

Network Working Group
Request for Comments: 5148
Category: Informational

T. Clausen
LIX, Ecole Polytechnique, France
C. Dearlove
BAE Systems Advanced Technology Centre
B. Adamson
U.S. Naval Research Laboratory
February 2008

Jitter Considerations in Mobile Ad Hoc Networks (MANETs)

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document provides recommendations for jittering (randomly modifying timing) of control traffic transmissions in Mobile Ad hoc NETWORK (MANET) routing protocols to reduce the probability of transmission collisions.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Applicability Statement	4
4. Protocol Overview and Functioning	4
5. Jitter	5
5.1. Periodic Message Generation	5
5.2. Externally Triggered Message Generation	6
5.3. Message Forwarding	7
5.4. Maximum Jitter Determination	8
6. Security Considerations	9
7. References	10
7.1. Normative References	10
7.2. Informative References	10
Appendix A. Acknowledgements	11

1. Introduction

In a wireless network, simultaneous packet transmission by nearby nodes is often undesirable. This is because any resulting collision between these packets may cause a receiving node to fail to receive some or all of these packets. This is a physical problem, which occurs before packets can be inserted into the receiver queue. Depending on the characteristics of the medium access control and other lower layer mechanisms, in particular whether retransmission of unacknowledged packets is supported, this may cause at best increased delay, and at worst complete packet loss. In some instances, these problems can be solved in these lower layers, but in other instances, some help at the network and higher layers is necessary.

This document considers the case when that help is required, and provides recommendations for using jitter (randomly varying timing) to provide it. It is possible that the techniques described here could be implemented either by IP protocols designed for wireless networks or in conjunction with lower-layer mechanisms.

The problems of simultaneous packet transmissions are amplified if any of the following features are present in a protocol:

Regularly scheduled messages - If two nodes generate packets containing regularly scheduled messages of the same type at the same time, and if, as is typical, they are using the same message interval, all further transmissions of these messages will thus also be at the same time. Note that the following mechanisms may make this a likely occurrence.

Event-triggered messages - If nodes respond to changes in their circumstances, in particular changes in their neighborhood, with an immediate message generation and transmission, then two nearby nodes that respond to the same change will transmit messages simultaneously.

Schedule reset - When a node sends an event-triggered message of a type that is usually regularly scheduled, then there is no apparent reason why it should not restart its corresponding message schedule. This may result in nodes responding to the same change also sending future messages simultaneously.

Forwarding - If nodes forward messages they receive from other nodes, then nearby nodes will commonly receive and forward the same message. If forwarding is performed immediately, then the resulting packet transmissions may interfere with each other.

A possible solution to these problems is to employ jitter, a deliberate random variation in timing. Such jitter is employed in e.g., [2], [3], and [4], in which transmission intervals for regularly scheduled messages are reduced by a small, bounded and random amount in order to desynchronize transmitters and thereby avoid overloading the transmission medium as well as receivers. This document discusses and provides recommendations for applying jitter to control packet transmissions in Mobile Ad hoc NETWORKS (MANETs), with the purpose of avoiding collisions, with particular reference to the features listed above.

2. Terminology

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [1].

Additionally, this document uses the following terminology:

Node - A MANET router that implements a message sending protocol.

MANET interface - A network device participating in a MANET. A node may have one or more MANET interfaces.

Message - An entity carrying protocol information intended for exchange between nodes. Messages are transmitted over MANET interfaces embedded in packets.

Packet - An entity embedding zero or more messages for transmission over a MANET interface of the node.

Transmission - A packet being sent over a MANET interface of the node. A transmission can be due to either a message being generated or a message being forwarded.

Generation - Creation of a new message (rather than a received and forwarded message) for transmission over one or more MANET interfaces of the node. Typically, a node will generate messages based on a message schedule (periodic or otherwise) or as a response to changes in circumstances.

Forwarding - Retransmission of a received message (whether modified or unchanged) over one or more MANET interfaces of the node.

Collision - A specific instance of interference, where two or more nodes transmit a packet at the same time and within the same signal space (at the same frequency and/or encoding) such that

another, closely located, node that should receive and decode these packets instead fails to do so, and loses one or more of the packets.

3. Applicability Statement

The mechanisms described in this document are applicable to the control messages of any MANET protocol in which simultaneous transmissions by different nodes are undesirable, and that contains mechanisms, such as periodic control message transmission, triggered control message transmission, or control message forwarding, which either make a simultaneous transmission more likely, or cause one to be repeated when it occurs. This particularly applies to protocols using broadcast transmissions in wireless networks, where proactive MANET routing protocols such as [5] employ scheduled messages, where reactive MANET routing protocols such as [6] employ event-triggered messages, and where both employ message forwarding.

These mechanisms are intended for application where the underlying medium access control and lower layers do not provide effective mechanisms to avoid such collisions. Where these layers do provide effective mechanisms, the recommendations of this document are not needed.

The approach described in this document uses random variations in timing to achieve a reduction in collisions. Alternatives using, for example, pseudo-random variation based on node identity, may be considered, but are not discussed by this document.

Any protocol based on [7] and using the message forwarding mechanism facilitated by that structure is a particular candidate for application of at least some of these mechanisms.

The document has been generalized from the jitter mechanism used in the proactive MANET routing protocol OLSR (the Optimized Link State Routing Protocol) [5].

4. Protocol Overview and Functioning

This document provides recommendations for message transmission (and retransmission) that may be used by MANET routing protocols. It may also be used by other protocols that employ a periodic or triggered message schedule running over wireless interfaces. Using such simultaneous message transmissions from two (or more) adjacent nodes may cause delays, packet losses, and other problems. Any protocol using jitter as outlined here must specify its precise usage insofar as is necessary for interoperability.

5. Jitter

In order to prevent nodes in a MANET from simultaneous transmission, whilst retaining the MANET characteristic of maximum node autonomy, a randomization of the transmission time of packets by nodes, known as jitter, SHOULD be employed. Three jitter mechanisms, which target different aspects of this problem, SHOULD be employed, with the aim of reducing the likelihood of simultaneous transmission, and, if it occurs, preventing it from continuing.

Three cases exist:

- o Periodic message generation;
- o Externally triggered message generation;
- o Message forwarding.

For the first of these cases, jitter is used to reduce the interval between successive message transmission by a random amount; for the latter two cases, jitter is used to delay a message being generated or forwarded by a random amount.

Each of these cases uses a parameter, denoted MAXJITTER, for the maximum timing variation that it introduces. If more than one of these cases is used by a protocol, it MAY use the same or a different value of MAXJITTER for each case. It also MAY use the same or different values of MAXJITTER according to message type, and under different circumstances -- in particular if other parameters (such as message interval) vary.

Issues relating to the value of MAXJITTER are considered in Section 5.4.

5.1. Periodic Message Generation

When a node generates a message periodically, two successive messages will be separated by a well-defined interval, denoted MESSAGE_INTERVAL. A node MAY maintain more than one such interval, e.g., for different message types or in different circumstances (such as backing off transmissions to avoid congestion). Jitter SHOULD be applied by reducing this delay by a random amount, so that the delay between consecutive transmissions of messages of the same type is equal to (MESSAGE_INTERVAL - jitter), where jitter is the random value.

Subtraction of the random value from the message interval ensures that the message interval never exceeds MESSAGE_INTERVAL, and does

not adversely affect timeouts or other mechanisms that may be based on message late arrival or failure to arrive. By basing the message transmission time on the previous transmission time, rather than by jittering a fixed clock, nodes can become completely desynchronized, which minimizes their probability of repeated collisions. This is particularly useful when combined with externally triggered message generation and rescheduling.

The jitter value SHOULD be generated uniformly in an interval between zero and MAXJITTER.

Note that a node will know its own MESSAGE_INTERVAL value and can readily ensure that any MAXJITTER value used satisfies the conditions in Section 5.4.

5.2. Externally Triggered Message Generation

An internal or external condition or event may trigger message generation by a node. Depending upon the protocol, this condition may trigger generation of a single message (including, but not limited to, an acknowledgement message), initiation of a new periodic message schedule, or rescheduling of existing periodic messaging. Collision between externally triggered messages is made more likely if more than one node is likely to respond to the same event. To reduce this likelihood, an externally triggered message SHOULD be jittered by delaying it by a random duration; an internally triggered message MAY also be so jittered if appropriate. This delay SHOULD be generated uniformly in an interval between zero and MAXJITTER. If periodically transmitted messages are rescheduled, then this SHOULD be based on this delayed time, with subsequent messages treated as described in Section 5.1.

When messages are triggered, whether or not they are also periodically transmitted, a protocol MAY impose a minimum interval between messages of the same type, denoted MESSAGE_MIN_INTERVAL. In the case that such an interval is not required, MESSAGE_MIN_INTERVAL is considered to be zero. When MESSAGE_MIN_INTERVAL is non-zero, it is however appropriate to also allow this interval to be reduced by jitter. Thus, when a message is transmitted, the next message is allowed after a time (MESSAGE_MIN_INTERVAL - jitter). This jitter SHOULD be generated uniformly in an interval between zero and MAXJITTER (using a value of MAXJITTER appropriate to periodic message transmission).

It might appear counterintuitive to have a defined MESSAGE_MIN_INTERVAL, yet allow this to be reduced by jittering. For periodic messages, setting MESSAGE_INTERVAL, MAXJITTER and MESSAGE_MIN_INTERVAL such that (MESSAGE_INTERVAL-MAXJITTER) >

MESSAGE_MIN_INTERVAL would ensure at least MESSAGE_MIN_INTERVAL would elapse between two subsequent message transmissions. In a highly dynamic network with triggered messages, however, external circumstances might be such that external triggers are more frequent than MESSAGE_MIN_INTERVAL, effectively making MESSAGE_MIN_INTERVAL take the role of MESSAGE_INTERVAL as the "default" interval at which messages are transmitted. Thus, in order to avoid synchronization in this highly dynamic case, jittering SHOULD be applied to MESSAGE_MIN_INTERVAL. This also permits MESSAGE_MIN_INTERVAL to equal MESSAGE_INTERVAL, even when jitter is used.

When a triggered message is delayed by jitter, the node MAY also postpone generation of the triggered message. If a node is then triggered to generate a message of the same type while waiting, it can generate a single message. If however the node generates a message when it is triggered, and then receives a another trigger while waiting to send that message, then the appropriate action to take is protocol specific (typically to discard the earlier message or to transmit both, possibly modifying timing to maintain message order).

5.3. Message Forwarding

When a node forwards a message, it SHOULD be jittered by delaying it by a random duration. This delay SHOULD be generated uniformly in an interval between zero and MAXJITTER.

Unlike the cases of periodically generated and externally triggered messages, a node is not automatically aware of the message originator's value of MESSAGE_INTERVAL, which is required to select a value of MAXJITTER that is known to be valid. This may require prior agreement as to the value (or minimum value) of MESSAGE_INTERVAL, may be by inclusion in the message of MESSAGE_INTERVAL (the time until the next relevant message, rather than the time since the last message) or be by any other protocol specific mechanism, which may include estimation of the value of MESSAGE_INTERVAL based on received message times.

For several possible reasons (differing parameters, message rescheduling, extreme random values), a node may receive a message while still waiting to forward an earlier message of the same type originating from the same node. This is possible without jitter, but may occur more often with it. The appropriate action to take is protocol-specific (typically, to discard the earlier message or to forward both, possibly modifying timing to maintain message order).

In many cases, including [5] and protocols using the full functionality of [7], messages are transmitted hop-by-hop in

potentially multi-message packets, and some or all of those messages may need to be forwarded. For efficiency, this SHOULD be in a single packet, and hence the forwarding jitter of all messages received in a single packet SHOULD be the same. (This also requires that a single value of MAXJITTER is used in this case.) For this to have the intended uniform distribution, it is necessary to choose a single random jitter for all messages. It is not appropriate to give each message a random jitter and then to use the smallest of these jitter values, as that produces a jitter with a non-uniform distribution and a reduced mean value.

In addition, the protocol MAY permit control messages received in different packets to be combined, possibly also with locally generated control messages (periodically generated or triggered), as supported by [7]. However, in this case, the purpose of the jitter will be accomplished by choosing any of the independently scheduled times for these events as the single forwarding time; this may have to be the earliest time to achieve all constraints. This is because without combining messages, a transmission would be due at this time anyway.

5.4. Maximum Jitter Determination

In considering how the maximum jitter (one or more instances of parameter MAXJITTER) may be determined, the following points may be noted:

- o While jitter may resolve the problem of simultaneous transmissions, the timing changes (in particular the delays) it introduces will otherwise typically have a negative impact on a well-designed protocol. Thus, MAXJITTER SHOULD always be minimized, subject to acceptably achieving its intent.
- o When messages are periodically generated, all of the following that are relevant apply to each instance of MAXJITTER:
 - * it MUST NOT be negative;
 - * it MUST NOT be greater than MESSAGE_INTERVAL/2;
 - * it SHOULD NOT be greater than MESSAGE_INTERVAL/4.
- o If MESSAGE_MIN_INTERVAL > 0, then:
 - * MAXJITTER MUST NOT be greater than MESSAGE_MIN_INTERVAL;
 - * MAXJITTER SHOULD NOT be greater than MESSAGE_MIN_INTERVAL/2.

- o As well as the decision as to whether to use jitter being dependent on the medium access control and lower layers, the selection of the MAXJITTER parameter SHOULD be appropriate to those mechanisms. For example, MAXJITTER should be significantly greater than (e.g., an order of magnitude greater than) any medium access control frame period.
- o As jitter is intended to reduce collisions, greater jitter, i.e., an increased value of MAXJITTER, is appropriate when the chance of collisions is greater. This is particularly the case with increased node density, which is significant relative to (the square of) the interference range rather than useful signal range.
- o The choice of MAXJITTER used when forwarding messages MAY also take into account the expected number of times that the message may be sequentially forwarded, up to the network diameter in hops, so that the maximum accumulated delay is bounded.

6. Security Considerations

This document provides recommendations for mechanisms to be used in protocols; full security considerations are to be provided by those protocols, rather than in this document.

It may however be noted that introduction of random timing by these recommendations may provide some security advantage to such a protocol in that it makes the prediction of transmission times, and thereby intentional interference with a protocol functioning through selectively scheduling jamming transmissions to coincide with protocol message transmissions, more difficult.

7. References

7.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

- [2] Moy, J., "OSPF Database Overflow", RFC 1765, March 1995.
- [3] Marlow, D., "Host Group Extensions for CLNP Multicasting", RFC 1768, March 1995.
- [4] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [5] Clausen, T., Ed., and P. Jacquet, Ed., "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.
- [6] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", RFC 3561, July 2003.
- [7] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized MANET Packet/Message Format", Work in Progress.

Appendix A. Acknowledgements

The authors would like to acknowledge the MANET working group and the OLSRv2 Design team, in particular Joe Macker and Justin Dean (both NRL), for their contributions and discussions in developing and testing the concepts retained in this document, and Alan Cullen (BAE Systems) for his careful review of this specification. OLSRv1, as specified in [5], introduced the concept of jitter on control traffic, which was tested thoroughly by Gitte Hansen and Lars Christensen (then, both Aalborg University).

Authors' Addresses

Thomas Heide Clausen
LIX, Ecole Polytechnique, France

Phone: +33 6 6058 9349
EMail: T.Clausen@computer.org
URI: <http://www.ThomasClausen.org/>

Christopher Dearlove
BAE Systems Advanced Technology Centre

Phone: +44 1245 242194
EMail: chris.dearlove@baesystems.com
URI: <http://www.baesystems.com/>

Brian Adamson
U.S. Naval Research Laboratory

Phone: +1 202 404 1194
EMail: adamson@itd.nrl.navy.mil
URI: <http://www.nrl.navy.mil/>

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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, THE IETF TRUST 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.