance would be lower than that of a pure cellular model. In Figure 19, we present the results for the mobile scenarios when the dual mode of operation is used. It can be seen that when the mobility is low enough for the performance of the peer-to-peer network model not to be impacted significantly, the dual mode of operation does not kick in and the throughput observed is that of the peer-to-peer model. On the other hand, when the mobility increases, more number of flows are switched to the cellular model. Note that the throughput degrades because as more number of flows are shifted to the cellular mode, a greater fraction of the time is used to employ the cellular network model decreasing the benefits of using the peer-to-peer network model.

5. ISSUES AND SUMMARY

5.1 Issues

In this section we discuss some issues with the proposed set of approaches, including practical considerations that need to be addressed before they can be deployed.

- *Ad-hoc Network Protocols*: While this paper only makes a

case for adopting the ad-hoc network model in cellular wireless packet data networks, the question of whether or not to adopt protocols developed for the ad-hoc network environment is also an interesting one. There exists an abundant body of research in the area of ad-hoc network protocols, and any ability to reuse the same would be both desirable and potentially beneficial. There exist two probable arguments for why ad-hoc network protocols are more desirable than purely base-station located protocols: (i) For certain layers of the protocol stack, any amount of involvement from the base-station will not provide to the same degree the requisite intelligence to suit the ad-hoc network model. For example, vanilla-TCP as the transport layer has been shown to perform badly in ad-hoc networks [3, 12], and appropriate enhancements have been proposed. Such enhancements can be directly used to address problems due to the adoption of the ad-hoc network model, while any base-station located protocol may not be able to address the same issues. (ii) For each base-station located protocol, a control channel needs to be established between the base-station and the mobile-stations in order for the base-station to obtain the parameters to be input to the protocol, and to convey to the mobile-stations the decisions made by the protocol. Hence, there exists a fine balance between the performance benefits due to the centralized execution of the protocol, and the performance bottlenecks due to the control channels between the base-station and the mobile-stations. For protocols wherein the baseline performance difference between the centralized and the distributed versions is not significant enough to warrant a base-station located protocol, conventional ad-hoc network protocols are more appropriate to use.

- *Multiple-Cell Environment:* While we have primarily considered only a single-cell environment in this paper, we now discuss some issues pertaining to a multiple-cell environment. In a multiple-cell environment the cellular model will also be affected by any inter-cell mobility. All discussions and approaches presented thus far will still hold good as long as a mobile-station stays within a cell. When the mobile-station hand-offs into a neighboring cell, we envision the hand-off to take place as it would in a conventional cellular model, and then for the mobile-station to participate in the newly proposed approaches. Any overhead or performance degradation because of the inter-cell hand-off will be common to both the existing cellular model and the newly proposed approaches. However, one issue in question is when a hybrid station is serving a mobile-station and the mobile-station moves out of the cell. In the solution adopted in this paper wherein the base-station re-routes packets destined for the mobile-station to the hybrid station, this situation can be handled easily as the base-station is aware of the hand-off. However, if the hybrid station allocates network addresses for mobile-stations it serves, then when packets arrive at the hybrid station directly, it should be aware of the fact that the mobile-station is no longer within the same cell. A possible solution to this problem is for the base-station to periodically broadcast the current list of mobile-stations within the cell (or the incremental update since the last broadcast). The hybrid stations on receiving the broadcast can update its state accordingly. Also, this can help in the hybrid station using the network addresses that were allocated for mobile-stations that have left the cell, for other mobile-stations.
- *Security:* While source-destination pairs in a cellular net-

work communicate directly through the base-station, in the proposed approach they would have to rely on other peer stations to relay packets. This can potentially be a security issue. Although the base-station can be treated as a trusted entity, other peer stations in the network need not be necessarily trustworthy. Hence, the proposed approach might potentially suffer from the same security problems that exist in pure peer-to-peer networks. Authentication, confidentiality, integrity, availability, and non-repudiation are the key security issues that need to be solved anew for peer-to-peer networks [34]. However, by virtue of mobile-stations always having connectivity with the base-station (and hence the backbone infrastructure), authentication, availability, and non-repudiation can still be solved by the same solutions that exist in conventional cellular networks. Moreover, confidentiality in the proposed model is only as big a problem as in cellular networks, since transmissions are being done on a shared medium and other stations can snoop information although they are not being used as relays. However, integrity is a genuine problem that needs to be addressed specifically in the proposed approach, in order for it to be a practical solution. In particular, the issue is, when packets are being forwarded by other peer stations, how it can be ensured that the packets are not “corrupted” on the path to the destination. While sophisticated cryptography-based solutions can be explored, a simpler solution is more desirable. For example, in a simple solution that would detect any compromise of integrity, stations can randomly send duplicate data packets across to the base-station through a separate control channel dedicated for the purpose. Depending on whether the flow is intra-cell or inter-cell, either the destination or the base-station can verify if packets are indeed being corrupted by comparing the duplicates. If the duplicates differ, it can be inferred that one of the stations on the multi-hop path is compromising the integrity of the information being forwarded.

- *Pricing:* While the proposed approaches rely on cooperation between mobile-stations, a very reasonable question to ask is: why would the stations want to participate in the cooperative operations? While mobile-stations that have data to send might participate because of the perceived performance improvement, the same cannot be said about the hybrid stations and mobile-stations that do not have any data to send. One potential scheme to convince such stations to participate is to provide them with a share of the excess revenue generation because of the cooperation. Specifically, customers who are willing to “convert” their static hosts into hybrid stations can be provided with a share of the revenue generated from flows that they serve.

5.2 Summary

The scarce resources of wireless packet data networks are being severely exposed by the significant growth in mobile-user populace. While several approaches at different layers of the network protocol stack have been proposed to improve the performance, a new set of approaches proposed in literature has focused on alternate network models to the traditional cellular network models. A commonality between these approaches is that they employ peer-to-peer communication. In this paper, we study the true impact of using peer-to-peer communication in a cellular wireless packet data environment. We conclude that while peer-to-peer communication has its benefits in terms of better spatial reuse characteristics, the benefits get translated into *only* better throughput per unit power, with the throughput showing degradation. We identify and discuss

the reasons behind this observation. Finally, we also propose a set of three approaches - assisted scheduling, hybrid stations, and dual mode access, that when used in tandem with peer-to-peer communication translate its spatial reuse benefits into better throughput, and power consumption performance. The proposed approaches achieve the performance enhancement while providing fair service and being resilient to mobility.

6. REFERENCES

- [1] ADVENTIS. Beating the mobile data access bottleneck - commentary from Adventis. <http://www.adventis.com>, 2000.
- [2] G. Aggelou and R. Tafazolli. On the relaying capacity of next-generation GSM cellular networks. *IEEE Personal Communications Magazine*, 8(1):40–47, Feb. 2001.
- [3] V. Anantharaman and R. Sivakumar. A microscopic analysis of TCP performance over mobile ad-hoc networks. In *Proceedings of ACM SIGMETRICS* (poster paper), Marina Del Rey, CA USA, June 2002.
- [4] A. Bakre and B. Badrinath. Implementation and performance evaluation of indirect TCP. *IEEE Transactions on Computers*, 46(3):260–278, Mar. 1997.
- [5] J. Broch, D. B. Johnson, and D. A. Maltz. The dynamic source routing protocol for mobile ad hoc networks. Internet Draft; draft-ietf-manet-dsr-07.txt, Feb. 2002.
- [6] 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.
- [7] CAIDA. Analysis of Internet traffic. <http://www.caida.org/analysis/performance>.
- [8] I. Chlamtac and A. Farago. Making transmission schedules immune to topology changes in multi-hop packet radio networks. *IEEE/ACM Transactions on Networking*, 2(1):23–29, Feb. 1994.
- [9] D. J. Goodman. *Wireless Personal Communications Systems*. Addison-Wesley, Reading, MA USA, 1997.
- [10] J. Gronkvist. Traffic controlled spatial reuse TDMA in multi-hop radio networks. In *Proceedings of IEEE PIMRC*, Boston, MA USA, Sept. 1998.
- [11] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, Mar. 2000.
- [12] G. Holland and N. Vaidya. Analysis of TCP performance over mobile ad hoc networks. In *Proceedings of ACM MOBICOM*, Seattle, WA USA, Aug. 1999.
- [13] H.-Y. Hsieh and R. Sivakumar. IEEE 802.11 over multi-hop wireless networks: Problems and new perspectives. In *Proceedings of IEEE VTC Fall*, Vancouver, Canada, Sep. 2002.
- [14] H.-Y. Hsieh and R. Sivakumar. Performance comparison of cellular and multi-hop wireless networks: A quantitative study. In *Proceedings of ACM SIGMETRICS*, Boston, MA USA, June 2001.
- [15] H.-Y. Hsieh and R. Sivakumar. Internetworking WWANs and WLANs in next generation wireless data networks. In *Proceedings of 3G Wireless and Beyond*, San Francisco, CA USA, May 2002.
- [16] L. Kleinrock and J. Silvester. Spatial reuse in multihop packet radio networks. *Proceedings of IEEE*, 75(1):156–166, Jan. 1987.
- [17] J. C. Liberti and T. S. Rappaport. *Smart Antennas for Wireless Communications: IS-95 and Third Generation CDMA Applications*. Prentice Hall, Upper Saddle River, NJ USA, 1999.
- [18] Y.-D. Lin and Y.-C. Hsu. Multihop cellular: A new architecture for wireless communications. In *Proceedings of IEEE INFOCOM*, Tel-Aviv, Isarel, Mar. 2000.
- [19] S. Lu, T. Nandagopal, and V. Bharghavan. A wireless fair service algorithm for packet cellular networks. In *Proceedings of ACM MOBICOM*, Dallas, TX USA, Oct. 1998.
- [20] H. Luo, S. Lu, and V. Bharghavan. A new model for packet scheduling in multihop wireless networks. In *Proceedings of ACM MOBICOM*, Boston, MA USA, Aug. 2000.
- [21] R. Nelson and L. Kleinrock. Spatial TDMA: A collision-free multihop channel access protocol. *IEEE Transactions on Communications*, 33(9):934–944, Sept. 1985.
- [22] The network simulator ns-2. <http://www.isi.edu/nsnam/ns/>, Nov. 2000.
- [23] C. Qiao and H. Wu. iCAR: An integrated cellular and ad-hoc relay system. In *Proceedings of IC3N*, Las Vegas, NV USA, Oct. 2000.
- [24] S. Rappaport and L.-R. Hu. Microcellular communication system with hierarchical macrocell overlays: Traffic performance model and analysis. *Proceedings of the IEEE*, 82(9):1383–1397, Sept. 1994.
- [25] S. Seshan, H. Balakrishnan, and R. Katz. Handoffs in cellular wireless networks: The Daedalus implementation and experience. *Kluwer Journal on Wireless Personal Communications*, 4(2):141–162, Mar. 1997.
- [26] P. Sinha, N. Venkitaraman, R. Sivakumar, and V. Bharghavan. WTCP: A reliable transport protocol for wireless wide-area networks. In *Proceedings of ACM MOBICOM*, Seattle, WA USA, Aug. 1999.
- [27] 3G TR 25.924 V1.0.0. *3GPP TSG-RAN; Opportunity Driven Multiple Access*, Dec. 1999.
- [28] ETSI TR 101 957 V1.1.1. *BRAN; HiperLAN/2; Requirements and Architecture for Internetworking Between HiperLAN/2 and 3rd Generation Cellular Systems*, Aug. 2001.
- [29] N. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless LAN. In *Proceedings of ACM MOBICOM*, Boston, MA USA, Aug. 2000.
- [30] A. J. Viterbi. *CDMA: Principles of Spread Spectrum Communications*. Addison-Wesley, Reading, MA USA, 1995.
- [31] E. K. Wesel. *Wireless Multimedia Communications*. Addison-Wesley, Reading, MA USA, 1997.
- [32] J. H. Winters. Enhancements to wireless local area networks. Georgia Institute of Technology Seminar, July 2001.
- [33] X. Wu, S.-H. Chan, and B. Mukherjee. MADF: A novel approach to add an ad-hoc overlay on a fixed cellular infrastructure. In *Proceedings of the IEEE WCNC*, Chicago, IL USA, Sept. 2000.
- [34] L. Zhou and Z. Haas. Securing ad hoc networks. *IEEE Network Magazine*, 13(6):24–30, Nov.-Dec. 1999.