A Hybrid Network Model for Cellular Wireless Packet Data Networks *

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

***Abstract*— In this paper, we propose a hybrid network model called *Sphinx* for cellular wireless packet data networks. *Sphinx* uses a peer-to-peer network model in tandem with the cellular network model to achieve *higher throughput and lower power consumption*. At the same time, *Sphinx* avoids the typical pitfalls of the pure peer-to-peer network model including *unfair resource allocation, and throughput degradation due to mobility and traffic locality*. We present simulation results showing that *Sphinx* outperforms the cellular network model in terms of throughput and power consumption, and achieves better fairness and resilience to mobility than the peer-to-peer network model.**

I. Introduction

The increasing population of mobile Internet users in recent years has severely exposed the limitations of existing wireless packet data networks. Next generation wireless systems have projected data rates significantly lower when compared to their wireline counterparts. While this has inspired several research efforts toward improving wireless data performance, an interesting solution proposed by some recent research works is using a combination of the cellular network model with a peer-to-peer network model hitherto used only in a special class of wireless networks called ad-hoc networks. By using peer-to-peer communication in a conventional cellular network, performance improvements such as increased data rate [1], reduced transmission power [2], enhanced network capacity [3], better load balancing [4], and extended coverage area [5] have been demonstrated.

Notwithstanding the performance enhancements achieved when incorporating peer-to-peer communication in a cellular network, the shortcomings of such hybrid approaches lie in the adoption of the peer-to-peer network model on an "as-is" basis without addressing its limitations and impacts on the overall performance, namely *greater vulnerability to mobility* [6], *unfairness problems* [7], and *severe throughput degradation for certain kinds of traffic locality* [8].

In this paper we contend that while it is desirable to leverage the advantages of the peer-to-peer network model in terms of throughput and power consumption, its pitfalls in terms of unfairness and performance degradation due to mobility and traffic locality, need to be addressed and if possible eliminated. We present a hybrid network model called *Sphinx* that allows for the optimal use of the cellular and peer-to-peer network models in tandem. We show that *Sphinx* exhibits the high throughput and low power consumption characteristics of peer-to-peer networks, and at the same time achieves the fairness and mobility resilient characteristics of cellular networks. Furthermore, we believe that the fundamental design elements of *Sphinx* can be applied in any hybrid approach to alleviate the negative side-effects of using a peer-to-peer model.

The rest of the paper is organized as follows: In Section II we present the motivation for the hybrid network model. In Section III we describe the *Sphinx* hybrid network model. In Section IV we present simulation results evaluating the performance of *Sphinx*, and in Section V we discuss several design issues. Finally, in Section VI we discuss some related work and conclude the paper.

II. Motivation

A *cellular network* consists of a collection of mobile stations served by a central coordinating entity called the base-station. The base-station arbitrates the channel allocation between the mobile stations. In contrast, in a *peer-to-peer network*, mobile stations establish a network without the aid of a backbone infrastructure. Sources and destinations communicate with each other through multi-hop paths consisting of peer-stations in the network. In the followings, we identify the trade-offs between the two network models, and thus motivate the need for a hybrid network model.

A. Peer-to-Peer Networks vs. Cellular Networks

- **Throughput:** Stations in the peer-to-peer network use short-range transmissions. This has two consequences: (i) an increase in the degree of spatial reuse in the network, and (ii) an increase in the number of hops to be traversed end-to-end by a packet. Assuming that nodes are distributed randomly on a unit area disk, the number of simultaneous transmissions possible (spatial reuse) increases with the number of nodes n in the network, namely $O(n)$. On the other hand, the distance in number of hops between the source and destination increases as $O(\sqrt{n})$, resulting in a per-flow achievable rate that is of the order of $O(\frac{C}{\sqrt{n}})$, where C is the capacity of the underlying channel [9]. Compared to the per-flow rate achievable in a cellular network $O(\frac{C}{n})$, the peer-to-peer network has a higher throughput by a factor of $O(\sqrt{n})$.

*This work was supported in part by NSF (under grant #ANI-0117840) and Yamacraw (http://www.yamacraw.org).

- **Power Consumption:** In wireless communications, the received signal power is of the order of $O(\frac{1}{d^k})$ when the transmitter and the receiver are separated by the distance d. The attenuation factor k usually varies between 2 and 6 depending on the signal propagation model used. In a cellular network with cell radius d, the power used for a packet transmission is then of the order of $O(d^k)$. On the other hand, in a peer-to-peer network where transmission range is reduced to $O(\frac{d}{\sqrt{n}})$, and the number of hops is increased to $O(\sqrt{n})$, the power used for transmitting a packet end-to-end can be computed to be lower than that in the cellular network by a factor of $O(\sqrt{n^{k-1}})$.

B. Using Peer-to-Peer Communications in Cellular Networks

Although the above factors make the peer-to-peer model an attractive alternative to the cellular model, it suffers from three critical drawbacks when used in cellular packet data networks:

- **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, flows will primarily have either the source or the destination outside the cell. All such flows will then use the base-station as the destination even when functioning in the peer-to-peer mode. As the number of such flows increases, the contention in the local neighborhood of the base-station will increase resulting in throughput degradation. In the extreme scenario when all flows are non-local, the throughput in the peer-to-peer model may be even lower than that in the cellular model due to significantly reduced spatial reuse and inefficiency of the distributed medium access protocol used in the multi-hop peer-to-peer network model [8].
- **Fairness:** In cellular networks, each flow in the network gets a fair allocation of the network resources under the base-station's coordination. However, in peer-to-peer networks, the throughput observed by a flow will depend on the absolute and relative locations of the source and destination of the flow. Although sophisticated protocols can be used to reduce unfairness [7], due to lack of a central coordinating entity the fairness exhibited by the peer-to-peer model will not be as good as the cellular model. While such a service model is sufficient in military environments or emergency-relief applications where there are no other communication options, it will not suffice in a commercial environment where users are less likely to tolerate unfair allocation of resources.
- **Impact of Mobility:** In a peer-to-peer network, mobility of stations can result in two distinct phenomena: (i) network partitions, and (ii) route failures and re-computations. Both result in a degradation of the throughput [6]. However, in a cellular network, since the stations communicate directly with the base-station, and the base-station coverage area is much larger, they do not experience either of the above phenomena frequently. Hence, the flow throughputs are relatively less affected by mobility in cellular networks. In the rest of the paper, we use the term *mobility* to refer to the movement of mobile stations within the coverage area of the base-station in the cellular model.

The goal of the *Sphinx* hybrid network model is thus to combine the throughput and power consumption advantages of the peer-to-peer model with the fairness and resilience to mobility advantages of the cellular model.

III. The Hybrid Network Model

A. Overview

The *Sphinx* hybrid network model is based on a regular cellular infrastructure with a base-station that supports a dual mode of operation. The cellular and peer-to-peer modes are provided time-divisioned access to the channel.[1] When operating in the cellular mode, mobile stations communicate directly with the base-station as in a conventional cellular network, and do not interact with other mobile stations. In the peer-to-peer mode, mobile stations act as routers for other mobile stations in the network without the aid from the base-station. In the initial state of the system, all flows are served in the peer-to-peer mode by virtue of the higher throughput and lower power consumption. A flow served in the peer-to-peer mode is switched to the cellular mode in *Sphinx* whenever its performance (throughput) degrades (lower than the expected throughput of the flow in a pure cellular network) due to topology constraints or mobility.

Periodically, mobile stations monitor the throughputs for each of the flows which they act as the sources. If the throughput of a flow during the last measurement period is less than a threshold value, the mobile station requests the base-station to serve the flow in the cellular mode. Flows served in the cellular mode do not receive any service in the peer-to-peer mode and vice-versa. However, irrespective of which flows are selected to be served in the cellular mode, all mobile stations 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-computations. In either case, the flow will be served in the cellular mode thereon. Network partitions are detected by mobile stations experiencing route errors for a threshold amount of time. Even in the absence of route errors and network partitions, a flow served in the peer-to-peer mode can still be switched over to the cellular mode because of experiencing low throughput.

By default, flows served in the cellular mode are periodically reverted back to the peer-to-peer mode, and have their

[1] 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.

throughputs monitored. A flow selected for cellular service because of a network partition is reverted back when the network re-configures such that the flow is no longer partitioned (when the routing protocol discovers routes).

B. Algorithm

n	$\longrightarrow$ number of flows in the network
SF	$\longrightarrow$ set of flows currently operated in cellular mode
cT	$\longrightarrow$ time division allocation for cellular mode
rp	$\longrightarrow$ cellular mode repetition period
mp	$\longrightarrow$ throughput monitoring period
up	$\longrightarrow$ division update period
$Tp(i)$	$\longrightarrow$ route partition timer (timeout=pp) for flow i
$Ts(i)$	$\longrightarrow$ cellular mode sojourn timer (timeout=sp) for flow i
$M(i)$	$\longrightarrow$ mode of operation {`CELLULAR, PEER`} for flow i
$P(i)$	$\longrightarrow$ peer-to-peer mode connectivity {`PARTITION, CONNECT`} for flow i
$g(i)$	$\longrightarrow$ throughput over mp for flow i
$G(i)$	$\longrightarrow$ aggregate throughput for flow i
$R(i)$	$\longrightarrow$ reference throughput for flow i

Fig. 1. Variables Used in the Hybrid Model Algorithm

Fig. 1 lists the variables used in the algorithm of the hybrid network model, and Fig. 2 presents the algorithm at the mobile and base-station. The channel is time-divisioned into periods of length rp, which is further divided between the cellular and peer-to-peer modes (lines 1-2). Each mobile station periodically monitors the throughput in an interval of length mp, and requests the base-station to serve its own flow in the cellular or peer-to-peer mode (lines 3-6) based on the observed performance. The base-station periodically consolidates such request and broadcasts the latest time-division to use in the next period of length up based on the number of flows served in the cellular mode (lines 36-37). We assume the use of a separate control channel for the above information exchange. For clarity, in the following discussion we also assume that each mobile station acts as a source for exactly one flow, and hence the number of flows n is equal to the number of nodes in the network.

Every mp time units (set to multiples of rp), each mobile station monitors its per-flow throughput and decides if it needs to switch its mode of operation. Specifically, a mobile station i keeps track of the short-term throughput $g(i)$ (over the last mp period), and long-term throughput $G(i)$ (since its inception), for the flow it acts as the source. A mobile station switches its flow to the cellular mode only if both the short-term and long-term throughputs are lower than the reference throughput $R(i)$ it would have observed in a pure cellular network. The reference throughput is a lower bound on the throughput a mobile station desires to enjoy, and is decided when the mobile station initially joins the network (say, as part of the registration phase). For simplicity we assume that the reference throughput is $\frac{C}{2*n}$ for intra-cell flows, and $\frac{C}{n}$ for flows with the source or destination outside the cell. Therefore, *in the fairness model supported by Sphinx, every individual flow will receive a throughput that is no less than $R(i) - \delta$, where $R(i)$ is the expected throughput in the cellular network model, and δ is a small constant.*

```
At Mobile Station i
    Every rp time:
1       participate in cellular mode for cT period
2       participate in peer-to-peer mode for the remaining period
    Every mp time:
3       if M(i) is PEER and g(i) < R(i) and G(i) < R(i)
4           send request[i, JOIN] to the base-station
5       elseif M(i) is CELLULAR and G(i) > R(i) and P(i) is CONNECT
6           send request[i, LEAVE] to the base-station
    Selective Dequeue:
7       in cellular mode
8           if M(i) is CELLULAR
9               dequeue only packets belonging to flow i
10          else  do not dequeue any packets
11      in peer-to-peer mode
12          if M(i) is PEER
13              dequeue head-of-line packets
14          else  dequeue only packets not belonging to flow i
    Receive division[time t, set S]:
15      cT ← t
16      if i ∈ S
17          M(i) ← CELLULAR
18          start Ts(i) if not set
19      else
20          M(i) ← PEER
21          stop Ts(i) if set
    Callback from routing protocol with reason r:
22      if r is ROUTE-ERROR
23          start Tp(i) if not set
24      elseif r is ROUTE-OKAY
25          P(i) ← CONNECT
26          stop Tp(i) if set
27          if M(i) is CELLULAR and Ts(i) expired
28              send request[i, LEAVE] to base-station
    When partition timer Tp(i) expires:
29      P(i) ← PARTITION
30      if M(i) is PEER
31          send request[i, JOIN] to base-station
32      else  start route probes until P(i) is CONNECT
    When sojourn timer Ts(i) expires:
33      if P(i) is CONNECT
34          send request[i, LEAVE] to base-station

At Base-Station
    Every rp time:
35      participate in cellular mode for cT period
    Every up time:
36      cT ← rp * |SF|/n
37      broadcast division[cT, SF] to mobile stations
    Receive request[mobile station i, action a]:
38      if a is JOIN
39          SF ← SF + {i}
40      else  SF ← SF − {i}
```

Fig. 2. Algorithm of the Hybrid Network Model

A flow in the peer-to-peer mode may be switched to the cellular mode due to topology or mobility induced network partitions, as well as throughput degradation. Mobile stations discover a network partition via the callback from the peer-to-peer routing protocol (lines 22-28). A mobile station operating in the peer-to-peer mode considers itself partitioned from the destination and switches its flow to the cellular mode if such route errors last for more than a duration of length pp (lines 29-32). Note that even if route errors do not trigger a switch to the cellular mode, the mobile station may still switch to the cellular mode due to throughput degradation (lines 3-4).

When a mobile station i switches its flow to the cellular mode, a timer $Ts(i)$ is associated with the flow (line 18). The timeout sp is 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 33-34). Flows that are not partitioned anymore, but were switched to the cellular mode because of a partition, are also reverted back to the peer-to-peer mode (lines 27-28).

Every up time units, the base-station consolidates requests from mobile stations to join or leave the cellular mode. It then sends to the mobile stations the updated set of selected flows SF to operate in the cellular mode and the corresponding time-division cT to be used for the next up period (lines 36-37). Each mobile station, upon receiving the updated information from the base-station, sets its cT timer, updates its mode of operation $M(i)$, and appropriately configures its link layer for selective dequeue. A mobile station in the cellular mode can dequeue packets in its link buffer only if it is selected to operate in the cellular mode (lines 7-10). However, in the peer-to-peer mode all mobile stations dequeue and forward packets normally, except for those flows that can be dequeued in the cellular mode (lines 11-14).

When operating in the cellular mode, the base-station and mobile stations communicate directly as in a conventional cellular network. For example, the base-station can perform direct polling (as in WLANs) or broadcast a transmission schedule (as in WWANs) to the subset of mobile stations in SF. Flows coming into the cell from the distribution network are served under the same fairness scheme used for serving outgoing or intra-cell flows.

IV. Simulation Results

A. Simulation Model

We use the *ns-2* network simulator [10] for simulations presented in this paper. The followings are the elements of the simulation model: (i) **Topology:** All simulations are run using a topology with 100 nodes randomly distributed in a 1500m $\times$ 1500m grid. The average, minimum, and maximum degrees of a node in the network is approximately 7.3, 1, and 14 respectively with a standard deviation of 3. The density of the network is thus non-uniform, hence not causing any bias toward one of the two network models. (ii) **Physical Layer:** The signal propagation model used is a combination of a free space propagation model ($\frac{1}{r^2}$) and a two-ray ground reflection model ($\frac{1}{r^4}$). The channel data rate is set to 2Mbps and the transmission range to 250m. (iii) **Medium Access and Routing Layers:** The cellular model is simulated using the IEEE 802.11 MAC protocol in PCF mode and the peer-to-peer model is simulated using the DCF mode. We use the "Dynamic Source Routing" (DSR) protocol [6] for the routing layer in the peer-to-peer model. (iv) **Traffic:** In all of the simulations, TCP flows are used as the transport protocol by virtue of its prevalence in the Internet. A constant-bit-rate (CBR) application is used on top of TCP to have fine control over the offered load. We use 50 flows with randomly chosen source-destination pairs in the network. Although various loads were used, we present results only for a per-flow offered load of 48Kbps. (v) **Mobility:** We use the way-point mobility model [6] supported in *ns-2* for generating the scenarios with mobility. Node movements with zero-second pause time and maximum speed from 0 m/s to 20 m/s are simulated. As mentioned in Section II we consider only one base-station cell in the simulations, and node movements are confined within the cell. For results presented in this section, mobility is enabled only when considering the impact of mobility.

B. Impact of Traffic Locality

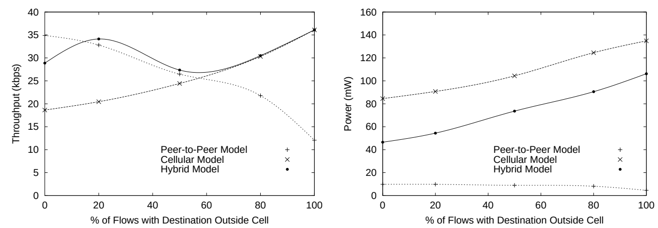

(a) Throughput (b) Power Consumption

Fig. 3. Impact of Traffic Locality

Fig. 3(a) shows the throughput performance of the three network models for different percentages of flows having either the source or destination outside the cell. The performance of the peer-to-peer model starts decreasing as the percentage of such flows increases, and falls down to much below the performance of the cellular network for a scenario where 100% of the flows have either their source or destination outside the cell. It is thus evident that using peer-to-peer communication in cellular networks may exhibit poor throughput performance for certain traffic pattern. While the throughput of *Sphinx* does decrease initially, it does not fall below the throughput of the cellular network. Instead, when the cellular model starts performing better than the peer-to-peer model, *Sphinx* adapts itself (by switching more flows to the cellular mode) and starts tracking the performance of the former. In essence, *Sphinx* adapts its point of operation between the pure peer-to-peer mode, and the pure cellular mode depending on the state of the network, and always achieves a performance level that is equal to or better than that of the other two models. Fig. 3(b) shows the average per-node power consumption for the same scenario. As evident, the power consumption for the cellular model is much higher than that in the peer-to-peer model. Moreover, as the number of non-local flows increases, the cellular model consumes more power due to throughput increase (as observed in Fig. 3(a)). On the other hand, the peer-to-peer model consumes less power due to reduced spatial reuse and packet transmissions (which account for throughput decrease). Because *Sphinx* adapts itself to track the higher throughput of the two models, as the number of non-local flows increases, more and more flows are switched to the cellular mode which inherently requires higher transmit power. Note that even when all the flows have the destination and source inside the cell, in *Sphinx* some flows are operated in the cellular mode (its power consumption is higher than a pure peer-to-peer model at the 0% point). The difference is due to the fairness model that *Sphinx* supports, as explicated in the next section.

C. Fairness

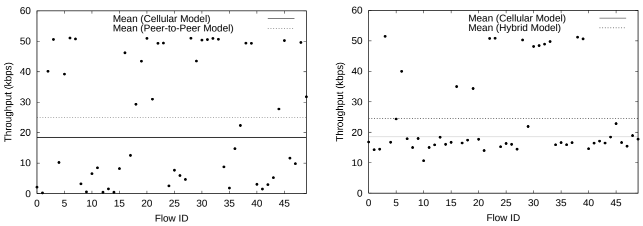

(a) Peer-to-Peer Model (b) Hybrid Model

Fig. 4. Throughput Distribution

In Fig. 4(a) and Fig. 4(b), we present the per-flow throughput distribution in the peer-to-peer model and *Sphinx* respectively. The average throughput of the cellular model is also shown in both figures for comparison. It can be seen that although the average throughputs of the *Sphinx* and the peer-to-peer model are both higher than that of the cellular model, *Sphinx* achieves a much better distribution of throughput with none of the flows observing throughputs much lower than the cellular model. This is an important property of *Sphinx* because some flows in the peer-to-peer model will "starve" (as discussed in Section II) despite the high average throughput exhibited by the network. *Sphinx* thus maximizes network capacity while ensuring no flows will suffer from throughput degradation as in a pure peer-to-peer model. The marginal throughput deviation experienced by some flows is the allocation deviation δ introduced in the algorithm.

D. Impact of Mobility

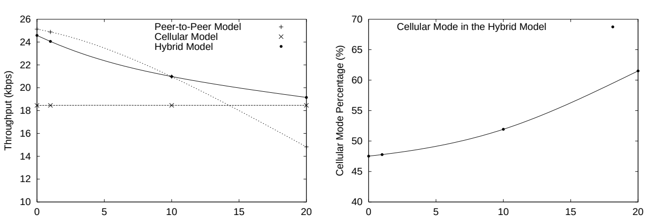

(a) Throughput (b) Mode of Operation

Fig. 5. Impact of Mobility

In Fig. 5(a), we present simulation results that illustrate the impact of mobility on the throughput in the three network models. The performance of the peer-to-peer network model degrades significantly with increasing mobility. As discussed in Section II, the throughput degradation is because of network partitions and route re-computations. The performance of the cellular model is not affected by mobility as the mobility simulated is intra-cell mobility. *Sphinx*, on the other hand, performs better than the cellular model in the low mobility scenarios (where a majority of flows function in the peer-to-peer mode), and performs better than the peer-to-peer model at high mobility rates when flows are switched to the cellular mode either because of network partitions or low throughputs due to frequent route re-computations. Fig. 5(b) shows the percentage of time *Sphinx* spends in the cellular mode. It shows that as mobility increases, *Sphinx* adapts its point of operation toward the cellular mode to achieve a better performance at high mobility rates.

E. Summary of Results

We have shown that the *Sphinx* hybrid model achieves better throughput, and power consumption performance than conventional cellular networks. At the same time, it manages to solve the unfairness, mobility induced performance degradation, and traffic-locality related throughput degradation that pure peer-to-peer networks suffer from.

V. Design Issues

In this section, we briefly discuss several design issues and alternative approaches for the algorithm used by *Sphinx*.

- **Throughput Monitoring:** *Sphinx* requires periodic throughput monitoring at each mobile station. While throughput monitoring can be performed either at the source or at the destination, in this paper we present an approach where the source is aware of its flow throughput. The source can obtain such information by monitoring the reverse traffic (e.g. ACKs in TCP traffic), or by having the destination feedback the information to the source periodically. Note that if the source cannot obtain the throughput of the flow it serves, the flow can still be switched to the appropriate mode by having the destination send the join or leave request directly to the base-station on behalf of the source.
- **Base-Station Centric vs. Mobile-Station Centric:** In the algorithm presented in Fig. 2, mobile stations are responsible for monitoring the throughput and making the decision as to which mode to operate in. While such design places additional overheads on the mobile station, the advantage is the reduced complexity at the base-station. An alternative approach is to let the base-station perform throughput monitoring and make the switching decisions for all flows in the network (the mobile station only needs to periodically feedback the throughput it observes). The overheads at the mobile stations will be greatly reduced at the expense of the increased complexity at the base-station. Interested readers are referred to [8] for a base-station centric approach for the hybrid network model.
- **Mode Multiplexing:** In the hybrid network model, flows are restricted to be served exclusively either in the cellular

or peer-to-peer mode by the use of selective dequeue. If flows are allowed to transmit in both modes, they will experience frequent out-of-order delivery at the destination, which will in turn adversely impact TCP's performance. However, note that during the transition from the peer-to-peer to cellular mode, packets transmitted in the peer-to-peer mode just before entering the cellular mode will still be delivered later than the packets transmitted during the cellular mode. A possible solution to address this problem and relieve the constraint is to use a transport protocol that is capable of multiplexing multiple paths (as opposed to TCP that is designed only for a single path) [11].

- **Multiple Channels:** The *Sphinx* hybrid network model is a generic approach that can be deployed in both local-area (WLANs) and wide-area (WWANs) wireless networks. Although we present the algorithm as one that multiplexes the cellular and peer-to-peer modes over a single channel (in an environment such as IEEE 802.11 WLAN), it can be extended to a multi-channel environment such as 3G WWAN. Several multi-channel MAC protocols for peer-to-peer networks have been proposed that show the ability to achieve performance improvement by using multiple channels (see, for example, related work presented in [12]). Therefore, if mobile stations use such multi-channel protocols in the peer-to-peer mode, *Sphinx* can be generalized in a multi-channel environment to a hybrid network model that optimally divides the number of channels (from the pool of channels) to use in the cellular and peer-to-peer modes.

VI. Related Work and Conclusions

A. Related Work

In this section, we discuss some related work that incorporates peer-to-peer communication in conventional cellular networks to enhance the system performance. In [1], a proposal under the 3GPP organization, a new relaying protocol called "Opportunity Driven Multiple Access" (ODMA) is used to maintain high data rates at the boundaries of a cell. Relaying seeds or terminals are deployed to relay traffic for mobile stations in the low-data-rate area (boundary) of the cell via multi-hop transmissions. However, the ODMA protocol does not address and thus suffers from throughput degradation due to mobility and traffic locality, which in the worst case can nullify the advantage of high speed relaying, producing decreased end-to-end throughput. Notwithstanding, ODMA incorporates peer-to-peer relay specific channels (ORACH and ODCH) in the 3G UTRA/TDD frame structure, which can be potentially used by *Sphinx* to increase the network performance.

In [2], the authors propose a new wireless network model called "Multi-hop Cellular Network" (MCN). Briefly, the model involves mobile stations farther away from the base-station communicating with the base-station using a multi-hop path consisting of other mobile stations. In their *MCN-p* model, the transmission power of the base-station and mobile stations are reduced to achieve throughput increase and power reduction. Because the transmission range of the base-station does not cover the whole cell, MCN suffers from the disadvantages of a pure peer-to-peer model.

In [4], the authors propose a network model called "Integrated Cellular and Ad-hoc Relay" (iCAR) that involves special stations called ad-hoc relay stations (ARSs) to complement a conventional cellular network model. The ARSs are strategically placed between regular cells in the network to help relay data in a congested cell through a neighboring non-congested cell. Although iCAR enhances traffic load balancing between cells, it requires the underlying channel to be split between the regular base-stations and the ARSs, thus reducing capacity in cells where mobile stations do not require an ARS or cannot be served by an ARS. Unavailability of ARSs and heavily loaded neighboring cells further pose problems of peer-to-peer communications in terms of unfairness and traffic locality.

B. Summary

In this paper, we present a new hybrid network model called *Sphinx* for wireless packet data networks which combines the advantages of the cellular and peer-to-peer network models. It has the following properties: (i) higher throughput, and (ii) lower power consumption than conventional cellular networks and, (i) better fairness, and (ii) more resilience to mobility and traffic locality than peer-to-peer networks. We present simulation results which substantiate our claims that *Sphinx* performs better than the cellular and peer-to-peer network models.

References

[1] 3GPP TSG-RAN, "Opportunity Driven Multiple Access," *3GPP Technical Report, 3G TR 25.924*, v1.0.0, Dec. 1999.

[2] 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.

[3] T. Harrold and A. Nix, "Capacity Enhancement Using Intelligent Relaying for Future Personal Communication Systems," in *Proceedings of IEEE VTC Fall*, Boston, MA USA, Sep. 2000.

[4] C. Qiao and H. Wu, "iCAR: An Intelligent Cellular and Ad-hoc Relay System," in *Proceedings of IC3N*, Las Vegas, NV USA, Oct. 2000.

[5] G. Aggelou and R. Tafazolli, "On the Relaying Capacity of Next-Generation GSM Cellular Network," *IEEE Personal Communications Magazine*, vol. 8, no. 1, pp. 40-47, Feb. 2001.

[6] J. Broch et al, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," in *Proceedings of ACM MOBICOM*, Dallas, TX USA, Oct. 1998.

[7] H. Luo and S. Lu, "A Topology-Independent Fair Queueing Model in Ad Hoc Wireless Networks," in *Proceedings of IEEE ICNP*, Osaka, Japan, Nov. 2000.

[8] H.-Y. Hsieh and R. Sivakumar, "On Using the Ad-hoc Network Model in Cellular Packet Data Networks," in *Proceedings of ACM MOBIHOC*, Lausanne, Switzerland, June 2002.

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

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

[11] H.-Y. Hsieh and R. Sivakumar, "A Transport Layer Approach for Achieving Aggregate Bandwidths on Multi-homed Mobile Hosts," in *Proceedings of ACM MOBICOM*, Atlanta, GA USA, Sep. 2002.

[12] S.-L. Wu, Y.-C. Tseng, C.-Y. Lin, and J.-P. Sheu, "A Multi-channel MAC Protocol with Power Control for Multi-hop Mobile Ad Hoc Networks," *The Computer Journal*, vol. 45, no. 1, pp. 101-110, Jan. 2002.