Mobile Ad Hoc Networking Working Group
INTERNET DRAFT
29 June 2005

Charles E. Perkins
Nokia Research Center

IP Flooding in Ad hoc Mobile Networks
draft-perkins-manet-bcast-02.txt

Status of This Memo

Abstract

An ad hoc mobile network is a collection of nodes, each of which communicates over wireless channels and is capable of movement. Nodes participating in such an ad hoc network communicate on a peer-to-peer basis. Flooding is often a desired form of communication in these networks, as it can enable both the dissemination of control information and the delivery of data packets. This document describes a method for sending packets to every node in an ad hoc network.

1. Introduction

This document makes a particular specification for a well-known flooding algorithm, as it can be used to disseminate IP packets across ad hoc networks. For the well-known flooding algorithm to work, the nodes flooding packets must ensure that each distinct packet that they send is uniquely identifiable, at least during the expected time taken for the flooded packet to disseminate though an ad hoc network. In this document, the method used for insuring uniqueness depends upon whether an IPv4 or IPv6 packet is being transmitted.

In IPv4, there are two kinds of broadcast address, and it seems that neither one of them is likely to present a good choice for the IP address to be used for network-layer flooding. The IPv4 address for "limited broadcast" is 255.255.255.255, and is not supposed to be forwarded. Since the nodes in an ad hoc network are asked to forward the flooded packets, the limited broadcast address should not used for network-layer flooding. The other available choice, the "directed broadcast" address, would presume a choice of routing prefix for the ad hoc network and thus is not a reasonable choice.

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 [1].

2. Applicability Statement

The specification in this document is immediately useful for for network-layer flooding -- i.e., when the TTL or hop count in the IP header is initialized to a value greater than 1. Note that there is another useful alternative for flooding in an ad hoc network. Namely, it is often the case information to be flooded requires the attention of some application at every node receiving the flooded information, before that node would further disseminate the flooded information. For instance, a routing protocol might need to carry out several kinds of operations before allowing the packet to be retransmitted. In those cases, it may be quite appropriate (or even preferable) to use the limited broadcast address, 255.255.255.255, with the understanding that the packet will not be retransmitted at the network layer by any of the node's neighbors. This specification can be used to guarantee uniqueness for such packets even when the application makes no modification to the payload before it it is retransmitted.

3. Flooding

In this specification, new multicast groups for flooding to all nodes of an ad hoc network are specified for use with network-layer flooding. These multicast groups are specified to contain all nodes of a contiguous ad hoc network, so that packets transmitted to the multicast address associated with the group will be delivered to all nodes as desired. In other words, any node that is reachable, is automatically granted membership in these multicast groups. For IPv6, the multicast address is specified to be "site-local". The names of the multicast groups are given as "ALL_IPv4_MANET_NODES" (TBD) and "ALL_IPv6_MANET_NODES" (TBD). This document does not specify transmissions to any directed broadcast address.

Every node maintains a list of those flooded packets which have already been received and retransmitted. The list contains, for each distinct flooded packet received, a value called the Flooded Packet Identifier (FPI). For IPv4, this FPI is composed of the source IP address, the IP ident value, and the fragment offset values obtained from the IP header of the flooded packet. For IPv6, the FPI is calculated as specified in section 4.

When a node receives a flooded packet, it checks its list for the FPI of the flooded packet's IP header [3]. If there is such a list entry with matching FPI, the node silently discards the flooded packet since it has already been received and forwarded. The node then checks to see whether it is enabled for retransmitting flooded packets. By default, all nodes in the ad hoc network are so enabled; however, this is not required (see section 6) and may be changed by configuration or by protocol action. If the node is not enabled for retransmitting flooded packets, it takes no further action. If there is no existing list entry containing the same FPI, and if the node has been enabled to forward flooded packets, the node retransmits the packet.

List entries SHOULD be kept for at least BROADCAST_RECORD_TIME before the node expunges the record. BROADCAST_RECORD_TIME is a configurable parameter, but it MUST be at least equal to NET_TRAVERSAL_TIME.

4. FPI computation for IPv6

DISCUSSION QUESTION: Is another cryptographic function better, or good enough but easier?

To obtain the FPI for IPv6 packets, a node uses MD5 [4] to perform the following calculation for the incoming flooded packet:

FPI = MD5 (IPv6 packet data).

The IP packet data includes all unpredictable IPv6 headers and extensions [2], as well as any higher-level protocol data. The source node for each flooded packet MUST ensure that this FPI is distinct from the FPI from every other flooded packet which the node has transmitted during the last BROADCAST_RECORD_TIME. In the unlikely event that the FPI value is identical to some such recently transmitted packet, the source node MUST add a Unique Identifier Destination Option to the flooded packet (see section 5).

DISCUSSION QUESTION: Should the same digestifying procedure be specified for IPv4 also?

5. Unique Identifier Destination Option

DISCUSSION QUESTION: The IPv6 fragment header can serve this purpose very well, at a cost of only another two bytes. Should that be used instead of the Unique Identifier option?

The Unique Identifier option is encoded in type-length-value (TLV) format as follows:

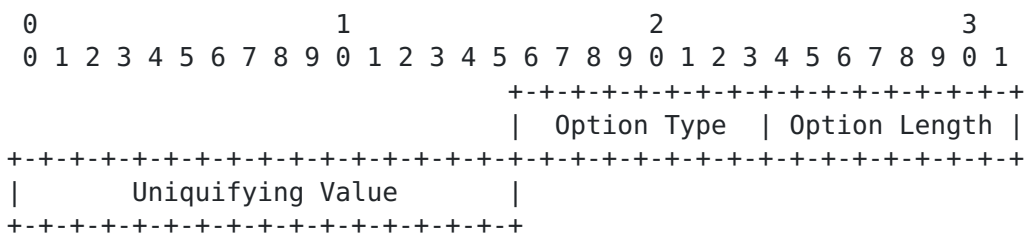

Option Type

TBD

Option Length

2

Uniquifying Value

The 16-bit Uniquifying Value is chosen to make the flooded packet FPI computation different than that for any other flooded packet from the same source node.

The Unique Identifier MUST be placed in the Destination Options before the Routing Header (and, thus, before the fragment header). This allows proper handling by all intermediate forwarding nodes.

6. Selective Retransmission for Flooded Packets

By default, each node in the ad hoc network is enabled to retransmit each distinct flooded packet that it receives. However, in some cases, there may be additional control signaling in place that is used to reduce the number of nodes that perform this retransmission, in order to reduce the overall bandwidth consumption and congestion which can be caused by excessive flooding. This document does not specify any such control protocol to disable or enable such node selection. However, an ad hoc network which employs such a node selection protocol can still be compliant with the flooding protocol specified in this document.

7. Configuration Parameters

This section gives default values for some important values associated with flooding operations. Mobile nodes in particular ad hoc networks may wish to change certain of the parameters, in particular the NET_DIAMETER and NODE_TRAVERSAL values. Choice of these parameters may affect the robustness of the flooding operation.

Parameter Name	Value
BROADCAST_RECORD_TIME	2 * NET_TRAVERSAL_TIME
NET_DIAMETER	35
NODE_TRAVERSAL_TIME	40 milliseconds
NET_TRAVERSAL_TIME	3 * NODE_TRAVERSAL_TIME * NET_DIAMETER / 2

NET_DIAMETER measures the maximum possible number of hops between two nodes in the network. NODE_TRAVERSAL_TIME is a conservative estimate of the average one hop traversal time for packets and should include queuing delays, interrupt processing times, medium access delays, and propagation delays. NET_TRAVERSAL_TIME is a conservative estimate of how long it should take for a message to traverse the entire ad hoc network.

8. Security Considerations

This draft specifies a general mechanism for flooding packets 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.

9. Acknowledgments

This flooding method is a codification of a well known algorithm which has been assumed for general use in various ad hoc protocols. Thus, the protocol specification in this draft should be considered the joint work of many engineers who have worked on producing ad hoc network protocols.

References

[1] 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.

[2] S. Kent and R. Atkinson. IP Authentication Header. Request for Comments (Proposed Standard) 2402, Internet Engineering Task Force, November 1998.

[3] J. Postel. Internet Protocol. Request for Comments (Standard) 791, Internet Engineering Task Force, September 1981.

[4] R. Rivest. The MD5 Message-Digest Algorithm. Request for Comments (Informational) 1321, Internet Engineering Task Force, April 1992.

A. Changes since the last revision

- Added applicability section.
- Changed author list
- Added discussion points
- Moved some text out of the introduction into the main body of the specification.
- Othe minor editorial corrections.

Author's Addresses

Questions about this memo can be directed to:

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

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement