# MobiCom Poster Abstract: Real-World Evaluation of Ring Flooding

**Wolfgang Kiess**[a] **Andreas Tarp**[a] **Martin Mauve**[a]
*kiess@cs.uni-duesseldorf.de* *andreas.tarp@uni-duesseldorf.de* *mauve@cs.uni-duesseldorf.de*
[a]Computer Networks Research Group, University of Düsseldorf, Germany

## I. Introduction

Flooding is a basic mechanism frequently used in mobile ad-hoc networks. In its simplest form flooding is realized by letting each node rebroadcast the flooded packet exactly once. To limit the scope of a flooded data packet, the sender of the packet may use *ring flooding*. For this, the packet's time-to-live (TTL) field is initially set to $n$. As the TTL of the packet expires after $n$ hops, it only reaches all those nodes that are at most $n$ hops away from the original sender. Ring flooding is used, e.g., to distribute information only relevant in a certain area such as emergency messages in car-to-car networks, or to do an expanding ring search during route discovery as in AODV [2].

This paper presents a preliminary experimental study of ring flooding in ad-hoc networks with up to 13 nodes. It examines 1) how reliable a flooded packet reaches all nodes 2) how long it takes the packet to reach the nodes and 3) how many nodes are reachable when flooding with a certain TTL.

## II. Experiments

We have performed one indoor and one outdoor experiment with ring flooding on static multihop topologies with IEEE 802.11b equipped nodes. The parameters that have been varied are initial TTL, jitter and packet size. Jitter was used to delay the rebroadcasting for a random time from the interval [0;jitter] in order to reduce collisions. All nodes used Linux, flooding was implemented with click [1], and the packets were traced with tcpdump. One node at the corner of the network acted as packet source. Each node repeated each packet exactly once and ignored duplicates, packets were dropped on TTL expiration.

The indoor experiment comprised 10 iPAQ5550 PDAs distributed in two rows over $15 \times 50$ meters. The flooded packets had a size of 100 bytes, we chose 0 and 10 ms maximum jitter. For each jitter interval, we flooded 10000 packets divided in sequences with TTL (1,3,5,7,9) and 10000 packets divided in sequences with TTL (1,2,4,8,16). The minimum spacing between two flood attempts was 120 ms.

For the outdoor experiment, 13 nodes (10 iPAQ5550, 3 notebooks) were distributed over an area of $110 \times 145$ meters on the university campus (Figure 1(a), 'moe' was the packet source). For all experiments, we used linear ring flooding, i.e. the source increased the initial TTL from one to 13 for each successive packet and then restarted from one. Packets had a size of 200 bytes and the used maximum jitter values were 0, 5, 10, and 15 ms. For each of these values, a total of 3000 packets were flooded in 6 runs. As minimum spacing between two flood attempts, we used 60 ms and increased this for higher TTLs and jitter. During run 13, the node 'notebook1' (see Figure 1(a)) failed due to a lack of battery power, dividing the experiment in two different topologies. We therefore eliminate the affected runs 8-15, leaving for each maximum jitter value 1000 packets on each topology.

## III. Results

**Reliability** is the percentage of packets that reach all nodes in the network. Although each node in the indoor experiment was theoretically reachable with at most three hops, packets with an initial TTL $\geq 5$ achieved the highest reliability as shown in Table 1. It is obvious that high reliability comes at the cost of letting each node repeat the packet. Interesting is the behavior for packets with an initial TTL $\geq 5$ for the [0;10] ms jitter runs. One of the runs had a reliability of only 54.6%. With this run, the reliability was 97.0%, without it, it was 99.7%. We assume that this stems from temporary interference. For the outdoor experiment, increased jitter seems to have a positive effect on reliability as shown in Figure 1(b). Nevertheless, more experiments are necessary to gain a better understanding of this effect.

| TTL | 0ms jitter | [0;10] ms jitter |
|---|---|---|
| 3 | 13.7% | 14.1% |
| 4 | 98.0% | 99.1% |
| $\geq 5$ | 99.7% | 97.0% (99.7%) |

Table 1: Reliability for the indoor experiments.

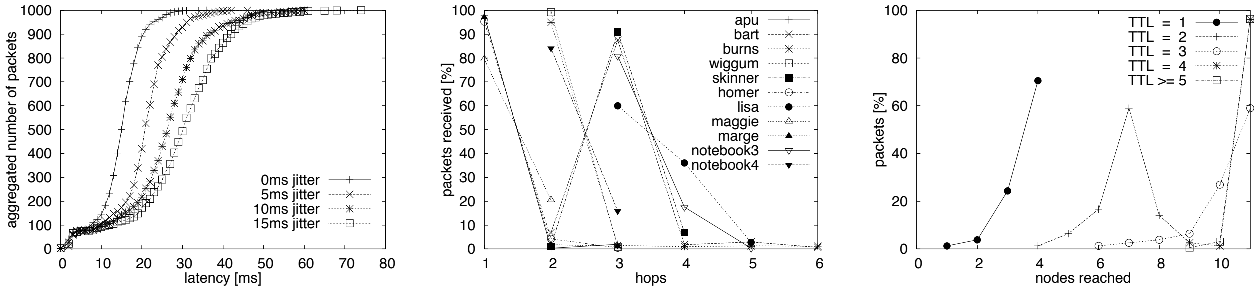


(a) Latency for all packets (b) Min. hops/node, TTL $\geq$ 6 (c) Number of nodes reached

Figure 2: Latency and neighborhood stability ([0;5] ms jitter) for topology two of the outdoor experiment.

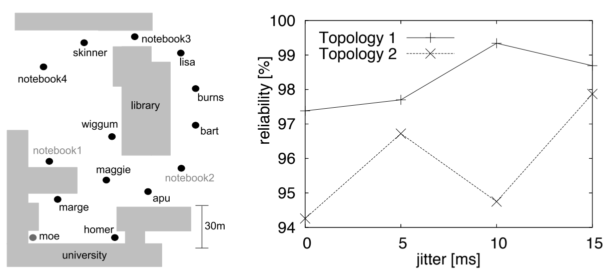


(a) Node setup (b) Reliability, TTL $\geq$ 6

Figure 1: The outdoor experiment.

**Latency** is the time from the initial broadcast until the last node receives the packet. For the outdoor experiments, the latency is shown in Figure 2(a) for topology two. The highest latency in the experiment, over 140ms occurred on topology one (not shown here). The reason for this was a rebroadcasting delay of 120 ms in 'notebook1'. Besides that, the highest latencies for both topologies were 42, 46, 67 and 74 ms for the increasing maximum jitter values.

**Received copies** are the number of copies a node receives of a flooded packet. As nodes at the border of the network do not receive packets with low initial TTL, we only consider packets with a sufficiently high TTL. For the outdoor experiment, a node received on average 3.8 copies on topology one and 3.3 copies on topology two (TTL $\geq$ 6). For the indoor network, the number of copies per node was 3.6 (TTL $\geq$ 5). This is interesting in relation to those copies sent but not received by any other node. While indoor runs with [0;10] ms jitter lost only 0.3% of all send events, 0.9% were lost for the runs without jitter. For the outdoor experiment (averaged over both topologies), this loss decreased slightly from 1.6% over 1.4% and 1.3% to 1.0% with increasing jitter. Due to the dependence on jitter, we suspect that this is an effect of collisions.

**Neighborhood stability** denotes how the allocation of a node to a *n*-hop neighborhood fluctuates. This is shown in Figure 2(b) giving the minimum number of hops it took the flooded packets to reach the nodes. Six of the nodes were firmly attached to a certain neighborhood with reception rates of over 90%. Four other nodes received between 75 and 90% over a certain number of hops. One node ('lisa') received 62% of the packets over three and 38% over four or more hops. Figure 2(c) shows how this influenced node reachability: while 60% of the packets flooded with TTL 2 reached exactly seven nodes, a 2-hop ring flood might reach as few as four or as much as nine nodes. A similar behaviour can be observed for the other TTL values.

## IV. Conclusions and Future work

Our experiments showed that flooding with the network diameter *d* is not enough if all nodes should be reached: in our setup, flooding was most successful for TTL $\geq d + 2$. Furthermore, the number of nodes reachable with a *n*-hop flood varied from attempt to attempt even in our static networks. Thus the definition of the often used *n*-hop neighborhood should be revised. In the next step, we plan to examine flooding in the presence of other traffic. Furthermore, repeating the experiment in a network simulator can provide a better understanding of neighborhood stability.

## References

[1] E. Kohler, R. Morris, B. Chen, and J. Jannotti. The Click Modular Router. *ACM Transactions on Computer Systems*, 18(3):263–297, August 2000.

[2] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc on-demand distance vector (AODV) routing. RFC 3561, July 2003.