

The Broadcast Storm Problem in a Mobile Ad Hoc Network *

Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu

Department of Computer Science and Information Engineering
National Central University
Chung-Li, 32054, Taiwan
Tel: 886-3-4227151 ext. 4512, Fax: 886-3-4222681
Email: yctseng@csie.ncu.edu.tw
URL: <http://www.csie.ncu.edu.tw/~yctseng/>

Abstract

Broadcasting is a common operation in a network to resolve many issues. In a mobile ad hoc network (MANET) in particular, due to host mobility, such operations are expected to be executed more frequently (such as finding a route to a particular host, paging a particular host, and sending an alarm signal). Because radio signals are likely to overlap with others in a geographical area, a straightforward broadcasting by flooding is usually very costly and will result in serious redundancy, contention, and collision, to which we refer as the *broadcast storm* problem. In this paper, we identify this problem by showing how serious it is through analyses and simulations. We propose several schemes to reduce redundant rebroadcasts and differentiate timing of rebroadcasts to alleviate this problem. Simulation results are presented, which show different levels of improvement over the basic flooding approach.

Keywords: broadcast, communication, mobile ad hoc network (MANET), mobile computing, wireless network.

1 Introduction

The advancement in wireless communication and economical, portable computing devices have made mobile computing possible. One research issue that has attracted a lot of attention recently is the design of mobile ad hoc network (MANET). A MANET is one consisting of a set of mobile hosts that may communicate with one another and roam around at their will. No base stations are supported in such

*This work is supported by the National Science Council of the Republic of China under Grants #NSC88-2213-E-008-013 and #NSC88-2213-E-008-014.

an environment. Due to considerations such as radio power limitation, channel utilization, and power-saving concerns, a mobile host may not be able to communicate directly with other hosts in a *single-hop* fashion. In this case, a *multihop* scenario occurs, where the packets sent by the source host are relayed by several intermediate hosts before reaching the destination host.

Applications of MANETs occur in situations like battlefields or major disaster areas where networks need to be deployed immediately but base stations or fixed network infrastructures are not available. Unicast routing in MANET has been studied in several articles [6, 7, 14, 15, 23]. A working group called “manet” has been formed by the Internet Engineering Task Force (IETF) to study the related issues and stimulate research in MANET [21].

This paper studies the problem of sending a broadcast message in a MANET. Broadcasting is a common operation in many applications, e.g., graph-related problems and distributed computing problems. It is also widely used to resolve many network layer problems. In a MANET in particular, due to host mobility, broadcastings are expected to be performed more frequently (e.g., for paging a particular host, sending an alarm signal, and finding a route to a particular host [6, 14, 15, 23]). Broadcasting may also be used in LAN emulation [2] or serve as a last resort to provide multicast services in networks with rapid changing topologies.

In this paper, we assume that mobile hosts in the MANET share a single common channel with *carrier sense multiple access* (CSMA), but no collision detection (CD), capability. Synchronization in such a network with mobility is unlikely, and global network topology information is unavailable to facilitate the scheduling of a broadcast. So one straightforward and obvious solution is broadcasting by *flooding*. Unfortunately, in this paper we observe that serious redundancy, contention, and collision could exist if flooding is done blindly. First, because the radio propagation is omnidirectional and a physical location may be covered by the transmission ranges of several hosts, many rebroadcasts are considered to be redundant. Second, heavy contention could exist because rebroadcasting hosts are probably close to each

To appear in the fifth annual International Conference on Mobile Computing and Networking (MobiCom'99), Seattle, Washington, Aug. 15-20 1999.

other. Third, collisions are more likely to occur because the RTS/CTS dialogue is inapplicable and the timing of rebroadcasts is highly correlated.

Collectively, we refer to these problems associated with flooding as the *broadcast storm* problem. Through analyses and simulations, we demonstrate how serious the storm is. Two directions to alleviate this problem is to reduce the possibility of redundant rebroadcasts and differentiate the timing of rebroadcasts. Following these directions, we develop several schemes, called *probabilistic*, *counter-based*, *distance-based*, *location-based*, and *cluster-based* schemes, to facilitate MANET broadcasting. Simulation results are presented to study the effectiveness of these schemes.

To the best of our knowledge, the broadcast storm problem has not been addressed in depth for MANET before. It is however worth of summarizing some results for broadcasting that are for other environments. Works in [3, 4, 9, 10, 11, 20] assume a packet-radio network environment. Most of these results rely on *time division multiple access* (TDMA, which requires timing synchronization) and certain levels of topology information. Their goal is to find a slot assignment. Obtaining an optimal assignment has been shown to be NP-hard [9]. The broadcast scheduling problem studied in [8, 16, 24, 25], although carries a similar name, is not intended to solve the problem addressed in this paper. Its goal is to assign a contention-free time slot to each radio station.

The rest of this paper is organized as follows. Section 2 defines and analyzes the broadcast storm problem. Mechanisms to alleviate the storm are proposed in Sections 3. Simulation results are in Section 4 and conclusions are drawn in Section 5.

2 Preliminaries

2.1 Broadcasting in a MANET

A MANET consists of a set of mobile hosts that may communicate with one another from time to time. No base stations are supported. Each host is equipped with a CSMA/CA (*carrier sense multiple access with collision avoidance*) [19] transceiver. In such environment, a host may communicate with another directly or indirectly. In the latter case, a *multihop* scenario occurs, where the packets originated from the source host are relayed by several intermediate hosts before reaching the destination.

The *broadcast problem* refers to the sending of a message to other hosts in the network. The problem considered here has the following characteristics.

- **The broadcast is spontaneous.** Any mobile host can issue a broadcast operation at any time. For reasons such as the host mobility and the lack of synchronization, preparing any kind of global topology knowledge is prohibitive (in fact this is at least as hard as the broadcast problem). Little or no local information may be collected in advance.

- **The broadcast is unreliable.**¹ No acknowledgement mechanism will be used.² However, attempt should be made to distribute a broadcast message to as many hosts as possible without paying too much effort. The motivations to make such an assumption are (i) a host may miss a broadcast message because it is off-line, it is temporarily isolated from the network, or it experiences repetitive collisions, (ii) acknowledgements may cause serious medium contention (and thus another “storm”) surrounding the sender, and (iii) in many applications (e.g., the route discovery in [6, 14, 15, 23]), a 100% reliable broadcast is unnecessary.

In addition, we assume that a host can detect duplicate broadcast messages. This is essential to prevent endless flooding of a message. One way to do so is to associate with each broadcast message a tuple (source ID, sequence number) as that in [6, 23].

Finally, we comment that we do not confine ourselves to the broadcasting of the **same** message.³ In this paper we focus on the flooding behavior in MANET — the phenomenon where the transmission of a packet will trigger other surrounding hosts to transmit the same (or modified) packet. We shall show that if flooding is used blindly, many redundant messages will be sent and serious contention/collision will be incurred. Our goal is to solve broadcast with efficiency in mind.

2.2 Broadcast Storm Caused by Flooding

A straight-forward approach to perform broadcast is by *flooding*. A host, on receiving a broadcast message for the first time, has the obligation to rebroadcast the message. Clearly, this costs n transmissions in a network of n hosts. In a CSMA/CA network, drawbacks of flooding include:

- **Redundant rebroadcasts:** When a mobile host decides to rebroadcast a broadcast message to its neighbors, all its neighbors already have the message.
- **Contention:** After a mobile host broadcasts a message, if many of its neighbors decide to rebroadcast the message, these transmissions (which are all from nearby hosts) may severely contend with each other.
- **Collision:** Because of the deficiency of backoff mechanism, the lack of RTS/CTS dialogue, and the absence of CD, collisions are more likely to occur and cause more damage.

¹A more strict one is *reliable broadcast* [1, 22], whose goal is to ensure all hosts receive the message. High-level acknowledgements between hosts are exchanged. Such protocols are typically accomplished at the application layer and is out of the scope of this paper. However, the result in this paper may serve as an underlying facility to implement reliable broadcast.

²The MAC specification in IEEE 802.11 [19] does not request acknowledgement on receipt of broadcast packets.

³For instance, the routing protocols in [6, 14, 15, 23] rely on broadcasting a UDP packet called *route request* to search for a route from a source to a particular destination. When propagating such a request, a host generally appends its ID to the message so that appropriate routing information can be collected.




Figure 1: Two optimal broadcasting schedules in MANETs. Connectivity between hosts is represented by links. White nodes are source hosts, and gray nodes are relay hosts.

Collectively, we refer to the above phenomena as the *broadcast storm problem*. The following discussion shows how serious the storm is through analyses.

2.2.1 Analysis on Redundant Rebroadcasts

We first use two examples to demonstrate how much redundancy could be generated. In Fig. 1(a), it only takes two transmissions for the white node to broadcast a message, whereas four transmissions will be carried out if no attempt is made to reduce redundancy. Fig. 1(b) shows an even serious scenario: only two transmissions are sufficient to complete a broadcast as opposed to seven transmissions caused by flooding.

The following analysis shows that rebroadcasts are very costly and should be used with caution. Consider the simple scenario in Fig. 2, where host **A** sends a broadcast message, and host **B** decides to rebroadcast the message. Let S_A and S_B denote the circle areas covered by **A**'s and **B**'s transmissions, respectively. The additional area that can benefit from **B**'s rebroadcast is the shaded region, denoted as S_{B-A} . Let r be the radii of S_A and S_B , and d the distance between **A** and **B**. We can derive that $|S_{B-A}| = |S_B| - |S_{A \cap B}| = \pi r^2 - \text{INTC}(d)$, where $\text{INTC}(d)$ is the intersection area of the two circles centered at two points distanced by d ,

$$\text{INTC}(d) = 4 \int_{d/2}^r \sqrt{r^2 - x^2} dx.$$

When $d = r$, the coverage area $|S_{B-A}|$ is the largest, which equals $\pi r^2 - \text{INTC}(r) = r^2(\frac{\pi}{3} + \frac{\sqrt{3}}{2}) \approx 0.61\pi r^2$. This shows a surprising fact that a rebroadcast can provide only 0 ~ 61% additional coverage over that already covered by the previous transmission.

Also, we would like to know the average value of $\pi r^2 - \text{INTC}(d)$. Supposing that **B** can randomly locate in any of **A**'s transmission range, the average value can be obtained by integrating the above value over the circle of radius x cen-




Figure 2: Analysis on the extra area that can benefit from a rebroadcast: **A** sends a broadcast packet and **B** decides to rebroadcasts the packet.




Figure 3: Analysis on redundancy: the expected additional coverage $EAC(k)$ (divided by πr^2) after a host heard a broadcast message k times.

tered at **A** for x in $[0, r]$:

$$\int_0^r \frac{2\pi x \cdot [\pi r^2 - \text{INTC}(x)]}{\pi r^2} dx \approx 0.41\pi r^2.$$

Thus, after the previous broadcast, a rebroadcast can cover only additional 41% area in average.

Now consider the scenario of having received a broadcast message twice: if host **C** decides to rebroadcast after it heard **A**'s and **B**'s broadcasts. The area that can benefit from **C**'s rebroadcast is $S_{C-(A \cup B)}$. Through simulations (by randomly generating **A** and **B** on **C**'s transmission range with grid estimation), we found that in average $|S_{C-(A \cup B)}| \approx 0.19\pi r^2$. This shows an even dimmer prospect of hoping rebroadcasts to propagate the message to new hosts.

In general, we would like to know the benefit of a host rebroadcasting a message after heard the message k times. The result can be easily obtained from simulation by randomly generating k hosts in a host X 's transmission range and calculating the area covered by X excluding those already covered by the other k hosts. Denote this value by $EAC(k)$ (EAC stands for *expected additional coverage*). Fig. 3 shows our simulation result. As can be seen, when $k \geq 4$, the expected additional coverage is below 0.05%.




Figure 4: Analysis on contention: the probabilities of having k contention-free hosts among n receiving hosts.

2.2.2 Analysis on Contention

To address the contention problem, consider the situation where host **A** transmits a broadcast message and there are n hosts hearing this message. If all these hosts try to rebroadcast the message, contention may occur because two or more hosts around **A** are likely to be close and thus contend with each other on the wireless medium.

Let's analyze the simpler case of $n = 2$. Let hosts **B** and **C** be the two receiving hosts. Let **B** randomly locate at **A**'s transmission range. In order for **C** to contend with **B**, it must locate in the area $S_{A \cap B}$. So the probability of contention is $|S_{A \cap B}|/\pi r^2$. Let x be the distance between **A** and **B**. Integrating the above formula over the circle of radius x from 0 to r , the expected probability of contention is

$$\int_0^r \frac{2\pi x \cdot \text{INTC}(x)/(\pi r^2)}{\pi r^2} dx \approx 59\%.$$

Clearly, the contention is expected to be higher as n increases. We derived a simulation by randomly generating n hosts in **A**'s transmission range. We observe the probability $cf(n,k)$ that k hosts among these n hosts experience no contention in their rebroadcasting (cf stands for *contention-free*). The results are shown in Fig. 4. We can see that the probability of all n hosts experiencing contention (i.e., $cf(n,0)$) increases quickly over 0.8 as $n \geq 6$. So the more crowded the area is, the more serious the contention is. On the other hand, the probability of having one contention-free host (i.e., $cf(n,1)$) drops sharply as n increases. Further, it is very unlikely to have more contention-free hosts (i.e., $cf(n,k)$ with $k \geq 2$). Note that having $k = n - 1$ contention-free hosts implies having n such hosts, so $cf(n, n - 1) = 0$.

2.2.3 Analysis on Collision

In a MANET, there is no base station or access point. Therefore, in this paper we exclude the use of the *point coordinate function* (PCF) described in the IEEE 802.11 MAC specification [19], and study mainly the behavior under the *distributed coordinate function* (DCF).

The CSMA/CA mechanism requires a host to start a *back-off* procedure right after the host transmitted a message, or when a host wants to transmit but the medium is busy and the previous backoff has been done. To perform a backoff, a counter is first set to an integer randomly picked from its current backoff window. If the *channel clear assessment* (CCA) mechanism of the host detects no channel activity during the past *slot* (a fixed period), the counter is decreased by one. When the counter reaches zero, the backoff procedure is finished.

Now consider the scenario where several neighbor hosts hear a broadcast from host **X**. There are several reasons for collisions to occur. First, if the surrounding medium of **X** has been quiet for enough long, all **X**'s neighbors may have passed their backoff procedures. Thus, after hearing the broadcast message (and having passed the DIFS period), they may all start rebroadcasting at around the same time. This is especially true if carriers can not be sensed immediately due to such as RF delays and transmission latency. Second, because the RTS/CTS forewarning dialogue is not used in a broadcast transmission, the damage of collision is more serious. Third, once collision occurs, without *collision detection* (CD), a host will keep transmitting the packet even if some of foregoing bits have been garbled. And the longer the packet is, the more the waste.

The above problem is not addressed in the ordinary IEEE 802.11 MAC activities, possibly because the one-to-many behavior is not considered therein. For all the above reasons, we believe that the broadcast storm problem deserves serious studies in a MANET environment.

3 Mechanisms to Reduce Redundancy, Contention, and Collision

One approach to alleviate the broadcast storm problem is to inhibit some hosts from rebroadcasting to reduce the redundancy, and thus contention and collision. In the following, we present five schemes to do so. These schemes differ in how a mobile host estimates redundancy and how it accumulates knowledge to assist its decision. Except the last scheme, which relies on some local connectivity information, all schemes operate in a fully distributed manner.

3.1 Probabilistic Scheme

An intuitive way to reduce rebroadcasts is to use probabilistic rebroadcasting. On receiving a broadcast message for the first time, a host will rebroadcast it with probability P . Clearly, when $P = 1$, this scheme is equivalent to flooding.

Note that to respond to the contention and collision problems addressed in Section 2.2.3, we should insert a small random delay (a number of slots) before rebroadcasting the message. So the timing of rebroadcasting can be differentiated.

3.2 Counter-Based Scheme

When a host tries to rebroadcast a message, the rebroadcast message may be blocked by busy medium, backoff procedure, and other queued messages. There is a chance for the host to hear the same message again and again from other rebroadcasting hosts before the host actually starts transmitting the message.

In Section 2.2.1 we have shown that $EAC(k)$, the expected additional coverage after heard the message k times, is expected to be lower when k increases. We can prevent a host from rebroadcasting when the expected additional coverage of the host's rebroadcast becomes too low. This is what the counter-based scheme is based on. Specifically, a counter c is used to keep track of the number of times the broadcast message is received. A counter threshold C is chosen. Whenever $c \geq C$, the rebroadcast is inhibited. The scheme is formally derived below.

- S1. Initialize counter $c = 1$ when a broadcast message msg is heard for the first time. In S2, if msg is heard again, interrupt the waiting and perform S4.
- S2. Wait for a random number of slots. Then submit msg for transmission and wait until the transmission actually starts.
- S3. The message is on the air. The procedure exits.
- S4. Increase c by one. If $c < C$, resume the interrupted waiting in S2. Otherwise $c = C$, proceed to S5.
- S5. Cancel the transmission of msg if it was submitted in S2. The host is prohibited from rebroadcasting msg . Then exits.

3.3 Distance-Based Scheme

In the previous scheme, a counter is used to decide whether to drop a rebroadcast or not. In this scheme, we will use the relative distance between hosts to make the decision.

For instance, suppose host \mathbf{H} heard a broadcast message from \mathbf{S} for the first time. If the distance, say d , between \mathbf{H} and \mathbf{S} is very small, there is little additional coverage \mathbf{H} 's rebroadcast can provide. If d is larger, the additional coverage will be larger. In the extreme case, if $d = 0$, the additional coverage is 0 too. Earlier, we have analyzed the relationship between the distance d and the additional coverage $\pi r^2 - INTC(d)$. So this can be used as a metric by \mathbf{H} to determine whether to rebroadcast or not.

Now, suppose that before a rebroadcast message is actually sent, host \mathbf{H} has heard the same message several times. Let d_{min} be the distance to the nearest host from which the

same message is heard. Then \mathbf{H} 's rebroadcast will provide additional coverage no more than $\pi r^2 - INTC(d_{min})$. In our distance-based scheme, we will use d_{min} as the metric to evaluate whether to rebroadcast or not. If d_{min} is smaller than some distance threshold D , the rebroadcast transmission of \mathbf{H} is cancelled. The scheme is formally derived below. In Section 4, we will test several possible values of D .

- S1. When a broadcast message msg is heard for the first time, initialize d_{min} to the distance to the broadcasting host. If $d_{min} < D$, proceed to S5. In S2, if msg is heard again, interrupt the waiting and perform S4.
- S2. Wait for a random number of slots. Then submit msg for transmission and wait until the transmission actually starts.
- S3. The message is on the air. The procedure exits.
- S4. Update d_{min} if the distance to the host from which msg is heard is smaller. If $d_{min} < D$, proceed to S5. Otherwise, resume the interrupted waiting in S2.
- S5. Cancel the transmission of msg if it was submitted in S2. The host is inhibited from rebroadcasting msg . Then exits.

Below, we comment on how to obtain the distance information. One possibility is to estimate from the signal strength on which a message is received. Specifically, let P_t and P_r be the power levels on which a message is sent and received, respectively. According to [26], $P_r = P_t (\frac{c_1}{d})^n c_2$, where n , c_1 , and c_2 are constants related to physical environment, the carrier's wavelength, and antenna gains, respectively. Since P_r and P_t can be measured, the distance d can be estimated from this formula.

Having understood the relationship between the distance and the power, we can even directly replace the role of distances by signal strengths by establishing a signal-strength threshold. As a comment, we note that signal strength information was also used in [12] to facilitate routing in a MANET.

3.4 Location-Based Scheme

Earlier we have used the number of times that a broadcast message has been heard or the distances to sending hosts as our rebroadcasting metrics. If we can acquire the locations of those broadcasting hosts, it is even possible to estimate the additional coverage more precisely. Such an approach may be supported by positioning devices such as GPS (Global Positioning System) receivers [17]. We note that location information was also used to facilitate route discovery in a MANET [5, 18].

Without loss of generality, let a host's location be $(0,0)$ (here we use xy -coordinate to facilitate our presentation; in fact, devices such as GPS receivers can provide 3-D locations in longitude, latitude, and altitude). Suppose a host has received the same broadcast message from k hosts located at




Figure 5: Scenarios of using convex polygons to determine whether to rebroadcast or not. (a) Host \mathbf{X} is within the triangle formed by three sending hosts. (b) \mathbf{X} is outside of the polygon. (c) Analysis of maximum loss of additional coverage on using the polygon test.

$(x_1, y_1), (x_2, y_2), \dots, (x_k, y_k)$. We can calculate the additional area that can be covered if the host rebroadcasts the message. Let $AC((x_1, y_1), (x_2, y_2), \dots, (x_k, y_k))$ denote the additional coverage divided by πr^2 . Then we can compare this value to a predefined coverage threshold A ($0 < A < 0.61$) to determine whether the receiving host should rebroadcast or not. The scheme is formally derived below.

- S1. When a broadcast message msg is heard for the first time, initialize AC to the additional coverage provided by the host's rebroadcast. If $AC < A$, proceed to S5. In S2, if msg is heard again, interrupt the waiting and perform S4.
- S2. Wait for a random number of slots. Then submit msg for transmission and wait until the transmission actually starts.
- S3. The message is on the air. The procedure exits.
- S4. Update AC . If $AC < A$, proceed to S5. Otherwise, resume the interrupted waiting in S2.
- S5. Cancel the transmission of msg if it was submitted in S2. The host is inhibited from rebroadcasting msg . Then exits.

One obstacle in using the above scheme is the cost of calculating AC , which is related to calculating many intersections among several circles. This problem is difficult already when there are four circles. One possibility is to use a grid-filling approximation to estimate its value.

An alternative is using a convex polygon to determine whether a rebroadcast should be carried out or not. For instance, suppose host \mathbf{X} received a broadcast message three times from hosts \mathbf{A} , \mathbf{B} , and \mathbf{C} . In Fig. 5(a), it shows that if \mathbf{X} is in the convex polygon formed by \mathbf{A} , \mathbf{B} , and \mathbf{C} , the additional coverage of \mathbf{X} 's rebroadcast is small or even none. On the contrary, as shown in Fig. 5(b), if \mathbf{X} is not in the polygon, it is likely that the rebroadcast will provide more additional coverage (shaded area). As a result, we can allow the host to rebroadcast only if it is not located within the convex polygon.

The following is to justify the reason of the above convex approximation through geometry calculation. We show that

if the polygon test prevents a host (\mathbf{X}) from rebroadcasting (\mathbf{X} is within the polygon), at most 22% of coverage will be lost in the extreme case. Observe that the additional coverage will be the largest when \mathbf{X} is located on some boundary of the polygon. Let \mathbf{A} and \mathbf{B} be the two end points of that boundary. We see that the additional coverage will be the largest if \mathbf{A} and \mathbf{B} are each separated from \mathbf{X} by the transmission range r as illustrated in Fig. 5(c). It is not hard to find that the size of shaded area in Fig. 5(c) is

$$4 \left[\int_0^{r/2} \sqrt{r^2 - x^2} \cdot dx - \int_{r/2}^r \sqrt{r^2 - x^2} \cdot dx \right] \\ = (\sqrt{3} - \frac{\pi}{3})r^2 \approx 0.22\pi r^2.$$

Now it is reasonable to say that if a host is in the convex polygon formed by the locations of previously sending hosts, the additional coverage that the host can provide is well below 22%.

3.5 Cluster-Based Scheme

The above schemes are based on *statistical* and *geometric* modeling to estimate the additional coverage of a rebroadcast. In the following, we show how to develop an approach based on *graph* modeling. Specifically, we will adopt the *clustering* concept by [15] to derive our scheme. Note that the clustering technique has been used to solve other problems in MANETs (e.g., traffic coordination [13], routing [13], and fault-tolerance [1]).

We first summarize the *cluster formation algorithm* proposed in [15]. It is assumed that a host periodically sends packets to advertise its presence. Thus any host can determine its connectivity with other hosts on its own. Each host has a unique ID. A *cluster* is a set of hosts formed as follows. A host with a local minimal ID will elect itself as a *cluster head*. All surrounding hosts of a head are *members* of the cluster identified by the head's ID. Within a cluster, a member that can communicate with a host in another cluster is a *gateway*. To take mobility into account, when two heads meet, the one with a larger ID gives up its head role. Fig. 6 shows an example of a clustered MANET.




Figure 6: An example of a clustered MANET. Clustering provides a host the connectivity information of its surroundings.

Now back to our broadcast storm problem. We assume that clusters have been formed in the MANET and will be maintained regularly by the underlying cluster formation algorithm. In a cluster, the head's rebroadcast can cover all other hosts in that cluster if its transmission experiences no collision. Apparently, to propagate the broadcast message to hosts in other clusters, gateway hosts should take the responsibility. But there is no need for a non-gateway member to rebroadcast the message. Our cluster-based scheme is formally developed as follows.

- S1. When the broadcast message *msg* is heard, if the host is a non-gateway member, the rebroadcast is inhibited and the procedure exits. Otherwise, the host is either a head or a gateway. Proceed to S2.
- S2. Use any of the probabilistic, counter-based, distance-based, and location-based schemes to determine whether to rebroadcast or not.

Note that the above scheme is derived by incorporating the schemes developed earlier into it. Step S1 is to eliminate non-gateway members from rebroadcasting. As a cluster may still have many gateway members, step S2 then further utilizes other knowledge (such as additional coverage) to reduce the number of rebroadcasts.

4 Performance Simulation

We have developed a simulator using C++. Central to the simulator is a discrete-event engine designed to simulate systems that can be modeled by processes communicating through signals. A simplified version of MAC specification in IEEE Std 801.11 is referenced to simulate CSMA/CA behavior among hosts.

The parameters that are fixed in our simulations are the transmission radius (500 meters), the packet size (280 bytes), the transmission rate (1M bits per second), and DSSS physical layer timing (PLCP overhead, the slot time, DIFS, back-off window size, as suggested in IEEE Std 801.11).

A geometric area called a *map* which contains one hundred mobile hosts is simulated. A map can be of size 1×1 , 3×3 , 5×5 , 7×7 , or 10×10 units, where a unit is of length 500 meters (equal to the transmission radius). The arrival

rate is one broadcast per second, and the broadcasting host is randomly picked.

The performance metrics to be observed are:

- *REachability (RE)*: the number of mobile hosts receiving the broadcast message divided by the total number of mobile hosts that are reachable, directly or indirectly, from the source host.
- *Saved ReBroadcast (SRB)*: $(r - t)/r$, where r is the number of hosts receiving the broadcast message, and t is the number of hosts actually transmitted the message.
- *Average latency*: the interval from the time the broadcast was initiated to the time the last host finishing its rebroadcasting.

The simulation results of the probabilistic, counter-based, distance-based, and location-based schemes are shown in Fig. 7, Fig. 8, Fig. 9, and Fig. 10, respectively. Each point in these figures represents our result obtained from a simulation run containing 10,000 broadcast requests. Fig. 7(a) shows the observed *RE* and *SRB* associated with the probabilistic scheme. In a small map (which implies a dense host distribution), a small probability P is sufficient to achieve high reachability, while a larger P is needed if the host distribution is sparse. The amount of saving (*SRB*) decreases, roughly proportionally to $(1 - P)$, as P increases. Also, the performance of broadcasting by flooding can be found at the position where the probability $P = 1$. Fig. 7(b) shows the broadcast latency at various P values. Interestingly, a MANET with sparse hosts tend to complete broadcasting in a faster speed than one with denser hosts.

Fig. 8 shows the performance of the counter-based scheme. From Fig. 8(a), we can see that the reachability *RE* in fact reaches about the same level as that of the flooding scheme when the counter threshold $C \geq 3$. However, various levels of saving (*SRB*) can be obtained over the flooding scheme, depending on the density of hosts in a map. For instance, the scheme in higher density maps (e.g., 1×1 , 3×3 , or 5×5) can offer $27 \sim 67\%$ saving at $C = 3$, while in lower density maps (e.g., 7×7 , 9×9 , or 11×11) $8 \sim 20\%$ saving. When C is larger (say 6) and the map is sparse (say 11×11), the amount of saving decreases sharply. This is because the number of neighbors of a host tends to be small (2.4 neighbors per host in an 11×11 map), and thus it is less likely that a host will receive the same broadcast message more than C times. So a threshold C of 3 or 4 is an appropriate choice.

The performance of the distance-based scheme is shown in Fig. 9. Note that the values of threshold D in the figure are chosen purposely in a way so as to make a reasonable comparison between the distance-based scheme and the counter-based scheme. We mainly use the additional coverage as our measurement. For instance, when $C = 2$, we know that $EAC(2) \approx 0.187$ from Fig. 3. So we should choose a value of D to match the value of addition coverage, i.e.,

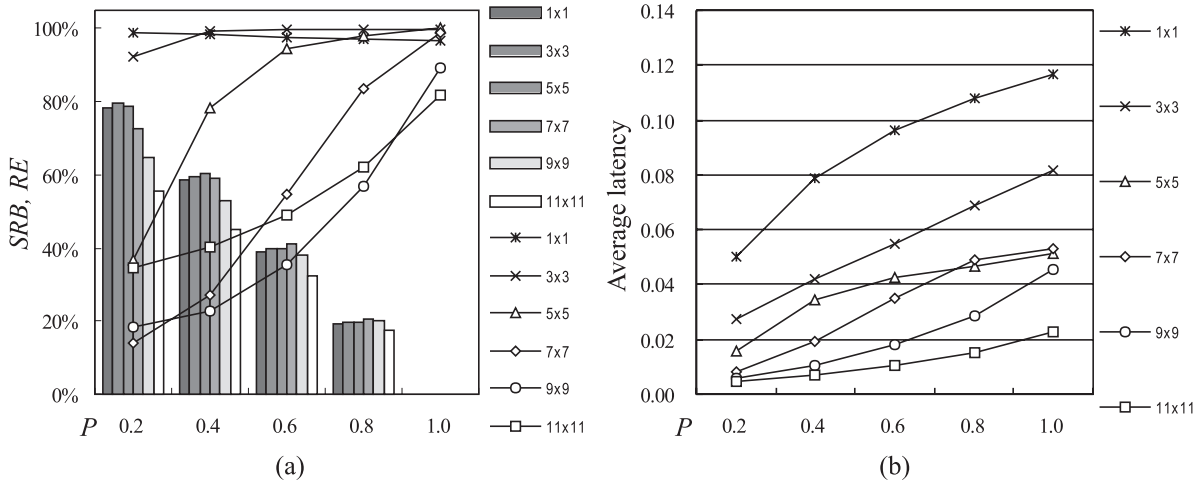


Figure 7: Performance of the probabilistic scheme. (a) Probability P vs. reachability RE (shown in lines) and saved rebroadcast SRB (shown in bars). (b) Probability P vs. average latency.

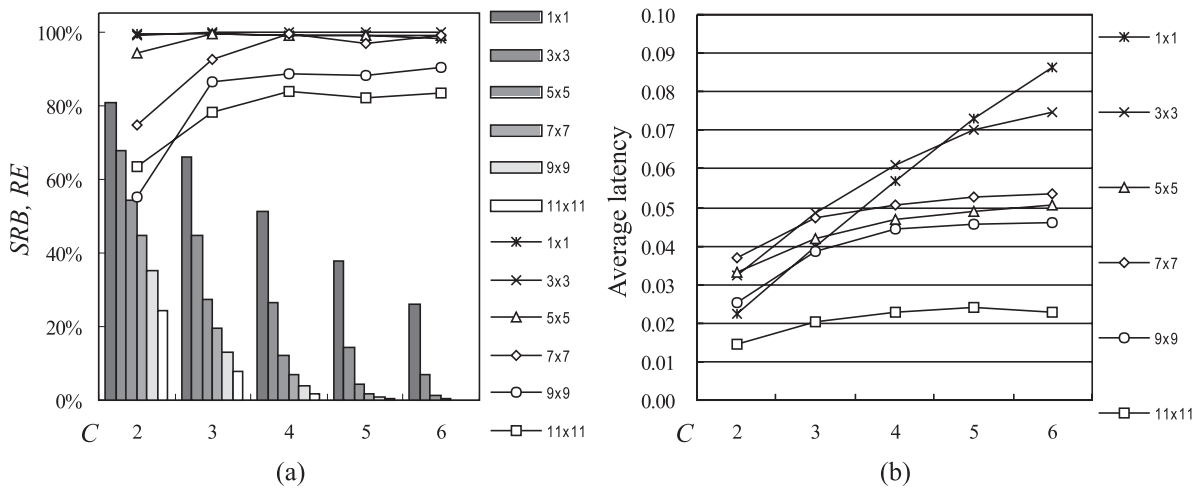


Figure 8: Performance of the counter-based scheme. (a) Counter threshold C vs. reachability RE (shown in lines) and saved rebroadcast SRB (shown in bars). (b) Counter threshold C vs. average latency.

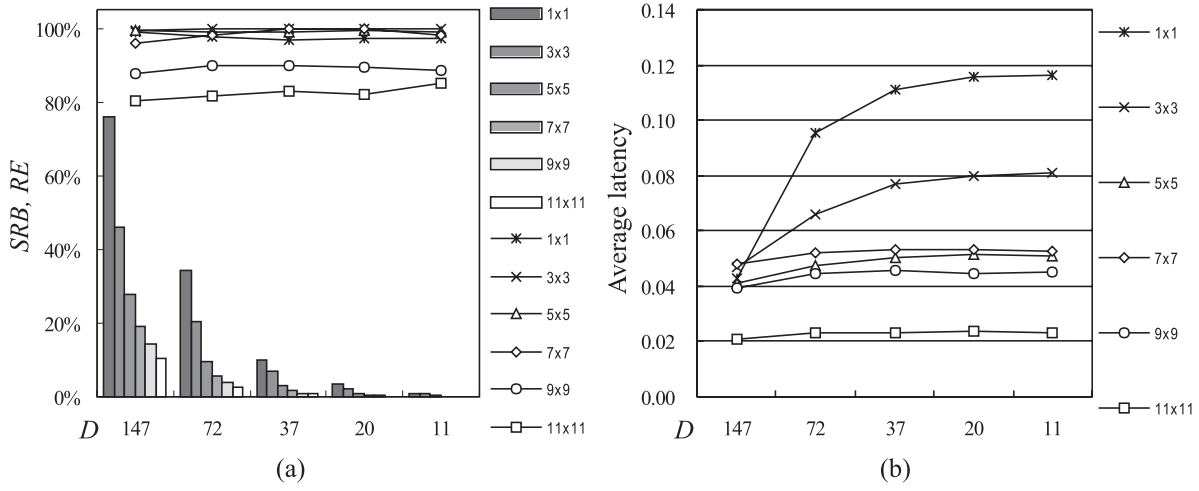


Figure 9: Performance of the distance-based scheme. (a) Distance threshold D vs. reachability RE (shown in lines) and saved rebroadcast SRB (shown in bars). (b) Distance threshold D vs. average latency.

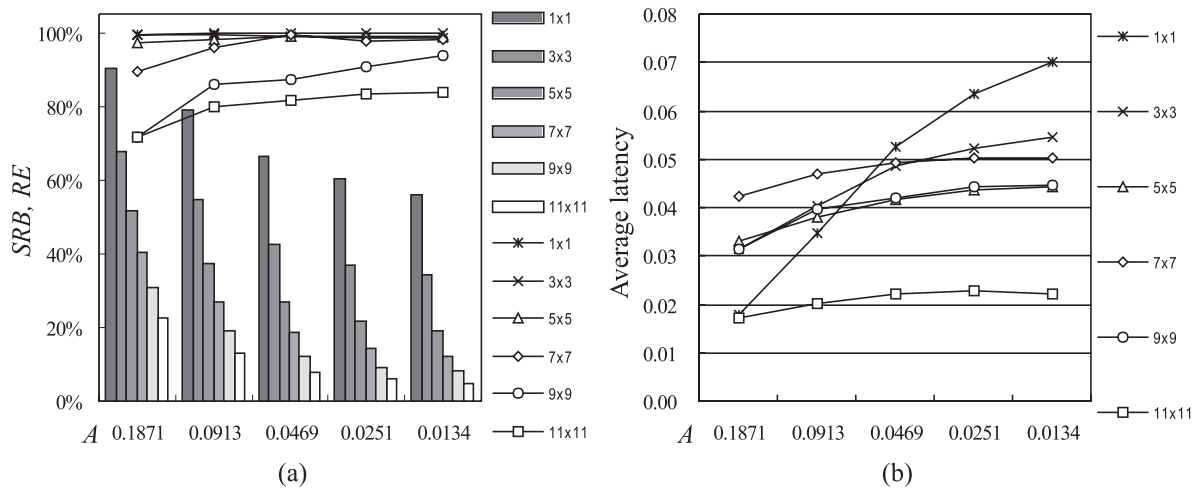


Figure 10: Performance of the location-based scheme. (a) Coverage threshold A vs. reachability RE (shown in lines) and saved rebroadcast SRB (shown in bars). (b) Coverage threshold A vs. average latency.

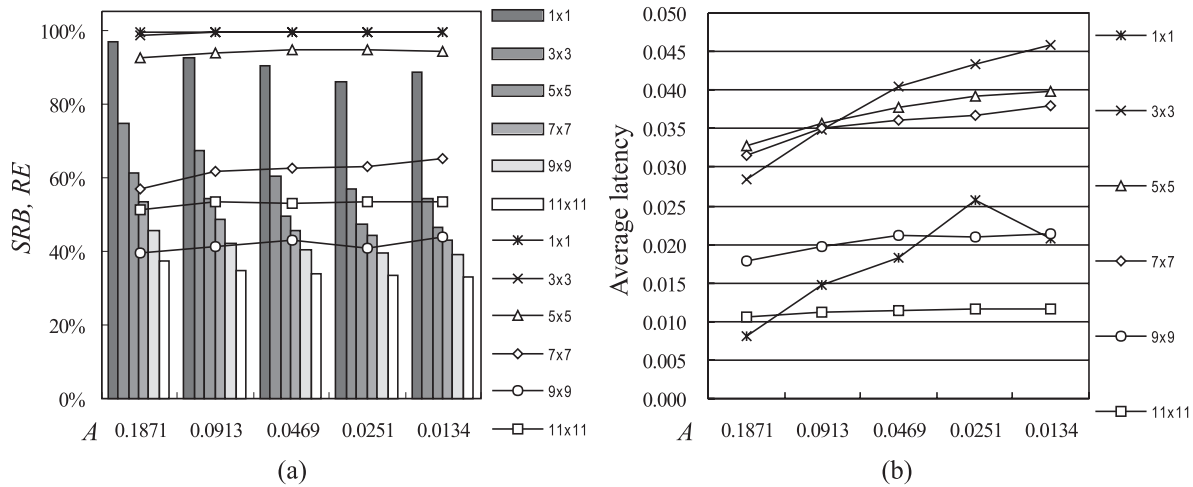


Figure 11: Performance of the cluster-based scheme by applying the location-based scheme in its step S2. (a) The coverage threshold A vs. reachability RE (shown in lines) and saved rebroadcast SRB (shown in bars). (b) The coverage threshold A vs. average latency.

$(\pi r^2 - \text{INTC}(D)) / (\pi r^2) \approx 0.187$. This gives a $D = 147$. The other D values along the x-axis in Fig. 9 are derived in a similar way.

By comparing Fig. 8(a) and Fig. 9(a), we see that although the distance-based scheme can provide a better reachability, not much of rebroadcasts are saved. In fact, the SRB values are worse than those of the counter-based scheme. This also incurs a higher broadcast latency than that of the counter-based scheme, as shown in Fig. 8(b) and Fig. 9(b). The reason that the distance-based scheme saves less among of rebroadcasts than the counter-based scheme does is as follows. In the distance-based scheme, a host may have heard a broadcast message so many times but still rebroadcast the message because none of the transmission distances are below a given distance threshold, where the rebroadcast would have been canceled if the counter-based scheme is used.

Fig. 10 illustrates the performance of the location-based scheme at various threshold values of A . Note that the values of A in the figure are chosen purposely for a similar reason as that described in the distance-based scheme. Overall, this scheme reveals the best performance in all aspects over all other schemes because it utilizes the exact information to calculate the additional coverage. It provides high levels of reachability, while remaining a good amount of saving. Even in sparse maps, it still provides a relatively higher level of saving. Because of the saving, the broadcast latency is also the best among all schemes.

Fig. 11 shows the performance of the cluster-based scheme where the location-based scheme is incorporated in its step S2. Compared to the original location-based scheme, the cluster-based scheme apparently saves much more rebroadcasts and leads to shorter average broadcast latencies. Unfortunately, the reachability is unacceptable at sparser areas (such as 7×7 , 9×9 , and 11×11 maps). This is probably




Figure 12: An example of hidden terminal problem in the cluster-based scheme. Hosts **A**, **C**, and **D** are gateways, **B** and **E** are heads. Non-gateway members are not shown.

because when the number of hosts participating in rebroadcasting is reduced (mostly by step S1 of the cluster-based scheme), the collisions caused by the hidden terminal problem will significantly reduce the chance of successful transmissions. For example, Fig. 12 illustrates a scenario where a broadcast message was propagated through gateway **A** to head **B**. After **B** rebroadcast the message, gateways **C** and **D** will try to rebroadcast the message. Unfortunately, because **C** and **D** can not hear each other, the two rebroadcasts from **C** and **D** are very likely to collide at **E**. Such loss of a message may significantly reduce the reachability, especially in a sparse map.

It is also worth mentioning that there is an anomaly that the average latencies in the 1×1 map are pretty low (see Fig. 11(b)). This is because in our simulations there are in average only 1.2 cluster heads in the map, and thus when there is only one head, the broadcast will be completed very quickly because there is no gateway.

We have also studied how our schemes perform under different load conditions. We varied the arrival rates at 1, 20, 40, 60, 80, and 100 broadcasts per second. Due to space

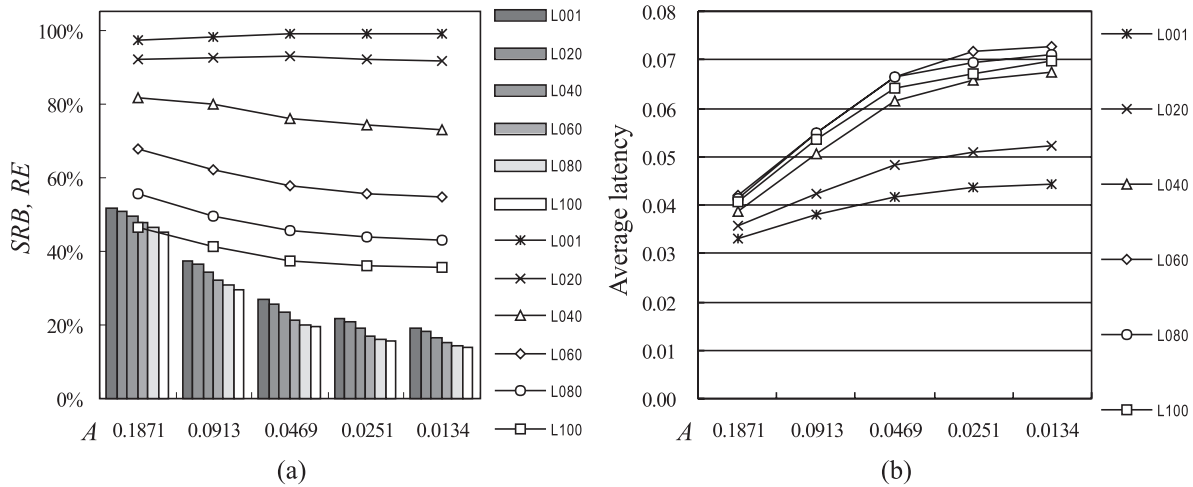


Figure 13: Performance of the location-based scheme in a 5×5 map under different load conditions. (a) Coverage threshold A vs. RE (shown in lines) and SRB (shown in bars). (b) Coverage threshold A vs. average latency.

limit, only the result of the location-based scheme in a 5×5 map is shown in Fig. 13. From the figure, we observe that a heavier load will result in a lower reachability. This is probably due to heavier collisions. At higher load, the reachability even degrades as the coverage threshold A becomes smaller, which contradicts with our intuition. This is because when the threshold A is low, less rebroadcasts are saved, which leads to more collisions. This also suggests that one consider dynamically adjusting the threshold A so that the scheme can adapt to different load conditions.

5 Conclusions

This paper has identified an important issue in a MANET, the broadcast storm problem. We have demonstrated, through analyses and simulations, how serious this problem could be. Several schemes, namely probabilistic, counter-based, distance-based, location-based, and cluster-based schemes, have been proposed to alleviate this problem. Simulation results based on different threshold values are presented to verify and compare the effectiveness of these schemes. As compared to the basic flooding approach, a simple counter-based scheme can eliminate many redundant rebroadcasts when the host distribution is dense. If location information is available through devices such as GPS receivers, the location-based scheme is the best choice because the scheme can eliminate even more redundant rebroadcasts under all kinds of host distributions without compromising the reachability. Future research could be on how to incorporate these schemes with other MANET protocols, such as reliable broadcast, multicast, and routing protocols.

References

- [1] S. Alagar, S. Venkatesan, and J. Cleveland. Reliable broadcast in mobile wireless networks. In *Military Communications Conference, MILCOM Conference Record*, volume 1, pages 236–240, 1995.
- [2] R. Bar-Yehuda, O. Goldreich, and A. Itai. Efficient emulation of single-hop radio network with collision detection on multi-hop radio network with no collision detection. *Distributed Computing*, 5(2):67–72, 1991.
- [3] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45(1):104–126, Aug. 1992.
- [4] R. Bar-Yehuda, A. Israeli, and A. Itai. Multiple communication in multihop radio networks. *SIAM Journal on Computing*, 22(4):875–887, Aug. 1993.
- [5] S. Basagni, I. Chlamtac, V. R. Syrotiuk, and B. A. Woodward. A distance routing effect algorithm for mobility (DREAM). In *Proc. IEEE/ACM Intl. Conf. on Mobile Computing and Networking MOBICOM*, pages 76–84, 1998.
- [6] J. Broch, D. B. Johnson, and D. A. Maltz. The dynamic source routing protocol for mobile ad hoc networks, 1998. Internet Draft.
- [7] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proc. IEEE/ACM Intl. Conf. on Mobile Computing and Networking MOBICOM*, pages 85–97, 1998.

- [8] I. Chlamtac and A. Faragó. Making transmission schedules immune to topology changes in multi-hop packet radio networks. *IEEE/ACM Trans. on Networking*, 2(1):23–29, Feb. 1994.
- [9] I. Chlamtac and S. Kutten. On broadcasting in radio networks - problem analysis and protocol design. *IEEE Trans. on Commun.*, COM-33(12):1240–1246, Dec. 1985.
- [10] I. Chlamtac and S. Kutten. Tree-based broadcasting in multihop radio networks. *IEEE Trans. on Commun.*, COM-36(10):1209–1223, Oct. 1987.
- [11] I. Chlamtac and O. Weinstein. The wave expansion approach to broadcasting in multihop radio networks. *IEEE Trans. on Commun.*, COM-39(3):426–433, Mar. 1991.
- [12] R. Dube, C. D. Rais, K.-Y. Wang, and S. K. Tripathi. Signal stability based adaptive routing (SSA) for ad hoc mobile networks. Technical Report CS-TR-3646, University of Maryland, College Park, Aug. 1996.
- [13] M. Gerla and J. T.-C. Tsai. Multicluster, mobile, multimedia radio network. *ACM-Baltzer Journal of Wireless Networks*, 1(3):255–265, 1995.
- [14] Z. J. Haas and M. R. Pearlman. The zone routing protocol (ZRP) for ad hoc networks, 1998. Internet Draft.
- [15] M. Jiang, J. Li, and Y. C. Tay. Cluster based routing protocol (CBRP) functional specification, 1998. Internet Draft.
- [16] J.-H. Ju and V. O. K. Li. An optimal topology-transparent scheduling method in multihop packet radio networks. *IEEE/ACM Trans. on Networking*, 6(3):298–306, June 1998.
- [17] E. D. Kaplan, editor. *Understanding GPS: Principles and Applications*. Artech House, Boston, MA, 1996.
- [18] Y.-B. Ko and N. H. Vaidya. Location-aided routing (LAR) in mobile ad hoc networks. In *Proc. IEEE/ACM Intl. Conf. on Mobile Computing and Networking MOBICOM*, pages 66–75, 1998.
- [19] LAN MAN Standards Committee of the IEEE Computer Society, editor. *IEEE Std 802.11-1997, Wireless LAN Medium Access Control MAC and Physical Layer (PHY specifications)*. IEEE, Nov. 1997.
- [20] C. Lee, J. E. Burns, and M. H. Ammar. Improved randomized broadcast protocol in multi-hop radio networks. Technical Report GIT-CC-93-14, College of Computing, Georgia Institute of Technology, Feb. 1993.
- [21] J. P. Macker and M. S. Corson. Mobile ad hoc networking and the IETF. *ACM Mobile Computing and Communications Review*, 2(3):7–9, July 1998.
- [22] E. Pagani and G. P. Rossi. Reliable broadcast in mobile multihop packet networks. In *Proc. IEEE/ACM Intl. Conf. on Mobile Computing and Networking MOBICOM*, pages 34–42, 1997.
- [23] C. E. Perkins and E. M. Royer. Ad hoc on demand distance vector (AODV) routing, 1998. Internet Draft.
- [24] S. Ramanathan and E. L. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE/ACM Trans. on Networking*, 2(2):166–177, April 1993.
- [25] A. Sen and M. L. Huson. A new model for scheduling packet radio networks. *ACM/Baltzer Wireless Networks*, 3:71–82, 1997.
- [26] E. K. Wesel. *Wireless Multimedia Communications: Networking Video, Voice, and Data*. Addison-Wesley, One Jacob Way, Reading, Massachusetts 01867 USA, 1998.