

A family of Encounter-based broadcast Protocols for Mobile Ad-hoc Networks

D.E. Cooper P. Ezhilchelvan I. Mitrani

School of Computing Science, University of Newcastle
Newcastle upon Tyne, NE1 7RU, U.K.

[d.e.cooper, paul.ezhilchelvan, isi.mitrani]@ncl.ac.uk

Abstract. A family of message propagation protocols for highly mobile ad-hoc networks is defined, and is studied analytically and by simulation. The coverage of a message (the fraction of nodes that receive it), can be made arbitrarily close to 1, at a moderate cost of extra message traffic. Under certain simplifying assumptions, it is shown that a high coverage is achieved by making a total of $O(n \ln n)$ broadcasts, where n is the number of nodes, and the time to propagate a message is $O(\ln n)$. The effect of various parameters on the protocol performance is examined.

Keywords: Ad-hoc networks, message propagation, broadcasting.

1 Introduction

Recent advances in the technologies of mobile devices and wireless communication have given rise to an increasingly popular form of networking, called Mobile Ad-hoc networking. A Mobile Ad-hoc network (MANET) consists of small, versatile and powerful mobile computing devices (nodes). It is typically formed at short notice and does not make use of any fixed networking infrastructure. A distinguishing feature of a MANET is that the nodes are not just the sources of message traffic but also engage in forwarding messages to final destinations; given that the nodes can be highly mobile, a MANET is a dynamic network characterized by frequent and hard-to-predict topological changes.

An application of a mobile network usually involves user collaboration towards achieving a common goal, in situations where access to base stations is unavailable or unreliable (e.g., command and control or disaster relief). The success of such collaborative undertakings depends to a large extent on the provision of reliable multicast [5]. That is, a message originating at any node should reach all other nodes within a reasonably short period of time. Unfortunately, both the nature of the devices (limited memory and power), and their mobility, imply that a guaranteed reliable multicast is normally harder to achieve in a MANET. Our objective, therefore, is to devise and evaluate multicast protocols which aim to maximize the probability of delivering a message to all nodes, while keeping the propagation time as low as possible.

Existing work in this area has concentrated on minimizing the number of broadcasts carried out while propagating a message. Several protocols have been proposed, where the nodes maintain pro-actively, or construct on demand, distributed state information about the network topology. That state information is then used for the purpose of improving coverage with small overhead (see [2, 10, 4, 1]). When the degree of mobility is low, these protocols perform well, but when it is high, the network state information can become out-of-date quickly and the coverage achieved (i.e., the fraction of nodes that receive a message) can be poor [9, 11].

A topology-independent and stateless protocol that seems to work better in highly mobile networks is ‘flooding’. Every node broadcasts every message once, either immediately upon receipt or after a random interval (for a study of basic flooding, see Ho et al [6]; an optimized version was examined in Ni et al [8]). The coverage achieved by flooding depends not only on the mobility pattern, but also on the ‘density’ of nodes (usually defined as the average number of nodes within a disc of radius equal to the wireless range). When the density is low, the flooding coverage tends to be poor.

We propose, and study, a family of protocols which preserve the topology-independent nature of flooding, while being able to achieve coverage levels arbitrarily close to 1, for any node density. Of course a specific high coverage cannot be guaranteed in any given instance, but can be expected with high probability. These protocols are based on a notion of ‘encounter’, and are controlled by an ‘encounter threshold’ parameter. The cost paid for a high coverage is an increase in the message traffic, since messages are broadcast more than once by each node. Under certain simplifying assumptions, analysis shows that to achieve a coverage close to 1 in a network with n nodes, the total average number of broadcasts per message is on the order of $O(n \ln n)$. This is a moderate increase on the $O(n)$ broadcasts carried out in flooding. The propagation time of a message is on the order of $O(\ln n)$. Various aspects of the protocols’ performance are examined by simulation (such an investigation, using a less realistic simulator, was presented in [3]).

The estimate of the necessary threshold provided by the analytical model is deliberately pessimistic. Because of that, it is worth introducing schemes for reducing the number of unnecessary broadcasts during the propagation of a message. One such scheme, *Random Assessment Delay* (RAD), is investigated by simulation.

The model, and the message propagation protocols, are described in section 2. Some analytical results concerning the propagation time and the number of broadcasts are obtained in section 3. The outcomes of a number of simulation experiments are presented in section 4, while section 5 summarizes the results obtained and outlines avenues of further enquiry.

2 The model

The system under consideration consists of n mobile nodes which move within a given terrain. The nodes communicate with each other using wireless technology, but without any fixed network infrastructure support. That is, the nodes themselves are the sources as well as the forwarders of the message traffic, and thus form a mobile ad-hoc network. Each node has a unique identifier (MAC or IP address). It is assumed that nodes do not fail; however, due to their mobility, they may become disconnected, and reconnected, as they move out of and into each others wireless range. Thus, the structure of the network can change with time in an unpredictable manner. For simplicity, assume that the wireless ranges of all nodes are equal and remain constant during the period of interest.

The movement of each node is governed by some ‘mobility pattern’, which controls its current speed and direction. It is assumed that the n nodes are statistically identical, i.e. the rules of their mobility patterns are the same, and any random variables involved have the same distributions for all nodes.

We shall define a protocol whose principal objective is to deliver a message, originating at any node, to all other nodes with high probability. A secondary objective is to minimize, as far as possible, the memory requirements at each node. In fact, what will be defined is not a single protocol, but a family of protocols depending on an integer parameter, τ .

Node i ($i = 1, 2, \dots, n$) advertises its presence by broadcasting, at regular intervals, a signal carrying its identifier and saying, essentially, ‘hello, this is node i ’. It also listens for similar signals from other nodes and maintains a list, $\{j_1, j_2, \dots, j_k\}$, of the nodes, other than itself, that it can hear. That list is called the ‘current neighbourhood’ of node i . At any moment in time, any current neighbourhood may be empty, or it may contain any number of other nodes.

The current neighbourhood of node i changes when a node which was in it, say j_1 , moves out of range, or when a node which was not in it, say j_{k+1} , moves into range. The latter event is called an ‘encounter’; that is, node i is said to encounter node j_{k+1} . Note that, since ‘hello’ signals are not assumed to be synchronized among the nodes, if node i encounters node j , node j does not necessarily encounter node i at the same time. Also note that, if node j leaves the current neighbourhood of node i and at some later point enters it again, then that entry constitutes an encounter. Nodes do not maintain a history of their current neighbourhoods, in order to keep their memory requirements low.

Now consider a message propagation protocol where each node behaves as follows:

- 1. Upon receiving or originating a new message, m , store it, together with an associated counter, $c(m)$, which is set to zero. Add the sending node to the current neighbourhood, unless already present. If the current neighbourhood contains nodes other than the sending one, broadcast m and increment $c(m)$ by 1.

- 2. At every encounter thereafter, if $c(m) < \tau$, broadcast m and increment $c(m)$ by 1.
- 3. When $c(m) = \tau$, remove m from memory (but keep its sequence number in order to remember that it has been handled).

Thus, every node receiving a message broadcasts it at τ consecutive encounters (one of which may be the message arrival), and then discards it. There are no acknowledgements. The integer τ is called the ‘encounter threshold’. The above protocol, with encounter threshold τ , will be referred to as ‘ τ -propagation’.

When $\tau = 1$, the 1-propagation protocol behaves like flooding (except that the broadcast is delayed until the next encounter if the current neighbourhood contains only the sender). At the other extreme, if $\tau = \infty$, we have an ∞ -propagation protocol whereby messages are kept forever and broadcast at every encounter. Assuming that the mobility pattern is such that every node eventually encounters every other node, ∞ -propagation achieves coverage 1. Of course, ∞ -propagation is not a practical option, but we shall see in section 3 that it can provide some useful insights.

It should be pointed out that τ -propagation trades memory capacity and probability of reaching all nodes against message traffic. Because past histories are not kept and exchanged, messages may be sent again to nodes who have already received them. By increasing the value of τ , the coverage can be made to approach 1, at the cost of having to store more messages for longer periods, and making more broadcasts.

In this paper, we place greater emphasis on evaluating the ability of τ -propagation to achieve high coverage, than on minimizing the message traffic overheads. That is why we assume the following:

- The overheads of collision resolution are negligible.
- Hello signals are sent and monitored at the MAC level; the information necessary to maintain the neighbourhood list is obtained at no extra cost to the higher level protocol.
- Encounters last long enough for a message to be received, i.e. the processing and propagation times of hello and broadcast messages are small enough for the encountered node to remain in the range of the encountering node.

The performance measures of interest are:

- (i) The average response time of τ -propagation, defined as the interval between the arrival (origin) of a message and the moment when no node can propagate it further.
- (ii) The average propagation time of a message, defined as the interval between its arrival and the moment when either all nodes have received it, or no node can propagate it further.
- (iii) The coverage of a message, i.e. the fraction of nodes that have received it by the end of its propagation time.

All of these performance measures are stated in terms of averages. However, each simulation experiment is repeated 50 times with different random number streams. Thus, observing a coverage of 1 implies that *all 50 runs* achieved a coverage of 1.

It is important to be able to choose the value of τ so as to achieve high coverage, without unduly increasing the response and propagation times. This question will be addressed in the following sections.

3 Analytical approximation

Consider an idealized system with n mobile nodes who never cease to propagate the messages they receive (∞ -propagation). Let T be the random variable representing a message propagation time, i.e., the interval between the origin of a message at some node, and the first instant thereafter at which all nodes have received it. If messages are not discarded, and every node eventually encounters every other node, T is finite with probability 1. It is then of interest to estimate its average value, $E(T)$. That quantity will also be used in choosing a suitable value for τ , when designing a practicable τ -propagation protocol.

An estimate for $E(T)$ will be obtained under the following simplifying assumptions:

- (a) Each node experiences encounters at intervals which are exponentially distributed with mean ξ .
- (b) At each encounter, a node meets one other node (in other words, the intervals between hello signals are infinitesimal).
- (c) The node encountered is equally likely to be any of the other nodes; that is, the probability that node i will next encounter node j , $j \neq i$, is equal to $1/(n - 1)$, regardless of past history.

Assumption (a) can be justified by remarking that the interval until the next encounter experienced by a given node — say node 1 — is the smallest of the intervals until its next encounters with node 2, node 3, \dots , node n . It is known (e.g., see [7]) that, under weak assumptions, the interval until the first of many random occurrences is approximately exponentially distributed. The value of ξ depends on the density of nodes, on the speed with which they move, and on the mobility pattern. It may be difficult to determine ξ analytically, but in practice it can be estimated by monitoring the system and taking measurements.

Assumption (b) is deliberately pessimistic, in order to give the estimate the character of an upper bound. If a node encounters more than one other node at the same time, then the propagation will proceed faster. In fact, it will be seen in the experiments that at high densities this assumption is *very* pessimistic.

Assumption (c) is loosely based on the fact that all nodes are statistically identical, and move independently of each other. If the starting positions of the nodes are uniformly distributed, the assumption is justifiable at the first encounter, although it may well be violated in subsequent ones. However, this assumption provides the simplification necessary for analytical tractability. Its

effect on the performance measures will be evaluated in the simulation experiments.

Let $X = \{X(t); t \geq 0\}$ be the Markov process whose state at any given time is the number of nodes that have already received the message. The initial state of X is $X(0) = 1$ (only the originating node has received it). This is a pessimistic simplification; the original neighbourhood may in fact contain other nodes, in which case $X(0)$ would be greater than 1 and the propagation would be faster.

The random variable T is the first passage time of X from state 1 to state n .

Suppose that X is in state k , i.e. k nodes have received the message and $n - k$ have not. If any of the former k nodes encounters any of the latter $n - k$, the process will jump to state $k + 1$. Since each node experiences encounters at rate $1/\xi$, and the probability of encountering any other node is $1/(n - 1)$, the transition rate of X from state k to state $k + 1$, $r_{k,k+1}$, is equal to

$$r_k = \left[\frac{k}{\xi} \right] \left[\frac{n - k}{n - 1} \right]. \quad (1)$$

In other words, the average time that X remains in state k is

$$\frac{1}{r_k} = \frac{(n - 1)\xi}{k(n - k)}. \quad (2)$$

Hence, the average first passage time from state 1 to state n is given by

$$E(T) = (n - 1)\xi \sum_{k=1}^{n-1} \frac{1}{k(n - k)}. \quad (3)$$

This last expression can be simplified by rewriting the terms under the summation sign in the form

$$\frac{1}{k(n - k)} = \frac{1}{n} \left[\frac{1}{k} + \frac{1}{n - k} \right].$$

The two resulting sums are in fact identical. Therefore,

$$E(T) = \frac{2(n - 1)\xi}{n} \sum_{k=1}^{n-1} \frac{1}{k} = \frac{2(n - 1)\xi H_{n-1}}{n}, \quad (4)$$

where H_n is the n th harmonic number. When n is large, the latter is approximately equal to

$$H_n \approx \ln n + \gamma,$$

where $\gamma = 0.5772\dots$ is Euler-Mascheroni's number. Also, when n is large, $(n - 1)/n \approx 1$ and $\ln(n - 1) \approx \ln n$.

We have thus arrived at the following estimate, valid under assumptions (a), (b) and (c):

Proposition 1 *In a large mobile network where messages are not discarded, the average propagation period for a message is approximately equal to*

$$E(T) \approx 2\xi(\ln n + \gamma) . \quad (5)$$

An immediate corollary of Proposition 1 is that, during the propagation period T , the originating node experiences an average of $2(\ln n + \gamma)$ encounters. Other nodes, who receive the message later on, tend to experience fewer encounters. Thus, choosing the encounter threshold, τ , to have the value

$$\tau = 2\lceil \ln n + \gamma \rceil , \quad (6)$$

should ensure that, when the protocol terminates, most nodes will have received the message. This suggestion will be tested experimentally.

Note 1. An attractive aspect of equation (6) is that the only parameter appearing in it is the number of nodes, n . The mobility pattern and the node density do not matter, because assumptions (a), (b) and (c) were deliberately chosen to be pessimistic. In practice, lower thresholds than those provided by (6) will be satisfactory, especially at higher densities. This will be illustrated by the simulation experiments.

Note 2. Since, under τ -propagation, every node that receives a message broadcasts it τ times, the total number of broadcasts per message is on the order of $O(n\tau)$. Hence, if τ is chosen according to (6), the total number of broadcasts per message is on the order of $O(n \ln n)$.

4 Experimental results

A number of experiments were carried out, using the Glomosim simulation package. The objectives were two-fold: First, to evaluate the effect of various parameters on the performance of τ -propagation, as measured by (a) the coverage achieved, and (b) the response and propagation times. The second objective was to investigate some schemes for reducing the protocol overheads, measured by the average number of *redundant broadcast per node*. A broadcast is considered to be redundant if all of its targets have already received the message).

The following factors were kept fixed:

The terrain is a square of dimensions $(1000 \text{ m}) \times (1000 \text{ m})$.

The number of nodes is 64.

The wireless range of all nodes is either 50 m or 180 m. These two ranges correspond to node densities of 0.5 and 6.5 respectively.

The node speed is 2 m/sec.

The mobility pattern is ‘Random Waypoint’: Initially, the nodes are distributed uniformly on the square; thereafter, each node chooses a random destination (also uniformly distributed on the square) and moves towards it at a given speed; upon reaching the destination, the node pauses for a given interval (1 ms in our case), selects a new random destination and so on.

Each run starts at time 0 with a message originating at node 1, and terminates when no node can propagate the message further. For each set of parameter values, the simulation ran 50 times, with different random number seeds, and the performance observations were averaged.

RAD Enhanced Encounter Propagation

One technique that is used to reduce redundant broadcasts is Random Assessment Delay (RAD). Having decided to broadcast a message, either as an originator or at an encounter epoch, a node waits for a random period of time, called the ‘RAD interval’. An encounter which occurs during a RAD interval does not generate a new RAD interval. If, during a RAD interval, a node hears another broadcast of the same message, then the broadcast it had planned is canceled (see [11]). The idea is that if several nodes in a given neighbourhood are in possession of the message and decide to broadcast it, one of them will do so first (the one with the shortest RAD interval), and then the others can keep quiet.

Remember that, under the τ -propagation protocol described in section 2, a node ceases to broadcast message m when the number of encounters, recorded in $c(m)$, reaches the value τ . Now, the addition of RAD may affect the way $c(m)$ is incremented. During a RAD interval associated with message m , the node counts the number of broadcasts of m that it hears. Denote that number by r . If $r = 0$ at the end of the RAD interval, then the node broadcasts m and increments $c(m)$ by 1. If $r > 0$, there are two possibilities:

1. Do not broadcast m and increase $c(m)$ by 1;
2. Do not broadcast m and increase $c(m)$ by r .

These two policies were simulated, but their effects on both performance and overheads turned out to be very similar. Therefore, only the results obtained with policy 1 are displayed.

Figures 1 and 2 show the coverage achieved as a function of the encounter threshold, τ , for node densities 0.5 and 6.5, respectively. The different graphs in each figure correspond to different distributions of the RAD interval. More precisely, the RAD interval is distributed uniformly on $(0, RAD)$, for different values of RAD . Thus, the case $RAD = 0$ corresponds to the original protocol, without the enhancement.

The figures quantify the extent to which the coverage can be improved by increasing τ : at low densities, where flooding performs poorly ($\tau = 0$), the improvement is very considerable; at high densities, flooding performs well and the gain of increasing τ is correspondingly smaller.

Consider the analytical predictions concerning τ . For the assumed terrain area and wireless range, the densities 0.5 and 6.5. For the case of 64 nodes, the encounter threshold given by equation (6) is $\tau = 10$, and the figures indicate that it does, indeed, achieve coverages close to 1. In fact, when the density is high, the threshold provided by equation (6) is rather conservative. This is because, for those densities, assumption (b) in section 3 is too pessimistic.




Fig. 1. Coverage vs Encounter Threshold; Density = 0.5

At the low density of 0.5, the introduction of RAD has no appreciable effect on the coverage. This is not surprising, since an encounter rarely involves more than one other node and so the RAD counter r is usually 0. At the high density of 6.5 and low threshold values, some coverage is lost by RAD because nodes who could usefully broadcast have in fact decided to stop too early.

The gains achieved by the RAD enhancement are shown in figures 3 and 4. At the low density of 0.5, a reduction of about 35% in redundant broadcasts per node can be achieved, without loss of coverage. Increasing the average RAD interval generally leads to bigger reductions, but the incremental gains become lower (the largest gain is achieved in moving from $RAD = 0$ to $RAD = 100$).

At the high density of 6.5, the picture is similar, but now there is a genuine trade-off between coverage and redundancy at low thresholds. At high thresholds, the reduction in redundancy is again achieved without loss of coverage.

Figures 5 and 6 show the average response time and the average propagation time as functions of τ , for the same two densities, 0.5 and 6.5. Remember that the propagation time is the period until all nodes receive the message (or the protocol terminates, if earlier), whereas the response time is the period until the protocol terminates. A noteworthy aspect of these figures is that, while the response time tends to keep growing with τ (as expected), the propagation time increases up to a point, and then decreases. To explain that behaviour, note that when the threshold is low, the coverage is less than 1 and therefore the propagation time is equal to the response time. When the threshold is large, a coverage of 1 is reached and the propagation completes, before nodes have

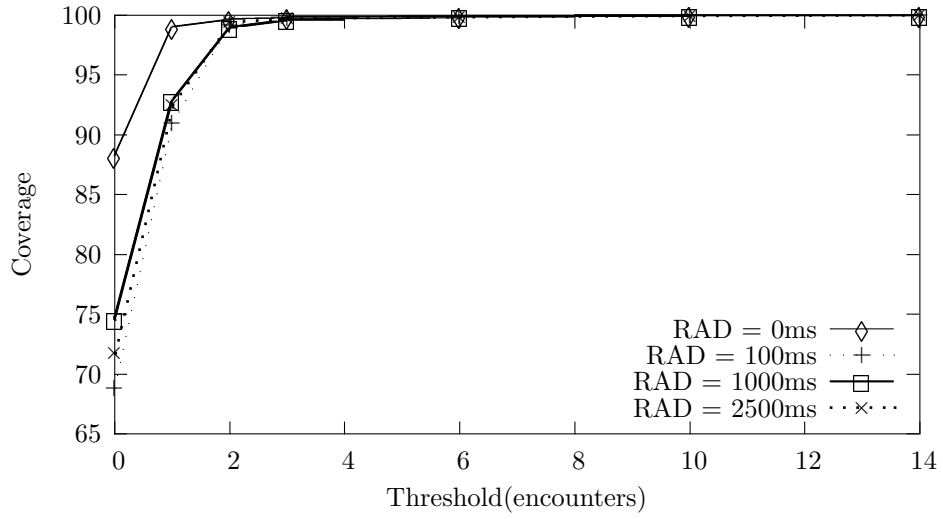


Fig. 2. Coverage vs Encounter Threshold; Density = 6.5

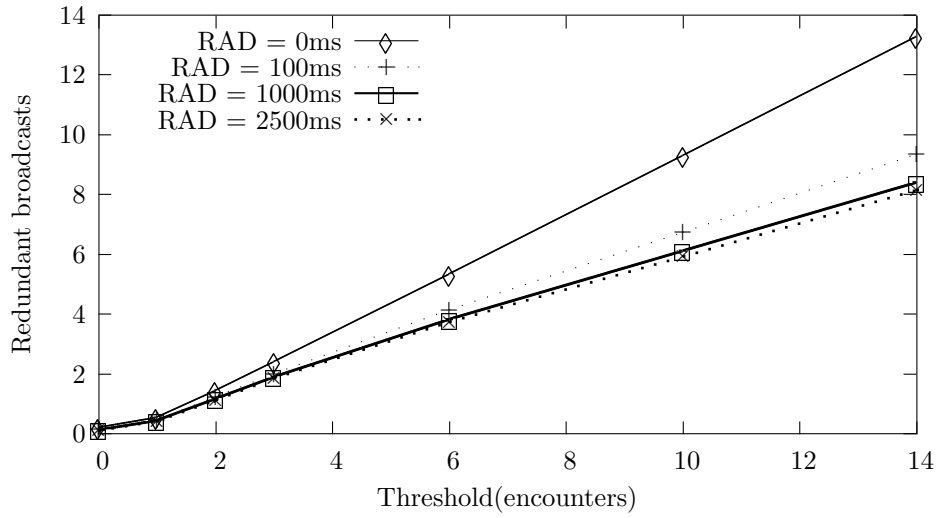


Fig. 3. Redundant broadcasts vs Encounter Threshold; Density = 0.5




Fig. 4. Redundant broadcasts vs Encounter Threshold; Density = 6.5

stopped broadcasting. Moreover, further increases in τ tend to speed up the propagation, but prolong the response time.

Increasing the average RAD interval does, of course, lead to longer average response and propagation times. However, the effect is negligible at low density, and becomes significant at high density mainly for the response time, and only when the threshold is low.

The process of propagating a message among the nodes in a network where the speed and threshold ($\tau = 14$) are fixed, while the density is varied in the range 0.5 – 6.5, is illustrated in figure 7. For the purpose of this figure only, the wireless range was kept at 50 m, and the density was changed by varying the number of nodes between 64 and 828; a different simulator was used. The graphs show how the rate of propagation changes as more and more nodes are covered. At high densities, it takes longer to cover the last 5% of the nodes than the first 95%. This phenomenon is due to the fact that some nodes on the periphery of the terrain can be relatively more difficult to reach than the others. It is less pronounced at lower densities, but is still in evidence: the last 20% of the nodes take about as long to cover as the first 80%.

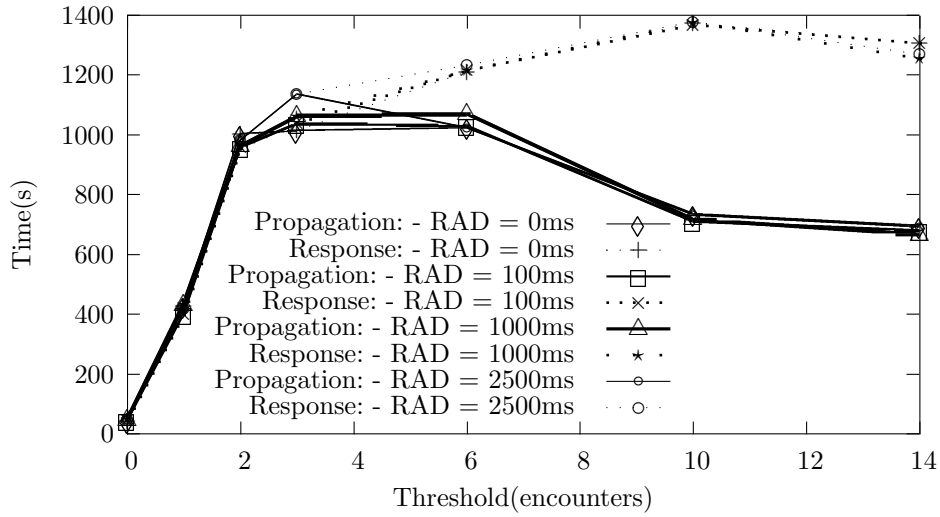


Fig. 5. Response/Propagation Time vs Encounter Threshold; Density =0.5

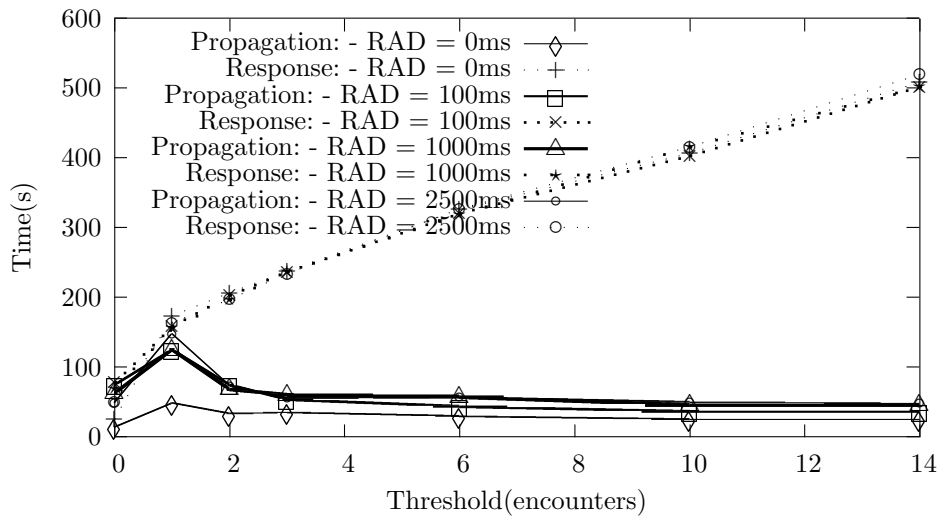


Fig. 6. Response/Propagation Time vs Encounter Threshold; Density =6.5




Fig. 7. Process of propagation: speed = 60 ms^{-1} ; $\tau = 14$

5 Conclusions

The main contributions of this paper can be summarized as follows:

- 1. Introduction of the τ -propagation family (section 2).
- 2. Mobility-independent estimate for the value of τ that achieves high coverage (equation (6)).
- 3. Quantitative performance results obtained by experimentation (section 4).

A more adaptable family of propagation protocols may be designed by introducing a FIFO buffer for messages. Messages would be kept in the buffer, and re-broadcast, until either they are displaced by new messages or they reach an encounter threshold. The number of times a message is broadcast by a node would then change dynamically in response to changing conditions. That number could also be adjusted by keeping track of repeated receptions of the same message. A time-out interval can be introduced, to force the discarding of a message if the node does not experience a sufficient number of encounters. The encounter threshold may be adjusted dynamically, as a function of the number of nodes already encountered, and possibly as a function of the mobility pattern. All these are worthy topics for future research.

It should be pointed out that the Random Waypoint mobility model is particularly favourable to the assumption that each node is equally likely to be encountered. Other mobility models are more likely to violate that assumption.

That is why it is important to investigate the behaviour of models with more pronounced locality properties, such as 'Manhattan Grid' and 'Drunken Walk'.

Acknowledgement

This work was carried out as part of the research project PACE (Protocols for Ad-hoc Collaborative Environments), funded by the UK Engineering and Physical Sciences Research Council. We are also grateful to Einar Vollset, for some helpful discussions.

References

1. C. Chiang and M. Gerla, "On-Demand Multicast in Mobile Networks", Procs., *IEEE ICMP*, pp. 260-270, Austin, 1998.
2. J. Garcia-Luna-Aceves and E. Madruga, "A Multicast Routing Protocol for Ad-hoc Networks", Procs., *Infocom*, pp.784-792, New York, 1999.
3. D.E. Cooper, P. Ezhilchelvan and I. Mitrani, "High Coverage Broadcasting for Mobile Ad Hoc Networks", Procs., *Networking'04*, Athens, 2004.
4. T. Gopalsamy, M. Shinghal, D. Panda and P. Sadayappan, "A Reliable Multicast Algorithm for Mobile Ad-hoc Networks", Procs., *ICDCS*, Vienna, 2002.
5. V. Hadzilacos and S. Tueg, "Fault-Tolerant Broadcasts and Related Problems", in *Distributed Systems* (Ed. S. Mullender), Addison-Wesley, 1993.
6. C. Ho, K. Obraczka, G. Tsudik and K. Viswanath, "Flooding for Reliable Multicast in Multihop Ad-hoc Networks", Procs., *Discrete Algorithms and Methods for Mobile Computing (DIAL-M)*, pp. 64-71, Seattle, 1999.
7. I. Mitrani, *Probabilistic Modelling*, Cambridge University Press, 1998.
8. S. Ni, Y. Tseng, Y. Chen and J. Sheu, "The Broadcast Storm Problem in a Mobile Ad-hoc Network", Procs., *ACM Int. Conf. on Mobile Computing and Networking (MOBICOM)*, pp. 151-162, Seattle, 1999.
9. E. Pagani and G.P. Rossi, "Reliable Broadcast in Mobile Multihop Packet Networks", Procs., *ACM Int. Conf. on Mobile Computing and Networking (MOBICOM)*, pp. 34-42, Budapest, 1997.
10. P. Sinha, R. Sivakumar and V. Bharagavan, "MCEDAR: Multicast Core Extraction Distributed Ad-hoc Networking", Procs., *Wireless Communications and Networking Conference (WCNC)*, pp. 1313-1317, New Orleans, 1999.
11. B. Williams and T. Camp, "Comparison of Broadcasting Techniques for Mobile Ad-hoc Networks", Procs., *ACM Int. Symp. on Mobile Ad-hoc Networking and Computing (MOBIHOC)*, Lausanne, 2002.