An Analysis of the Optimum Node Density for Ad hoc Mobile Networks

Elizabeth M. Royer*, P. Michael Melliar-Smith†, and Louise E. Moser†
* Department of Computer Science
† Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
eroyer@cs.ucsb.edu, {pmms,moser}@alpha.ece.ucsb.edu

Abstract—An ad hoc mobile network is a collection of nodes, each of which communicates over wireless channels and is capable of movement. Wireless nodes have the unique capability of transmission at different power levels. As the transmission power is varied, a tradeoff exists between the number of hops from source to destination and the overall bandwidth available to individual nodes. Because both battery life and channel bandwidth are limited resources in mobile networks, it is important to ascertain the effects different transmission powers have on the overall performance of the network. This paper explores the nature of this transmission power tradeoff in mobile networks to determine the optimum node density for delivering the maximum number of data packets. It is shown that there does not exist a global optimum density, but rather that, to achieve this maximum, the node density should increase as the rate of node movement increases.

I. Introduction

In 1978 Kleinrock and Silvester published their well-known paper "Optimum Transmission Radii for Packet Radio Networks" [6]. The paper provides an analytical analysis that explores the tradeoff between increased transmission radius, resulting in fewer hops to reach a destination, and the effective bandwidth lost at each node as a result of the increase in transmission range. The paper shows that the optimum number of neighbors for a given node is 5.89 (rounded to six), and concludes that a node's transmission radius should be adjusted so that it has six neighbors.

While this result may be valid for stationary networks, it does not consider the ramifications of node movement on the optimum transmission power. Mobile networking has existed since the early days of the DARPA packet radio network [5]; however, it has recently seen a large increase in popularity and usage due to improvements in both wireless devices and applications. These improvements include increased power and portability of wireless products and the development of applications suitable for those devices. With the increase in usage of mobile, wireless products has come a desire for communication between these devices, and hence a need for mobile routing protocols. As mobile networking becomes more popular, it is important to understand the characteristics of this type of communication so that users can communicate in an optimal manner without wasting battery life or bandwidth.

This paper examines the effects of transmission power on mobile network performance. Specifically, it explores what, if any, effect mobility has on the optimum transmission radius determined for stationary networks. While increasing the transmission radius does reduce the effective bandwidth seen at individual nodes, increasing the connectivity of the network may be more important as node mobility increases. To investigate this topic, the Ad hoc On-Demand Distance Vector (AODV) routing protocol [7] is used for route establishment. While it is likely that different routing protocols will have different saturation levels and route characteristics, the results obtained with AODV can be generalized to most on-demand ad hoc routing protocols. The remainder of this paper is organized as follows. Section II describes the basic mechanism of AODV's unicast routing. Section III then describes the simulations performed. Section III-A provides an in-depth description of the new mobility model used, and Section III-B presents the results from the investigation. Section IV discusses related work, and Section V concludes the paper.

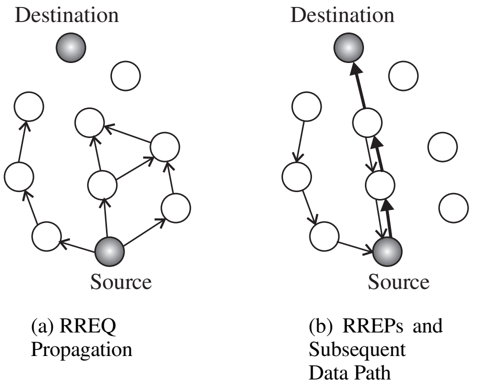

Fig. 1. Route Discovery.

II. Ad hoc On-Demand Distance Vector Routing

The Ad hoc On-Demand Distance Vector (AODV) routing protocol is a reactive protocol designed for use in ad hoc mobile networks. AODV initiates route discovery whenever a source node needs a route, and it maintains this route as long as it is needed by the source. Each node maintains a monotonically increasing sequence number that is incremented whenever there is a change in the local connectivity information for the node. These sequence numbers ensure that the routes are loop-free.

A. Route Discovery

Route discovery follows a route request/route reply query cycle. When a source node needs a route to a destination, it broadcasts a *Route Request* (RREQ) packet across the network, and then sets a timer to wait for the reception of a reply. This RREQ contains the IP address of the destination and last known sequence number for that destination. Nodes receiving the packet can respond to the RREQ either if they are the destination, or if they have an unexpired route to the destination whose corresponding sequence number is at least as great as that contained in the RREQ. If either of these two conditions are satisfied, the node responds by unicasting a *Route Reply* (RREP) back to the source node. Otherwise, the node rebroadcasts the RREQ, as indicated in Figure 1(a). Additionally, each node receiving the RREQ creates a *reverse route* entry for the source node in its route table.

As the RREP is forwarded to the source node, intermediate nodes that receive the RREP create a *forward route* entry for the destination in their route tables before transmitting the RREP to the next hop. Figure 1(b) illustrates the propagation of the RREPs back to the source node. Once the source node receives a RREP, it can begin using the route to send data packets, as shown by the darker arrows in Figure 1(b).

If the source node does not receive a RREP before its discovery timer expires, it rebroadcasts the RREQ. It attempts discovery up to some maximum number of attempts. If it does not discover a route after this maximum number of tries, the session is aborted.

B. Route Maintenance

AODV also provides a *Route Error* (RERR) message for notifying neighboring nodes of link breaks. A RERR message only needs to be sent when a link break occurs in an active route. An active route is a route that has recently been used to send data packets. When a link break in such an active route occurs, the node upstream of the break expires the route entry of each destination which is now unreachable. It then determines if any of its neighbors route through it in order to reach the invalidated destinations. If there exists one or more such neighbors, it broadcasts a RERR message. The RERR message contains a list of each destination which has become unreachable due to the link break. It also contains the last known sequence number for each listed destination, incremented by one.

When a neighboring node receives the message, it expires any routes to the listed destinations that use the source of the RERR message as the next hop. Then, if the node has a record of one or more nodes that route through it to reach one of the destinations, it rebroadcasts the message. Once a source node receives the RERR, it invalidates the listed routes as described. It can then reinitiate route discovery if it still requires the route.

III. Simulations

The simulations have been performed using the GloMoSim Network Simulator developed at UCLA [1]. This simulator models the OSI seven layer network architecture and includes models of IP routing and UDP. The simulator also provices network node mobility, thereby enabling simulation of mobile ad hoc networks.

The MAC layer protocol used in the simulations is the IEEE standard 802.11 Distributed Coordination Function (DCF) [3]. This standard uses Request-To-Send (RTS) and Clear-To-Send (CTS) control packets for unicast data transmissions between neighboring nodes. A node wishing to unicast a data packet to its neighbor broadcasts a short RTS control packet. When its neighbor receives the packet, it responds with a CTS packet. Once the source node receives the CTS, it transmits the data packet. After receiving this data packet, the destination then sends an acknowledgment (ACK) to the source, signifying reception of the data packet. The use of the RTS-CTS control packets reduces the potential for the hidden-terminal problem [10]. Broadcast data packets and RTS control packets are sent using the unslotted Carrier Sense Multiple Access protocol with Collision Avoidance (CSMA/CA) [3]. When a node wishes to broadcast a packet, it first senses the channel. If it does not detect an on-going transmission, it broadcasts its packet. On the other hand, if it does detect a transmission, it calculates a backoff time and then waits this amount of time before reattempting the transmission.

The propagation model used is the free space model with threshold cutoff included in the GloMoSim simulation package. The free space model has a power signal attenuation of $1/d^2$, where d is the distance between nodes. The radio model used also has capture capability, where it can lock on to a strong signal during interference from other signals, and still receive the packet. Other interfering packets with weaker signal strength are dropped. The data rate for the simulations is 2 Mb/sec.

Four different node mobilities between 0m/s and 10m/s are modeled. The average number of neighbors in each simulation is varied by adjusting the transmission range. Each mobility/transmission range combination is run for 10 different initial network configurations, and the results are averaged to produce the data points. Each simulation simulates 240 seconds and models a network of 100 nodes in a 1000m$\times$1000m area. Because network saturation must be achieved to determine the maximum throughput of the network, the number of sources is set at 40, and each source sends twelve 512-byte data packets per second. This combination of send rate and number of sources results in network saturation.

AODV does not guarantee packet delivery; however, it does find good routes for IP's best-effort delivery. Because data packets are not buffered for retransmission, these packets are likely to be lost. If a collision involving a data packet occurs at a node and the packet cannot be captured, the packet is lost. Additionally, if a link in an active route breaks, the source node attempts to repair the link by performing at most three route discoveries. If it does not receive a RREP after this number of attempts, it assumes the destination has become unreachable. It therefore aborts the session and discontinues the transmission of data packets for that session. For this reason, it is important to examine the number of packets that the destination nodes receive, instead of the ratio of the number of packets received to the number of packets sent (packet delivery ratio). The packet delivery ratio for some simulations may be artificially high if a large number of sessions are aborted due to lack of a route.

A. The Mobility Model

The mobility model originally used for the simulations was the Random Waypoint model [2]. At the beginning of a simulation, the Random Waypoint mobility model randomly places the nodes within the predefined simulation area. Each node then selects a destination within that area and a speed from a uniform

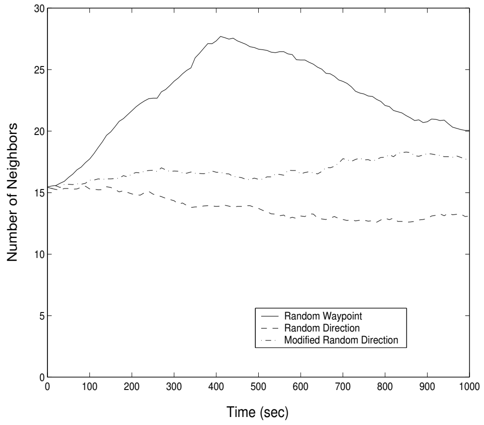

Fig. 2. Average Neighbors per Node at 1 m/s Mobility.

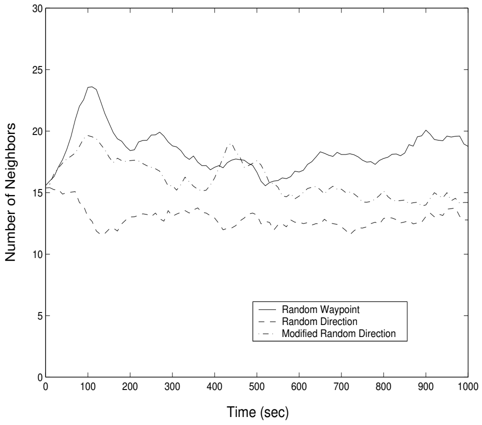

Fig. 3. Average Neighbors per Node at 5 m/s Mobility.

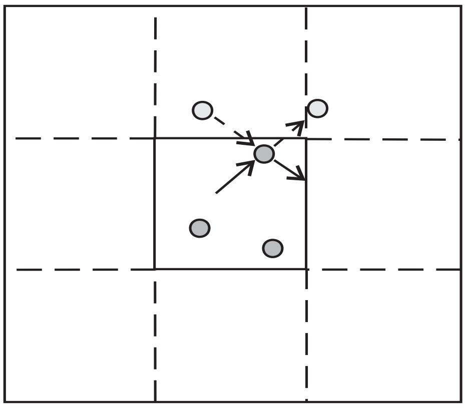

Fig. 4. Simulated Movement in the Random Direction Mobility Model.

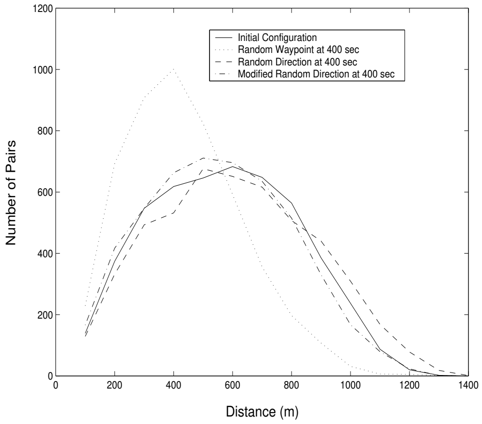

Fig. 5. Number of Node Pairs at a Given Distance.

distribution of user-specified speeds. The node then travels to its selected destination at the selected speed. Once it reaches the destination, it is stationary for some pre-defined pause time. At the end of the pause time, it selects a new destination and speed combination, and then resumes movement. This model causes continuous changes in the topology of the network.

While running the simulations, an interesting behavior of the mobility model was noticed. The average number of neighbors seen at a given node periodically increases and decreases as the simulation progresses, where the frequency of the change is relative to the speed of the nodes. The solid lines in Figures 2 and 3 illustrate the average number of neighbors per node in a 100 node network during 1000 seconds of simulated time. This fluctuation is due to the inherent characteristics of the mobility model. Because a node must select a destination in the simulation area, the node is most likely to travel in the direction in which there are the most destinations from which to choose. This predisposes nodes to choose destinations that are either in the middle of the simulation area, or that they reach by traveling through the middle. This characteristic creates a situation in which nodes converge in the center of the area, and then disperse, and then re-converge, etc., resulting in *density waves*.

Because a fairly constant number of neighbors per node was necessary throughout the simulation, we developed our own mobility model, called the Random Direction model. Instead of selecting a destination within the area, the nodes select a *direction* in which to travel, where a direction is measured in degrees. At the beginning of the simulation, each node selects a degree between 0 and 359, and then finds a destination on the boundary in this direction of travel. It then selects a speed, and travels to that destination at the given speed. Once it reaches the destination, it rests for the given pause time, and then selects a new degree between 0 and 180. The direction is limited because the node is already on the boundary, and nodes do not pass through the boundary. Thus the degree determines a direction relative to the wall of the boundary area on which the node is located. The node then identifies the destination on the boundary in this line of direction, selects a new speed, and resumes travel. This model can be thought of as examining a microcell of a larger area, as illustrated in Figure 4. A node travels to the boundary and then selects a new direction, which is similar to a node that continues past the boundary, while a new node enters the microcell at that location traveling in a different direction. The average number of neighbors per node for this model is illustrated in Figures 2 and 3. As is evident from the figures, the Random Direction mobility model causes many fewer fluctuations in the node distribution than the Random Waypoint model. Also in these figures is the average number of neighbors per node using a variation

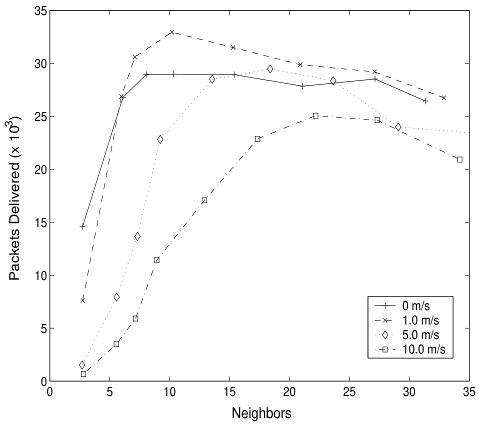

Fig. 6. Number of Packets Delivered.

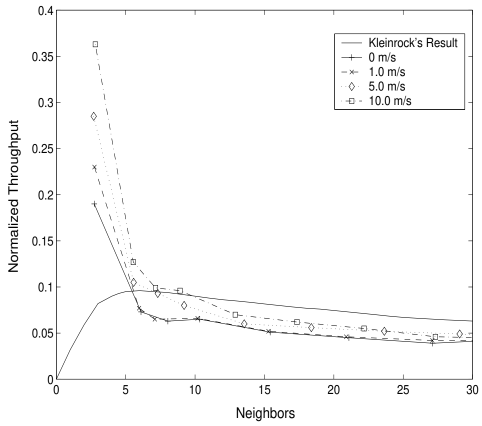

Fig. 7. Normalized Throughput.

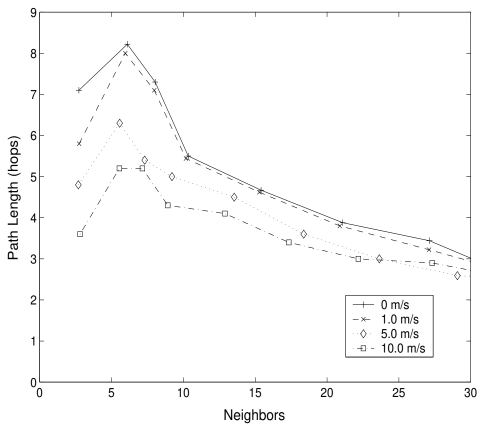

Fig. 8. Average Path Length.

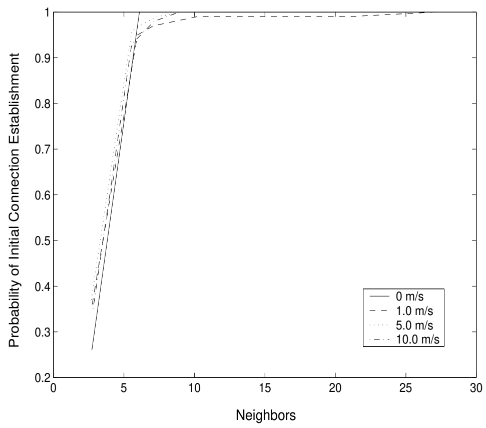

Fig. 9. Probability of Being Able to Establish Initial Route.

of the Random Direction model, called the Modified Random Direction model. In this model, nodes select a direction degree as before, but they may choose their destination anywhere along that direction of travel. They do not need to travel all the way to the boundary.

Figure 5 illustrates the average number of pairs of nodes with the given distance of separation. The figure represents nodes moving at 1 m/s. Initially, the nodes are evenly distributed throughout the simulation area. As the simulation progresses, the characteristics of the Random Waypoint model show that after 400 seconds of simulation time, the node density is much higher than the initial configuration. However, the Random Direction model is able to maintain a node density that is similar to the initial node distribution.

B. Results

The number of data packets received by their respective destination is shown in Figure 6. The figure shows that for small transmission radius and low connectivity, few data packets are delivered due to lack of a route. As the connectivity increases, however, the number of packets delivered rapidly increases, until the curves level off. As the transmission power continues to increase, the network becomes saturated, and consequently there is a gradual decline in the number of data packets delivered. Network saturation results in an increased number of collisions, as well as reduced channel access which leads to buffer overflows. There does not appear to be a global optimal number of neighbors for all mobilities. For 0m/s mobility, the results show this optimum to be around seven or eight neighbors per node, which differs only slightly from what Kleinrock proved for a stationary network [6]. As mobility increases, however, the optimum shifts to higher connectivity. The faster nodes move, the more frequently link breaks occur. Hence, even though the effective bandwidth seen at individual nodes suffers due to the increased transmission power and collisions, the number of packets delivered still increases, relative to shorter transmission ranges. This is because link breaks are less frequent and routes are maintained for longer periods of time.

Figure 7 represents Kleinrock's normalized throughput achieved by the simulations, and compares these values to an approximation derived by Kleinrock in Figure 3 of his paper [6]. The network throughput is given by the number of successful transmissions divided by the average path length,

$$\gamma = \frac{n}{Ne\overline{h}}$$

where n is the total number of nodes, N is the average node degree, e is the natural logarithm, and $\overline{h}$ is the average path length. To normalize, the value $\frac{\gamma}{\sqrt{n}}$ is plotted.

The normalized throughput seen by the individual nodes in the simulation differs greatly from the theoretical value, particularly for a small number of neighbors. In Kleinrock's paper, it is assumed that, as the transmission power decreases and the number of neighbors similarly decreases, connectivity in the network is still maintained. The throughput suffers as a result of the increased number of hops to reach the destination. However, the simulations show that for very low node densities, the network does not, in fact, remain connected. Numerous nodes or groups of nodes become disconnected from the node majority. The result of the disconnected operation is that many of the sessions abort because routes to the destination are unavailable. The only sessions that are able to complete are those with a small path length. Hence, the normalized throughput appears artificially high. Because of this effect, the utility of this metric for mobile networks is doubtful.

Figure 8 illustrates the average path length of active routes. It verifies that the lower transmission powers do have shorter path lengths, because those are the only routes that are able to be completed. Once the network is fully connected, the path length increases, and then decreases once again as the transmission power increases and fewer hops are needed to connect the source and the destination.

The probability of being able to establish an initial route between a source and destination is shown in Figure 9. For the sparsely connected network, the probability is fairly low as the network suffers from partitioning. As the node density increases, however, the probability rapidly increases for all mobilities, until it eventually stablizes at one.

IV. Related Work

Sanchez, Manzoni and Haas performed a related study in [9]. In this work, they investigate the critical transmission range in ad hoc networks, which is the minimum transmission range of the transceivers that is required to achieve full network connectivity. They present an algorithm to calculate this minimum transmission range, and then they study the effect of mobility on that value.

Two related papers study the problem of adjusting the transmission power in order to find a balance between the achieved throughput and power consumption. In [8], the authors present algorithms which adaptively adjust the transmission power of the nodes in response to topological changes, with the goals of maintaining a connected network while using minimum power. Through simulation, they show that an increase in throughput, together with a decrease in power consumption, can be achieved by managing the transmission levels of the individual nodes.

Similarly, the work presented in [4] presents a study on power management in ad hoc networks by determining the effects of transmission power on network throughput and power consumption. Nodes are grouped into clusters, where nodes within a cluster adapt their power levels to reach at most the farthest node within the cluster. The authors either allow or disallow nodes to dynamically adjust their power levels per transmission to be the maximum amount of power needed to reach the destination node within the cluster. The authors show that both methods reduce power consumption while increasing throughput, with the former scheme achieving greater performance improvements.

V. Conclusions

As mobile networking gains popularity, it is important to understand the characteristics of these networks so that they can be tuned to achieve optimum performance. The transmission power of wireless nodes is a key component in determining the interconnection pattern of the network. For wireless transmission, a tradeoff exists between decreasing the number of hops between sources and destinations and decreasing the effective bandwidth available to individual network nodes.

It has been shown that as the mobility speed of nodes increases, it is desirable to increase transmission power in order to achieve delivery of data packets to their destinations. Moreover, there does not exist a global optimum connectivity level for varying node mobilities. For a stationary network, the optimum connectivity is seven or eight neighbors per node, which is similar to Kleinrock's conclusion of six neighbors per node. As the mobility of the nodes increases, however, an increase in the transmission power of the nodes results in a larger percentage of data packets reaching their destination, and hence a higher network throughput.

While it is likely that different reactive ad hoc routing algorithms will have different optimal connectivity levels, the results presented here can be generalized for any such algorithm. The simulations performed for this study assume an open air, free space environment. In the physical world, terrain and atmospheric conditions exist which effect network connectivity. Further study is needed to determine the effects of these variables on the optimal node transmission levels.

References

[1] L. Bajaj, M. Takai, R. Ahuja, K. Tang, R. Bagrodia, and M. Gerla. GloMoSim: A Scalable Network Simulation Environment. Technical Report CSD Technical Report, #990027, UCLA, 1997.

[2] J. Broch, D. A. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva. A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols. *Proceedings of the* 4^{th} *Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 85–97, Dallas, Texas, October 1998.

[3] IEEE Standards Department. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Standard 802.11–1997*, 1994.

[4] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao. Power Management for Throughput Enhancement in Wireless Ad hoc Networks. *Proceedings of the IEEE International Conference on Communications (ICC)*, pages 1503–1513, New Orleans, LA, June 2000.

[5] J. Jubin and J. Tornow. The DARPA Packet Radio Network Protocols. *Proceedings of the IEEE*, 75(1):21–32, 1987.

[6] L. Kleinrock and J. Silvester. Optimum Transmission Radii for Packet Radio Networks or Why Six is a Magic Number. *Proceedings of the IEEE National Telecommunications Conference*, pages 4.3.1–4.3.5, Birmingham, Alabama, December 1978.

[7] C. E. Perkins and E. M. Royer. The Ad hoc On-Demand Distance Vector Protocol. C. E. Perkins, editor, *Ad hoc Networking*, pages 173–219. Addison-Wesley, 2000.

[8] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks using Transmit Power Adjustment. *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, pages 404–413, Tel Aviv, Israel, March 2000.

[9] M. Sanchez, P. Manzoni, and Z. J. Haas. Determination of Critical Transmission Range in Ad hoc Networks. *Proceedings of the Multiaccess, Mobility and Teletraffic for Wireless Communications (MMT) Conference)*, Venice, Italy, October 1999.

[10] F. A. Tobagi and L. Kleinrock. Packet Switching in Radio Channels: Part-II - The Hidden Terminal Problem in Carrier Sense MultipleAccess Models and the BusyTone Solution. *IEEE Transactions on Communications*, 23(12):1417–1433–20, December 1975.