

Mobile Ad Hoc Networking Working Group
INTERNET DRAFT
8 July 2006

Charles E. Perkins
Nokia Research Center
Thomas Clausen,
Philippe Jacquet
INRIA Rocquencourt, France

Multicast With Minimal Congestion Using Connected Dominating Sets
[draft-perkins-manet-smurf-00.txt](#)

Status of This Memo

This document is a submission by the Mobile Ad Hoc Networking Working Group of the Internet Engineering Task Force (IETF). Comments should be submitted to the manet@ietf.org mailing list.

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

Flooding is useful for many reasons in ad hoc networks, but causes congestion that can block traffic. Selecting multi-point relays (MPRs) reduces congestion due to flooding, but has several disadvantages. It is possible to select a subset of the MPRs to do the same job more effectively. In this document, signaling is specified for identifying a relatively small connected dominating set that can still manage the flooding operation.

1. Introduction

Flooding is useful for many reasons in ad hoc networks, but causes congestion that can block traffic within such a network. The goal of this document is to specify a flooding algorithm that greatly reduces the number of packet retransmissions needed to flood the whole network. Since there are many fewer retransmissions, the amount of congestion due to flooding will be greatly reduced.

In order to reduce the total number of retransmissions, we can reduce the number of nodes that perform the retransmissions, as long as all nodes still receive at least one transmission of the flooded packet. If a particular node can determine its one-hop and two-hop neighborhoods, then it can identify a subset of its neighboring nodes that still be able to reach every node of the two-hop neighborhood. The nodes in such a subset are called "multi-point relays" (MPRs) in the OLSR specification [7]. There are various ways to pick enough MPRs so that, when each MPR retransmits a packet, the two-hop neighborhood will be completely covered. These neighboring MPRs will themselves have selected some of their other neighbors to be MPRs. When the neighboring MPRs retransmit, then some of their neighbors will again retransmit, until eventually the whole network is covered.

Selecting multi-point relays MPRs, as specified in OLSR [7] is beneficial to reduce congestion due to flooding, but has several disadvantages. Another related concept is that of a "connected dominating set" (CDS). There are many distributed algorithms for computing a CDS, with a variety of interesting properties. The set of MPRs that retransmit a packet originated from a source node within an ad hoc network is also a CDS. Otherwise, the set of MPRs would not be able to cover the network by retransmission hop-by-hop.

In this document, signaling is specified for identifying a relatively small connected dominating set that can still manage the flooding operation.

The names of the multicast groups used for flooding are given as "ALL_IPv4_MANET_NODES" and "ALL_IPv6_MANET_NODES" [5]. The backbone nodes selected for flooding packets can, however, be used for

flooding to any designated multicast address as long as the members of that multicast group form a connected subgraph of the ad hoc network. Specifications for future multicast groups intended for flooding or very dense dissemination should specifically indicate that their members should communicate across the SMURF backbone. Future documents may offer improved algorithms for the purpose of identifying backbone nodes, while still utilizing the protocol in this document for periodic updates to the necessary neighborhood information.

DISCUSSION: should there be a version field to handle unforeseen future incompatibilities?

[2.](#) SMURF Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [2].

This section defines other terminology that is not already defined in [3].

MPR

A "multi-point relay" -- that is, a neighboring node that can relay flooded packets, thus assisting with the dissemination of those packets throughout the network

broadcast

Transmit to the IPv4 Limited Broadcast address, 255.255.255.255, or else for IPv6 to the link-local AllNodesMulticast address.

neighborhood, one-hop neighborhood

The neighborhood (also called the one-hop neighborhood) of a node is the set of nodes which are directly reachable by that node.

two-hop neighborhood

The two-hop neighborhood of a node is the set of nodes which are either directly reachable by that node, or directly reachable by one of the neighbors of that node.

dominating set

A subset of a graph is a dominating set if every node in the graph is either a member of the subset, or else a neighbor of at least one node in the subset.

connected set

A set with more than one element is connected if every node in the subset is a neighbor of some other node in the subset. A set with fewer than two elements is connected.

CDS (connected dominating set)

A subset of a graph is a CDS if it is a dominating set, and it is connected.

flooding backbone

The flooding backbone of an ad hoc network is a CDS which is used for flooding packets throughout the network, considering every node in the ad hoc network as a node in a graph with edges defined by the (one-hop) communication links that are available between the nodes in the ad hoc network.

backbone

in this document, ``backbone'' always means the flooding backbone

[3. Overview](#)

Each node listens for advertisements from its neighbors. From the advertisements, the node updates its stored information about its

one-hop neighborhood and its two-hop neighborhood. The sender of the advertisement is a member of the node's one-hop neighborhood, and the sender's neighbors are either members of the node's one-hop neighborhood or its two-hop neighborhood.

There are two kinds of periodic local topology advertisements, which are used to enable each node to construct a representation of its two-hop neighborhood and subsequently select MPRs. The first is called the "full advertisement", which has complete lists of all the necessary data. The second is called the "incremental advertisement", which only has information that has changed since the last full advertisement. It is expected that the incremental advertisement will be much smaller than the full advertisement.

A node's neighborhood is considered to be composed of three mutually exclusive subsets:

- undistinguished first hop neighbors
- MPRs
- the node itself

Likewise, a node's two hop neighborhood is considered to be composed of four mutually exclusive subsets:

- second hop neighbors
- undistinguished first hop neighbors

- MPRs
- the node itself

If a node is described as being in any one of these subsets, then by convention and agreement it does not belong to any of the other subsets. This eliminates overlap of information in the local topology advertisements.

The following information is present in full advertisements:

- advertisement sequence number
- reluctance for offering service as a backbone node
- first hop neighbors
- MPRs
- flags

The following information is present in incremental advertisements:

- advertisement sequence number of the last full advertisement
- reluctance for offering service as a backbone node
- addresses of new first hop neighbors
- addresses of new MPRs
- addresses of lost nodes
- flags

Any address in the lost node list has to be deleted from the neighbor list or the MPR list of the last full advertisement, and the nodes in these two lists reclassified if they were classified differently in the last full advertisement.

A neighboring node cannot be considered for selection as an MPR until the link to that node has been verified as bidirectional. For each neighbor, two conditions must be verified:

1. an advertisement has been recently heard from the neighbor
2. the node itself must appear as a member of the one-hop neighbors reported by that neighbor.

Once each node has constructed its two-hop neighborhood, it uses any reasonable algorithm to identify a subset of neighbors which are to be designated as MPRs. One algorithm that works well is the so-called "greedy algorithm":

1. Initialize the set of uncovered two-hop neighbors to contain all known two-hop neighbors
2. Pick the neighbor that is within range of the largest number of uncovered nodes in the two-hop neighborhood.

3. Mark the two-hop neighbors reachable by that neighbor's local broadcasts as already covered.
4. If there are no more uncovered two-hop neighbors, stop.
5. Otherwise, iterate steps (2) -- (4).

Once the nodes have identified the MPRs in their neighborhoods, each node must also determine whether to prune itself from the flooding backbone (i.e., the connected dominating set of nodes which will take responsibility for disseminating flooded data).

In order for a node to prune itself to the flooding backbone, it uses the following algorithm [1]:

1. It forms a "CDS identifier" by appending its IP address to its last advertised "reluctance" value.
2. If it is the node with a CDS identifier lower than any of its neighbors, then it must continue to participate as a member of the flooding backbone.
3. Otherwise, if it not a MPR of its neighbor with the lowest CDS identifier, then it may prune itself from the flooding backbone.

A node advertises the status of its decision to prune itself from the CDS backbone in its advertisement messages.

4. Full Advertisement

The Full Advertisement has the following format:

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-----+-----+-----+-----+
                                |   Type   |   Length   |
+-----+-----+-----+-----+
|           Sequence #           |Rel| #1-hop | #MPRs | reserved|C|
+-----+-----+-----+-----+
:           List of IP addresses of 1-hop neighbors           :
+-----+-----+-----+-----+
:           List of IP addresses of MPR neighbors             :
+-----+-----+-----+-----+

```

Type

Type == 1.

Length

Length of the full advertisement data, in bytes, including all fields except for the Type and Length.

Sequence #

Advertisement sequence number

Rel

Reluctance (2 bits). A node selects a low value for its reluctance, if it is willing to be identified as an MPR, and thus potentially as a member of the connected dominating set.

#1-hop

Number of one-hop neighbors

#MPRs

Number of one-hop neighbors reclassified as MPRs

'C'

The 'C' flag is set if the advertising node has selected itself as a member of the connecting dominating set for flooding

1-hop list

List of IP addresses of 1-hop neighbors (#1-hop total addresses)

MPR list

List of IP addresses of MPR neighbors (#MPRs total addresses)

The IP header TTL for the Full Advertisement packet MUST be 1.

DISCUSSION: is 255 better?

DISCUSSION: Is the Length field needed, or should we spend the bits on enlarging the other fields?

5. Incremental Advertisement

The Incremental Advertisement has the following format:

Internet Draft

Low-Congestion Flooding

8 July 2006

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                     +---+---+---+---+---+---+---+---+---+---+
                                     |   Type   |   Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Sequence #           |Rel| #1-hop|#MPRs| #lost |reservd|C|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:           List of IP addresses of 1-hop neighbors           :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:           List of IP addresses of MPR neighbors             :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:           List of IP addresses of lost neighbors            :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type

Type == 2.

Length

Length of the incremental advertisement data, in bytes,
including

Sequence #

Advertisement sequence number

Rel

Reluctance (2 bits). A node selects a low value for its
reluctance, if it is willing to be identified as an MPR, and
thus potentially as a member of the connected dominating set.

#1-hop

Number of new one-hop neighbors

#MPRs

Number of new nodes classified as MPRs

#lost

Number of nodes no longer in the one-hop neighborhood (and, thus, not possibly MPRs) as shown in the last full advertisement

Perkins

Expires 8 January 2006

[Page 8]

Internet Draft

Low-Congestion Flooding

8 July 2006

'C'

The 'C' flag is set if the advertising node has selected itself as a member of the connecting dominating set for flooding

1-hop list

List of IP addresses of new 1-hop neighbors (#1-hop total addresses)

MPR list

List of IP addresses of new MPR neighbors (#MPRs total addresses)

lost list

List of IP addresses of former neighbors that have not exhibited local connectivity recently (#lost total addresses)

The IP header TTL for the Incremental Advertisement MUST be 1.

DISCUSSION: is 255 better?

DISCUSSION: Is the Length field needed, or should we spend the bits on enlarging the other fields?

6. Security Considerations

This draft specifies a general mechanism for identifying a flooding backbone in an ad hoc network. It does not make any provision for securing the contents of the flooded data, either to protect against tampering or to protect against unauthorized inspection of the data.

It does not make any provision for ensuring that the nodes of the flooding backbone are operating as expected. It does not make any provision to guarantee the correctness of information presented within the advertisement messages. The use of local broadcast for the advertisement messages precludes the use of point-to-point security associations such as might be used for constructing Authentication Header digests.

In the future, if further security measures are to be undertaken, one could attempt to utilize public keys for the purpose of authenticating the flooded data, as long as the public keys were known to be properly associated with the node in question (to avoid man-in-the-middle attacks).

[7.](#) IANA considerations

New multicast addresses have to be allocated. A new UDP port number has to be allocated for the neighborhood messages. ICMP is another possibility.

[8.](#) Acknowledgments

Thomas Clausen and Emmanuel Baccelli were frequent contributors to the conversation which inspired the production of this document. The algorithm for identifying the flooding backbone was published as Research Report 4597 [[1](#)] by Cedric Adjih et al.

References

- [1] Cedric Adjih, Philippe Jacquet, and Laurent Veinnot. Computing Connected Dominating Sets with Multipoint Relays. Technical report, Institut National de Recherche en Informatique et en Automatique (INRIA). INRIA RR 4597, October 2002.
- [2] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. Request for Comments (Best Current Practice) [2119](#), Internet Engineering Task Force, March 1997.

- [3] Ed. J. Manner and Ed. M. Kojo. Mobility Related Terminology. Request for Comments (Informational) [3753](#), Internet Engineering Task Force, June 2004.
- [4] C. Perkins. Minimal Encapsulation within IP. Request for Comments (Proposed Standard) [2004](#), Internet Engineering Task Force, October 1996.
- [5] Charles E. Perkins. IP Flooding in Ad hoc Networks (Work In Progress). Internet Draft, Internet Engineering Task Force. [draft-perkins-manet-bcast-02.txt](#), June 2005.
- [6] J. Postel. Internet Protocol. Request for Comments (Standard) [791](#), Internet Engineering Task Force, September 1981.
- [7] Ed. T. Clausen and Ed. P. Jacquet. Optimized Link State Routing Protocol (OLSR). Request for Comments (Experimental Protocol) [3626](#), Internet Engineering Task Force, June 2004.

Author's Addresses

Questions about this memo can be directed to:

Perkins Expires 8 January 2006 [Page 10]

Internet Draft Low-Congestion Flooding 8 July 2006

Charles E. Perkins
Networking Technology Laboratory / Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94303
+1 650 625 2986
+1 650 625-2502 (fax)
Charles.Perkins@nokia.com

Thomas Heide Clausen
INRIA Rocquencourt, BP 105,
78153 Le Chesnay Cedex
France
+33 1 3963 5133
Thomas.Clausen@inria.fr

Full Copyright Statement

Full Copyright Statement

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.