

Study of the Capacity of Multihop Cellular Networks*

Antonis Panagakis, Elias Balafoutis, and Ioannis Stavrakakis

Dept. of Informatics and Telecommunications
University of Athens, Athens, Greece
{grad0260,balaf,istavrak}@di.uoa.gr

Abstract. Recently, the application of the peer to peer networking paradigm (typical for an ad hoc network) has been proposed for wireless local area networks (WLANs), instead of the traditional cellular networking paradigm. In this paper the performance of a WLAN employing the peer to peer networking paradigm is studied via simulations; the results indicate that the direct application of the peer to peer networking paradigm in a WLAN leads to a substantially decreased throughput for the traffic directed to the Access Point (AP). The study also reveals that the cumulative receiving throughput of nodes located at the periphery of relatively small circular areas around the AP is substantially higher. Thus, the capacity of the multihop cellular network may be enhanced by employing the peer to peer paradigm only outside a circular area around the AP and the cellular paradigm inside this circular area. Examples are provided of environments where the aforementioned idea of distributing the traditional AP functionality to a set of nodes at the periphery of a circular area around the AP can be effectively applied.

1 Introduction

The cellular networking paradigm is the traditional networking paradigm for WLANs. Recently, the application of the peer to peer networking paradigm (typical for an ad hoc network) has been proposed for WLANs ([7,9,6,2,3]). However, the direct application of the peer to peer networking paradigm in a WLAN demands that the AP be “downgraded” to an ordinary node, in the sense that the AP no longer has the control of the shared medium (as in typical cellular networks), and its transmission range is reduced. These can have important side-effects on the network’s performance, especially for the case of inter-cell traffic (traffic that traverses the AP). In this paper, the performance of a Multihop Cellular Network (MCN) ([6]) is studied via simulations for the case of inter-cell traffic and a more efficient approach for the employment of the peer to peer paradigm in WLANs is proposed. The main difference of a Multihop Cellular Network (MCN) compared to the cellular network is that peer

* This work has been supported in part by the IST Programme under contract IST-2001-32686 (BROADWAY).

to peer (between nodes) communications are allowed and the operation of the network is distributed. The transmission range of both the mobile nodes and the AP is reduced so that multi-pair peer to peer communication be possible unihibited by interference. On the other hand the transmission range should not be reduced to the extent that node connectivity be compromised¹. The benefits of the adaptation of the peer to peer paradigm within a WLAN include mainly reduced energy consumption and the possibility for multiple simultaneous transmissions over the shared medium (spatial reuse of the shared medium). In an effort to show the advantages of the MCN architecture due to the possibility for spatial reuse of the shared medium, past studies [6,2] have considered network traffic conditions in which all traffic is intra-cell. Under inter-cell traffic conditions though, the benefits of spatial reuse – if any – are questionable and the effectiveness of the MCN architecture is quite poor. In this paper the capacity of the MCN architecture is evaluated via simulations under inter-cell traffic conditions. The derived results indicate that the MCN architecture supports a substantially decreased throughput at the AP. The study also reveals that the cumulative receiving throughput of nodes located at the periphery of relatively small circular areas around the AP is substantially higher. Thus, the capacity of the multihop cellular network may be enhanced by employing the peer to peer paradigm only outside a circular area around the AP and the cellular paradigm inside this circular area in order to allow for a coordinated and high throughput (last) one-hop access to the AP of the nodes within this circular area.

The structure of this paper is as follows. In section 2 the capacity of the MCN is evaluated for the case of inter-cell traffic. In section 3 an enhancement is proposed which is shown to incorporate several of the advantages of the peer to peer communication into the cellular architecture. Related work is summarized in section 4.

2 Evaluation of the Capacity of the MCN Architecture

When all traffic is intra-cell, the MCN architecture is expected to be efficient due to the advantage of the spatial reuse of the shared medium in the peer to peer communications. No node is expected to be different than any other node in terms of the amount of traffic generated or received under the intra-cell traffic conditions. The uniform distribution of the peer to peer communications across the cell (a) maximizes the benefit of the spatial reuse by spreading evenly the one-hop transmissions (and minimizing the interference) and (b) results in a geographically evenly distributed traffic load avoiding the formation of problem-

¹ The reduced coverage area of the AP is referred to as the sub-cell; its radius is denoted as R_{sc} and is assumed to be equal to the transmission range of all the nodes. R_c denotes the radius of the geographical area that the AP is responsible for covering; in the cellular networking paradigm it is assumed that the transmission range of the AP is equal to R_c . In addition, neighboring APs are assumed of the same radius and not overlapping; that is, if L denotes the minimum distance between two APs then $L \geq 2R_c$.

atic bottlenecks; the latter occurs since a relatively low traffic load is directed to each receiving node and, thus, can be supported, even if (due to interference) its throughput is only a small portion of the medium's capacity.

However, in a WLAN architecture it is expected that the traffic would mostly be directed toward the node that supports accessing the web and nodes outside the cell. Thus, the long-term ratio of the inter-cell over the total (inter-cell and intra-cell) traffic is expected to be high and possibly very close to 1. The previous implies that the uniform distribution of the peer to peer transmission and traffic load are far from valid under realistic conditions, since one node, the AP, would be offered a high traffic load and have a high concentration of peer to peer transmissions in its vicinity. Due to the latter, the (receiving) throughput of the AP's shared medium would be drastically reduced and would be far from adequate in supporting the high traffic load offered to the AP (that would require a very high throughput). In this section the aforementioned issues are investigated by studying the performance of the MCN architecture under inter-cell traffic conditions.

2.1 Simulation Environment

The simulator used is the ns simulator [1] (version 2.1b7) with the multihop wireless extensions from the CMU Monarch group. At the medium access level the Distributed Coordination Function (DCF) mode of the IEEE 802.11 MAC standard is used with the default parameters of the simulator. The path loss model used is the two ray ground path loss model of the ns simulator according to which the signal attenuates in proportion to $1/d^2$ up to a certain distance (default value $\approx 86m$) and in proportion to $1/d^4$ beyond that. The only modifications applied concern the transmission power (in order to vary the transmission range) and the carrier sense threshold (CST) (in order to vary the range of interference). In the simulation results the ratio of the interference range over the transmission range is referred to as the interference factor and is denoted as a . For example, for the default values of the simulator $a \approx 2.2$ ². The UDP transport protocol is employed with a packet size of 512 bytes; the Distributed Source Routing (DSR) protocol is employed for packet routing. The nodes are assumed to be static so that any impact of node mobility on the derived results be avoided.

The simulated topology consists of an 11x11 (121-node) grid in which some (or all) nodes transmit to the node at the center of the grid, which plays the role of the AP. The side of the grid is 40m and the transmission range is set to 60m in order for the diagonal neighbors to be reachable in one hop. The investigation is limited to the case of the uplink (where the benefits in terms of power consumption due to the multihop operation are more important) and only inter-cell traffic is considered, that is, traffic directed to the AP.

² In this work the term interference range is used as in [5] to refer to the maximum distance for which a receiver can sense a transmitter; the term carrier sensing range (used for example in [4]) is rather more appropriate for this quantity (see [4] for a relevant discussion and [8] for details on the physical layer modeling in ns simulator).

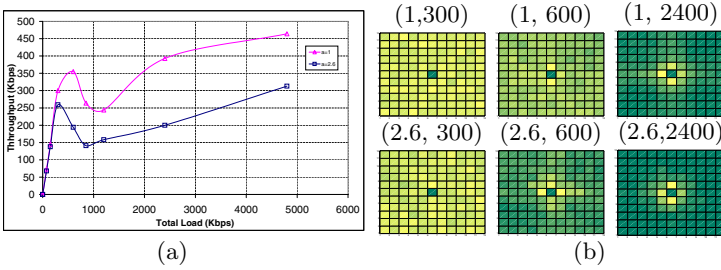


Fig. 1. (a) Total throughput vs total load for two different ranges of interference and (b) Throughput distribution for $(a, \text{total load (kbps)}) = (1, 300), (1, 600), (1, 2400), (2.6, 300), (2.6, 600), (2.6, 2400)$.

2.2 Simulation Results

The Effect of Interference. All nodes generate the same amount of traffic to the AP, which is assumed to be the node located in the center of the grid. Figure 1(a), illustrates the throughput achieved for different loads for the systems with $a = 2.6$ and $a = 1$, respectively. Two basic observations can be made from these results. The first is that the throughput of the MCN system is much lower than the channel capacity (1.3 Mbits)³. The second is that the throughput achieved for $a = 2.6$ is lower than that for $a = 1$. There is also a paradox observed in this figure regarding the shape of both throughput curves: initially, the throughput increases as the load increases, then decreases for a small range of loads and then increases again. Figure 1(b) attempts to provide for some insight into the aforementioned paradox. Each cell in each of the six grids shown in Figure 1(b) corresponds to a node in the simulated topology. The different colors in each cell represent the different ratios of packets sent by the source of the grid corresponding to that cell that have been successfully received at the destination. Only the traffic originated by each node is considered (packets that are relayed are not accounted for). The lighter the color, the larger the ratio of packets sent by the corresponding source that reach the AP. The very dark color at the center corresponds to the AP, which does not send any packets. In all figures the total load is evenly distributed among all sources. The first set of figures is obtained for a total load of 300Kbps. The received packet ratio at the destination for this load is 92% and 86% for $a = 1$ and $a = 2.6$, respectively; that is, the 92% and 86% of all packets reach the AP. As seen in the figure this ratio is almost evenly distributed among all nodes in the grid. The difference in throughput for the two systems is attributed to the level of interference, which is larger for $a = 2.6$. The second set of figures is obtained for total load of 600Kbps. The received packet ratio at the destination for this load is 60% and 32% for

³ The maximum throughput that can be achieved between two nodes using the IEEE 802.11 DCF, (with UDP traffic and packet size of 512 bytes) was found (via simulations) to be 1.3Mbits. This value may be considered as the channel capacity and is used for comparison with the throughput of the simulated system.

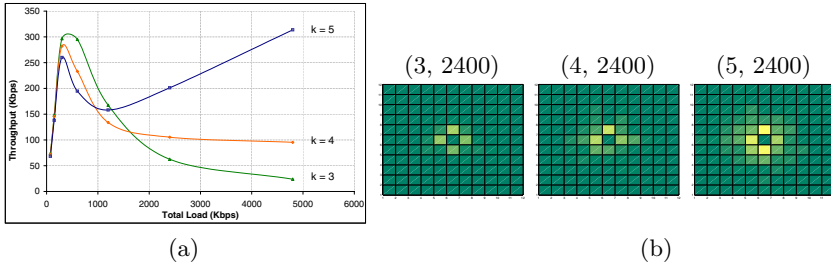


Fig. 2. (a) Total throughput vs total load for three different values of number of hops (3, 4 and 5) (b) Throughput distribution for (number of hops, total load (kbps))=(3, 2400), (4, 2400), (5, 2400).

$a = 1$ and $a = 2.6$, respectively. It is observed that mostly the nodes around the AP contribute to the total throughput, while the other nodes block each other's transmissions. However, this small fraction of nodes that do contribute to the system's throughput is rather low (5 Kbps per node) and the total throughput is low. In the last set of figures, which is obtained for total load of 2400Kbps, the situation is similar to the previous one. The received packet ratio for this load is 16% and 8% for $a = 1$ and $a = 2.6$, respectively. Again only the nodes around the AP contribute to the system's throughput but now their rate (20Kbps) is such that the total throughput becomes larger.

The Effect of the Number of Hops. In order to investigate the effect of the length of the path on the throughput only the nodes that lie k or less hops away from the AP transmit to the AP. Figure 2(a) illustrates the results for the system with interference factor $a = 2.6$ and for $k = 3$, $k = 4$ and $k = 5$. A general comment from Figure 2(a) is that the shape of the throughput curves remains the same for all values of k depicted in the figure; for light load situations (up to 300Kbps, which is far below the channel capacity), the impact of interference is negligible and the achieved throughput is about the same for the three cases ($k = 3, 4$ and 5). That is, the larger number of hops does not lead to a throughput reduction under light traffic load condition. When the interference becomes significant and a throughput reduction is observed, (at a load about 300Kbps in Figure 2(a)), the impact of the interference is larger on the larger hops paths and thus, the throughput decrease is larger for larger k . A "paradox" appears under heavy load conditions, where the aforementioned relative performance is reversed. Figure 2(b) attempts to shed some light into this behavior. Again, each cell in the figure corresponds to a node in the simulated topology. The different colors in each cell represent the different number of packets sent by each node that manage to reach the destination. The lighter the color, the larger the number of sent packets that reach the AP. This set of figures are derived for a system load of 2400Kbps (50, 30 and 20 Kbps per node for $k = 3, 4$ and 5 , respectively). By observing the lighter color for the nodes around the AP it can be concluded that despite the fact that the nodes send less packets (20 Kbps),




Fig. 3. Simulated topologies. The receivers are located at ring 1, ring 2 and ring 3 in (a), (b) and (c), respectively (,that is, 1, 2 and 3 hops away from the AP). Different colors are used to distinguish the different groups of nodes; in each group of nodes the node located nearer to the AP is the destination of the traffic originated by the other nodes belonging to the same group.




Fig. 4. Available bandwidth at the rings around the AP for $a = 2.6$.




Fig. 5. Available bandwidth at the rings around the AP for $a = 1$.

the system manages to deliver a larger number of packets when $k = 5$. It seems that for $k = 3$ and 4 the nodes block each other, while for $k = 5$, the nodes at the 5th hop block the nodes at the 4th and 3rd hop and leave space for the nodes at the first two hops to transmit to the AP. Similar observations can be made when comparing the cases for $k = 3$ and $k = 2$. Similar comments regarding the effect of the path length also apply for the case of $a = 1$, which is not presented here.

While the length of the paths does affect the performance of the system, the throughput remains significantly lower than the channel capacity even for a small number of hops around the AP and even for light loads. This fact highlights the inefficiency of the MCN architecture for the inter-cell traffic and confirms the related discussion in the beginning of this section.

Available Throughput at “Rings” around the AP. In an attempt to shed light into the observed performance degradation the throughput of the system is measured at rings of several (hops) distance from the access point. The ring located c hops away from the access point has m nodes where $(c, m) \in \{(1, 8), (2, 16), (3, 24), (4, 32), (5, 40)\}$ and is referred to as $ring_c$. In order to exclude the effect of the interference caused by the internal (with respect to each ring) nodes, these nodes do not transmit (or relay) any traffic. Thus, when the receiving nodes belong to $ring_c$, only nodes located $m, m > c$, hops away from




Fig. 6. The topology of a traditional cellular network may be decomposed into a doughnut-shaped topology and a small circular area around the AP.

the AP are considered as potential sources; the exact association of each source to one receiver is shown in Figure 3. Figures 4 and 5 illustrate the throughput for $a = 2.6$ and $a = 1$, respectively, for the cases where the receivers belong to $ring_c$, $c = 0, 1, 2, 3$. It can be observed that the throughput increases as c increases, which can be much larger than the capacity of the channel (1.3Mbps) due to spatial reuse.

The presented results indicate that although the achievable throughput relatively far from the AP can be much larger than the capacity of the channel, this throughput cannot be “transferred” to the AP due to the bottleneck formed around it. In the sequel an enhancement is proposed, which attempts to alleviate the AP bottleneck problem by exploiting the higher throughput achieved in the rings around it.

3 Enhancing the Efficiency of an MCN

Simulation results indicate that the achievable bandwidth at the AP for the inter-cell traffic in an MCN is low due to contention and interference. On the other hand, it is observed that the receiving throughput at nodes located at the periphery of a circular area around the AP is significantly higher. This is because when the destination nodes are nodes located on this periphery, the multihop peer to peer transmissions from nodes beyond this periphery benefit from spatial reuse and are more efficiently supported.

An enhancement could be based on the idea that the topology of a traditional cellular network be decomposed on a conceptual level into a small circular area around the AP and a doughnut-shaped topology. In a doughnut-shaped topology, where the traffic that would be directed to the AP in the traditional cellular environment is destined to nodes located on the inner boundary of the doughnut-shaped topology, it is expected that the multihop peer to peer transmissions would be effective, since spatial reuse of the shared medium would be possible along the entire path including the destination. In the real topology all traffic is destined to the AP and, thus, there would be the need to distribute (or pseudo-




Fig. 7. Proposed enhancement.

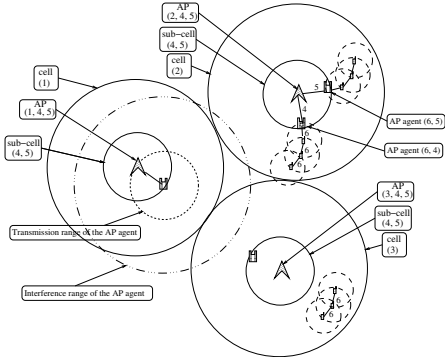


Fig. 8. An example of the generalization of the proposed enhancement for three neighboring cells.

distribute) the AP functionality at the inner boundary of the multihop doughnut-shaped area (see figure 6). To address the latter a dedicated channel may be used in order to pseudo-distribute the AP, that is, the AP would operate in a different channel than that used by the ordinary nodes for their peer to peer communications. In order for the communication between the nodes and the AP to be possible, some nodes, referred to as the AP agents, should operate on a time division basis in these two channels. Thus, at least two channels would be required in each cell. However, if one of them were reusable in adjacent cells (its use were not prohibited due to interference) the additional cost (in terms of required channels) would be rather low⁴.

The proposed enhancement is depicted in figure 7. The nodes located outside the sub-cell operate as in the simple MCN architecture in a channel (say ch_m). The difference is that the AP and the nodes located within its coverage area (the sub-cell) operate in a different channel (say ch_{AP}) by employing a centralized protocol, and thus, the AP keeps its distinguished role within the sub-cell. Some nodes within the sub-cell (the AP agents) operate on a time division basis between the two channels (ch_{AP} and ch_m) providing for connectivity between the AP and the nodes located outside the sub-cell. By employing the concept of the AP agents, and as far as the nodes located outside the sub-cell are concerned, the functionality of the AP is distributed to a set of nodes (the AP agents) that undertake the effort to concentrate the traffic directed to the AP in the multihop channel. Due to the fact that there is no interference between communications in and outside the sub-cell the maximum achievable throughput of the system (for inter-cell traffic) is expected to be greater than the throughput of the sim-

⁴ A potential enhancement such as the above requires that the AP agents be capable of operating on a time division basis between two channel. The exact description of such a mechanism, as well as the exact description of the mechanism used for determining which of the nodes should act as AP agents, is out of the scope of this paper.

ple MCN and approximately equal to the minimum between the capacity of the channel utilized within the sub-cell and the throughput of the system measured at the AP agents. The throughput of the system measured at the AP agents depends on the number and the location of the AP agents which in their turn depend on the transmission range of the AP.

The transmission range of the AP does not have to be the same as that of the other nodes, as in the simple MCN network. However, it should be noted that we restrict our presentation to the case where the dimensions of the sub-cell are such that the use of the channel utilized within the sub-cell does not prohibit (due to interference) the use of the same channel within the sub-cell of adjacent cells. It is straightforward to determine under which conditions this requirement is fulfilled. More specifically, and assuming that all cells are identical the use of a channel within the subcells of adjacent cells is permitted if (it is assumed that nodes located inside the sub-cell have a transmission range no greater than the transmission range of the AP) $L - 2R_{sc} \geq aR_{sc}$. Assuming that $L = 2R_c$ it follows that $\frac{R_c}{R_{sc}} \geq \frac{a}{2} + 1$ should hold. For typical cases (more precisely for the values of the interference ratio considered in this work ($a \leq 4$)) it is concluded that for $R_c \geq 3R_{sc}$ the use of an additional channel within the sub-cell does not affect adjacent cells. Thus, only the case of MCNs with at least three hops (and interference conditions no worse than those described in section 2.1) is considered in this work. It should be noted that the expansion of the sub-cell involves a tradeoff between the bandwidth that is available at the margins of the sub-cell (at the AP agents) and the total power consumption along of the paths between the nodes and the AP (since the last hop is longer).

The proposed enhancement may be generalized in order to combine the cellular and the peer to peer networking paradigms under the assumption that the AP is capable of operating simultaneously in multiple channels (at least two). This can be realized by adding additional antennas to the AP; the cost of such an approach is much less than the cost of adding additional APs since there are no extra infrastructure/administrative costs. In addition, a single channel is assumed outside the sub-cell for peer to peer communications. According to the proposed generalization one of the AP's operating channels is utilized exactly as in the cellular networking paradigm; it covers the entire cell and centralized control is employed by the AP. The rest of the channels are used within the sub-cell (whose dimensions are assumed to be subject to the same constraints as those discussed above, in order for the corresponding channels to be reusable in neighboring cells) and AP agents that operate on a time division basis between one of these channels and the channel used for multihop communications outside the sub-cell are employed.

The main difference compared to the MCN is that the peer to peer paradigm is not replacing the cellular one, but is used to form alternative multihop paths in order to offload the cellular system (by being utilized for intra-cell traffic) and provide power consuming alternative paths to the AP (for example, for traffic that can tolerate delay and losses) by exploiting the concept of the AP agents.

For the peer to peer communications the same channel can be utilized in all neighboring cells; interference occurring on cells' boundaries should not be of great concern since there always is available the cellular mode for nodes that face strong interference. Note that this does not hold in the case of MCN where the multihop path is the only one available. For the same reason (existence of the cellular mode), there is no need for the exact determination of the transmission range of the nodes in order for the peer to peer network to remain fully connected.

An example of the proposed generalization is illustrated in figure 8. Assume that there are three neighboring AP's, whose cells are of the same dimensions and do not overlap, that operate according to the cellular networking paradigm. That is, each AP uses a channel with such a transmission range that its entire cell is covered; let these channels be denoted as ch_1, ch_2, ch_3 . In order to double the capacity of the system three additional channels (say ch_4, ch_5, ch_6) would be required, since in each cell a different channel must be utilized. According to the proposed enhancement two of these channels (say ch_4 and ch_5) are employed by all three APs with a transmission range such that interference between nodes operating at these channels in different subcells is avoided, that is, the dimensions of the subcells corresponding to ch_4 and ch_5 are subject to the aforementioned restrictions. The other channel (ch_6) is employed in all three cells for the peer to peer communications between the nodes outside the sub-cell. In each sub-cell there are some AP agents for each channel utilized within the sub-cell, that is, there are some nodes that on a time division basis operate in (ch_4 and ch_6) or (ch_5 and ch_6). In this way the available capacity at the AP is tripled at the cost that the 2/3 of this capacity is not available with the same characteristics to all nodes (multihop paths are typically characterized by greater delay, lower availability but also lower power consumption). Given the diversity of the QoS requirements of the applications and the existence of (high demand) hot spots the proposed enhancement, which combines the cellular and the peer to peer networking paradigms, might be efficient in many environments.

4 Related Work

In [7] a cellular architecture is proposed in which some ad hoc relay stations are placed inside the cells and relay traffic dynamically from overloaded cells to lightly loaded cells in order to balance the system's traffic. In [9] a similar architecture is presented, in which an ad hoc overlay layer is added to the fixed cellular infrastructure for the forwarding of the traffic from a heavy loaded cell to neighboring cells. In [2] the multihop paradigm is adopted but with the coordination of a base station as in the cellular paradigm. The performance in this case is evaluated for intra-cell traffic and highlights the good spatial reuse characteristics of the peer to peer communication. In [6] the multihop cellular network architecture (MCN) is proposed and an analytical evaluation of the capacity of the network is presented. In [3] a simulation study of the MCN is provided for TCP traffic and in the presence of mobility, and three approaches to alleviate the bottleneck around the AP are presented. The effect of interference

in peer to peer communications is studied in [5] where the capacity of the ad-hoc networking paradigm is evaluated and scalability issues are addressed.

References

1. The network simulator ns-2, <http://www.isi.edu/nsnam/ns/>.
2. Hung-Yun Hsieh and Raghupathy Sivakumar. Performance comparison of cellular and multi-hop wireless networks: A quantitative study. In *Proc. of the ACM SIGMETRICS*, 2001.
3. Hung-Yun Hsieh and Raghupathy Sivakumar. On using the ad-hoc network model in wireless packet data networks. In *Proc. of the ACM MOBIHOC*, 2002.
4. Sang Bae Kaixin Xu, Mario Gerla. How effective is the ieee 802.11 rts/cts handshake in ad hoc networks? In *Proc. of the IEEE Globecom 2002*.
5. Jinyang Li, Charles Blake, Douglas S. J. De Couto, Hu Imm Lee, and Robert Morris. Capacity of ad hoc wireless networks. In *Proc. of ACM/IEEE MOBICOM*, pages 61–69, 2001.
6. Ying-Dar Jason Lin and Yu-Ching Hsu. Multihop cellular: A new architecture for wireless communications. In *Proc. of IEEE INFOCOM*, 2000.
7. Chunming Qiao and Hongyi Wu. icar: An integrated cellular and ad hoc relay system. In *Proc. of IC3N*, Las Vegas, NV, USA, October 2000.
8. Mineo Takai and Jay Martin. Effects of wireless physical layer modeling in mobile ad hoc networks. In *Proc. of the ACM MOBIHOC*, 2001.
9. Xiabin Wu, Gary S. H. Chan, and Biswanath Mukherjee. Madf: A novel approach to add an ad-hoc overlay on a fixed cellular infrastructure. In *Proc. of the IEEE WCNC*, 2000.