

INTERNET-DRAFT
IETF MANET Working Group
Expiration: 11 January 2006

Emmanuel Baccelli
Thomas Clausen
Julien Garnier
LIX, Ecole Polytechnique, France
1 July 2005

OLSR Passive Duplicate Address Detection
draft-clausen-olsr-passive-dad-00.txt

Status of this Memo

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 12, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This draft discusses ways to perform duplicate address detection with OLSR. Methods using passive detection through OLSR messages monitoring are described and briefly evaluated.

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Applicability	4
2. Duplicate Address Detection with OLSR	4
2.1. Mismatching Neighborhood in a HELLO Message	5
2.2. MPR Selection Abnormality in a HELLO Message	5
2.3. Link-State Mismatch in a TC Message	5
2.4. TC Sequence Number Mismatch	6
2.5. Interface Mismatch in an MID Message	6
3. Scope of Passive Mechanisms	6
4. Resolving Duplicate Address Conflicts	7
5. Authors' Addresses	8
6. References	9
7. Changes	9
8. IANA Considerations	9
9. Security Considerations	9
10. Copyright	9

1. Introduction

Usually, duplicate address detection is performed when configuring network interfaces in order to ensure that unique addresses are assigned to each interface in the network. Such mechanisms commonly operate with the premises that a node "intelligently" selects an address which it supposes to be unique, followed by a duplicate address detection cycle, through which it verifies that no other active interfaces on the same network has been or is in the process of being configured with the same address. Even assuming that such a mechanism is present in a MANET, allowing MANET nodes to initially configure their interfaces with addresses unique within the network, additional complications arise: two or more MANETs may merge to form a single network, and a formerly connected MANET may partition. Thus, unless it is ensured that all MANET interfaces are assigned globally unique addresses, addressing conflicts may at any point -- not just during initial network configuration.

In this draft, we investigate the task of performing duplicate address detection when otherwise independent OLSR networks merge. We benefit from the information already exchanged by OLSR, and identify a number of mechanisms through which a node may detect a conflict between the address assigned to one of its interfaces, and an address assigned to an interface on another node. The mechanisms proposed are, thus, entirely passive, creating no additional information exchange on the network.

1.1. 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](#) [5].

Several references are made to the OLSR terminology -- that was first described and employed in [1]. Among others, this document uses the following terminology:

- Node: a device capable of participating in a MANET.
- Neighbor Node: A node X is a neighbor node of node Y if node Y can hear node X.
- Multipoint Relay (MPR): A node which is selected by its neighbor, node X, to "re-transmit" all the broadcast messages

it emits.

- HELLO messages: A node performs link sensing (the discovery of its neighborhood) via the periodic exchange of HELLO messages with its neighbors.
- TC messages: A node shares link state information with the whole network through sending and receiving TC messages.
- MID messages: In case a node features several OLSR interfaces, it announces this fact to the other nodes in the network with an MID message.

1.2. Applicability

This draft describes and discusses ways to perform passive duplicate address detection with OLSR.

2. Duplicate Address Detection with OLSR

In this section, we present different mechanisms through which an OLSR node can detect if an address currently assigned to one of its interfaces is concurrently being used by an interface on another node. We note that none of the mechanisms presented here impose any additional information exchange between nodes beyond what is already performed by OLSR.

The duplicate address detection mechanisms are based on inspecting received OLSR control messages, as well as the receiving nodes state, to determine if an address on the receiving node is duplicated elsewhere in the network. More precisely, a node can inspect a received message to detect (i) if the message appears to have been sent from an interface the receiving node or (ii) if the message contains information about interfaces of the receiving node. In either of these cases, the information contained in the received OLSR message is compared to the state recorded in the receiving node, allowing the receiving node to detect a potential duplicate of one of its addresses.

With this in mind, the following subsections will inspect the three main OLSR message types: HELLO, TC and MID-messages.

2.1. Mismatching Neighborhood in a HELLO Message

With HELLO messages, an OLSR node declares its presence to its neighborhood and its view of this neighborhood. Therefore, if a node receives a HELLO message on one of its interfaces, where the HELLO message appears to come from the node itself, a potential address duplication may incur. Since HELLO messages are never forwarded in OLSR, an OLSR node should not receive a copy of a HELLO message with any of its own interface addresses as originator (note that this ignores the situation where a node has two radio interfaces running OLSR on the same channel). Therefore, should a node receive a HELLO message with one of its own interface addresses listed as originator, there's a likely collision: two adjacent nodes may have interfaces configured with the same address. This can be confirmed by inspecting the neighborhood being advertised in the suspected HELLO message: this HELLO message will include a neighborhood that is different from the node receiving the HELLO.

2.2. MPR Selection Abnormality in a HELLO Message

With HELLO messages, a node also announces which neighbors it selects as MPR. Therefore, another intuitive diagnostic on HELLO messages is to consider MPR selection: an MPR node must be selected from among neighbors with which a symmetric link exist. Thus, if a node A has a recorded asymmetric link with node B, and receives a HELLO from node B declaring its selection as MPR, then a conflict exists as indicated: a second node A', adjacent to B, has the same address as A.

This could, however, be a false conclusion. On establishment of the link between A and B, node A receives a HELLO from B, bringing node A to see the link to B as ASYM. In the next HELLO from node A, node B will see its own address listed and conclude that the link is symmetric. Node B may, then, select A as MPR and include this selection in the next HELLO message. In this way, node A will receive an MPR selection from a node with which it has only an asymmetric link, without this being an indication of address conflicts in the network.

2.3. Link-State Mismatch in a TC Message

With TC Messages, an OLSR node announces local link state to the whole MANET. Thus, if a node A receives a TC message, declaring the address of one of node A's interfaces as MPR selector, the originator of that TC-message must be a direct neighbor of node A. If it is not

the case, it may be an indication that the address of node A's interface is duplicated somewhere in the network.

2.4. TC Sequence Number Mismatch

TC messages feature a sequence number in order indicate how recent the link state information is, and detect duplicates. Therefore if a node A receives a TC message with the address of one of its own interfaces listed as originator address and with a sequence number very different from the sequence number that node A currently is using, it can be an indication that the address of this interface of node A is concurrently being assigned to another interface in the OLSR network.

2.5. Interface Mismatch in an MID Message

With an MID message, an OLSR node with multiple interfaces declares its interface configuration to the other nodes in the network. If a node A receives a MID messages, in which the address of one of its own interfaces is listed, the remaining addresses listed in the MID must also belong to node A. Alternatively, if node A, receives an MID-message containing one or more addresses belonging to node A but also listing addresses which do not belong to node A, then at least one address is assigned to more than one node.

3. Scope of Passive Mechanisms

Passive mechanisms, such as those described in this draft, are based on the monitoring of the control messages of the routing protocol. These aim at detecting anomalies in this traffic, that can hint to possible address collisions. However, this approach has a few shortcomings, both in terms of false alarms and in terms undetected duplications.

In the rare case of a totally symmetric "mirrored" MANET (A-B-C-D-C'-B'-A'), routing message monitoring may not be sufficient to detect the duplicate addresses. In this case, the duplicate nodes cannot detect the collision with each other since the routing messages produced by the left side of the network are identical to the routing messages produced by the right side of the network (because the topology is symmetric). Sequence number mismatch monitoring may help in this case, but it may also crash the network further, as such mismatches may invalidate the link state information with each TC transmission, alternatively from the right side and the left side of the

network.

Another example is with the sequence number mechanism. This technique is not completely reliable in order to detect duplicate addresses, as delayed delivery can cause an outdated control message that is received to be possibly wrongly interpreted as a case of address duplication. This category of false alarm is more likely to be caused by TC or MID messages rather than HELLO messages, as they feature only one hop scope, suppressing delays due to forwarding.

Such cases challenge the passive approach to DAD. Therefore other techniques maybe employed in addition to passive mechanisms in order to increase the reliability of the DAD. These techniques can be called active, or semi-passive, depending on how much additional overhead is produced by the mechanism.

Semi-passive techniques involve deeper analysis of the link state information traffic, such as tracking and processing the history of such traffic, in order to prevent errors. However, these techniques come with much more processing and memory needs, a fact that must be carefully evaluated.

Active techniques involve sending specific DAD information or messages, in addition to the routing control overhead. For instance, flooding a neighbor solicitation message is part of such a technique. These can be more efficient than passive waiting, but they nevertheless come with greater overhead, a fact that must also be carefully evaluated.

4. Resolving Duplicate Address Conflicts

The purpose of the mechanisms, described in this paper, is to detect when two or more interfaces in the network have been configured with the same address -- that a duplicate address conflict exists in the network. The logical next-step to having detected this situation is to resolve it -- to reconfigure nodes such that each interface participating in the OLSR network has a network-wide unique address.

Resolving a duplicate address conflict is, functionally, orthogonal to detecting a duplicate address conflict and, depending on the specificities of the network, different mechanisms can be employed. In this section, we briefly outline a few general approaches to resolving duplicate address conflict. The objective, however, remains to remove conflicting interfaces from the OLSR network, while disrupting the network operation as little as possible.

The simplest solution, once a duplicate address conflict is detected, is for a node to simply disable the local interface(s) which are

conflicting. If these interfaces then wish to enter the network again, a new initial autoconfiguration cycle must be initiated. The advantage of this method is its simplicity and fact that no lengthy election procedure must be completed before duplicate address conflicts are resolved. The disadvantage is, that when a conflict arises, all conflicting interfaces are potentially disabled without consideration to traffic (or even necessity: when two interfaces are conflicting, it suffices to disable one of them, not both).

A more elegant class of solutions to resolving a duplicate address conflict would be for the node(s) which detect a conflict to "negotiate" which interface should yield -- possibly based on metrics such as active traffic flows for a given interface etc. This negotiation would take form of a broadcast of information (a "CONFLICT" message), containing necessary information for a recipient to decide if it should yield and disable a given interface, or not.

5. Authors' Addresses

Thomas Heide Clausen,
Project PCRI
Pole Commun de Recherche en Informatique du plateau de Saclay,
CNRS, Ecole Polytechnique, INRIA, Universite Paris Sud,
Ecole polytechnique,
Laboratoire d'informatique,
91128 Palaiseau Cedex, France
Phone: +33 1 69 33 40 73,
Email: T.Clausen@computer.org

Emmanuel Baccelli
HITACHI Labs Europe/ Project PCRI,
Pole Commun de Recherche en Informatique du plateau de Saclay,
CNRS, Ecole Polytechnique, INRIA, Universite Paris Sud,
Ecole polytechnique,
Laboratoire d'informatique,
91128 Palaiseau Cedex, France
Phone: +33 1 69 33 40 73,
Email: Emmanuel.Baccelli@inria.fr

Julien Garnier,
Project PCRI
Pole Commun de Recherche en Informatique du plateau de Saclay,
CNRS, Ecole Polytechnique, INRIA, Universite Paris Sud,
Ecole polytechnique,

Laboratoire d'informatique,
91128 Palaiseau Cedex, France
Phone: +33 1 69 33 40 73,
Email: Julien.Garnier@polytechnique.fr

6. References

- [1] T. Clausen, P. Jacquet, ``[RFC 3626](#): Optimized Link State Routing Protocol," Request for Comments (Experimental), Internet Engineering Task Force, October 2003.
- [2] E. Baccelli, T. Clausen, J. Garnier, ``Duplicate Address Detection in OLSR Networks," WPMC Proceedings, September 2005.

7. Changes

This is the initial version of this specification.

8. IANA Considerations

This document does currently not specify IANA considerations.

9. Security Considerations

This document does not specify any security considerations.

10. Copyright

Copyright (C) The Internet Society (2004). 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.