

Vulnerability Analysis of the Optimized Link State Routing Protocol version 2 (OLSRv2)

Thomas Heide Clausen, Ulrich Herberg

► **To cite this version:**

Thomas Heide Clausen, Ulrich Herberg. Vulnerability Analysis of the Optimized Link State Routing Protocol version 2 (OLSRv2). [Research Report] RR-7203, INRIA. 2010. <inria-00456376>

HAL Id: inria-00456376

<https://hal.inria.fr/inria-00456376>

Submitted on 14 Feb 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

***Vulnerability Analysis of the Optimized Link State
Routing Protocol version 2 (OLSRv2)***

Thomas Clausen, Ulrich Herberg

N° 7203

February 2010

A large, light gray stylized 'R' logo is positioned to the left of the text. A horizontal gray brushstroke is located below the text.

R *apport
de recherche*

Vulnerability Analysis of the Optimized Link State Routing Protocol version 2 (OLSRv2)

Thomas Clausen*, Ulrich Herberg†

Thème : COM – Systèmes communicants
Équipes-Projets Hipercom

Rapport de recherche n° 7203 — February 2010 — 17 pages

Abstract: Mobile Ad hoc NETWORKS (MANETs) are leaving the confines of research laboratories, to find place in real-world deployments. Outside specialized domains (military, vehicular, etc.), city-wide community-networks are emerging, connecting regular Internet users with each other, and with the Internet, via MANETs. Growing to encompass more than a handful of “trusted participants”, the question of preserving the MANET network connectivity, even when faced with careless or malicious participants, arises, and must be addressed.

A first step towards protecting a MANET is to analyze the vulnerabilities of the routing protocol, managing the connectivity. By understanding how the algorithms of the routing protocol operate, and how these can be exploited by those with ill intent, countermeasures can be developed, readying MANETs for wider deployment and use.

This paper takes an abstract look at the algorithms that constitute the Optimized Link State Routing Protocol version 2 (OLSRv2), and identifies for each protocol element the possible vulnerabilities and attacks – in a certain way, provides a “cookbook” for how to best attack an operational OLSRv2 network, or for how to proceed with developing protective countermeasures against these attacks.

Key-words: OLSRv2, MANET, Vulnerability Analysis, Security

* LIX – Ecole Polytechnique, Thomas@ThomasClausen.org

† LIX – Ecole Polytechnique, Ulrich@Herberg.name

Analyse de vulnérabilité du protocole de routage Optimized Link State Routing Protocol version 2 (OLSRv2)

Résumé : Les réseaux mobiles MANETs (Mobile Ad hoc NETworks) sortent des laboratoires de recherche pour être déployés dans le monde réel. Outre les applications spécialisées (militaires, véhiculaires etc.), des réseaux communautaires urbains émergent pour connecter des simples utilisateurs d'Internet à d'autres utilisateurs et à Internet via MANETs. Pour supporter un nombre croissant d'utilisateurs au-delà d'une poignée de participants de confiance, la question de préserver la connectivité des réseaux MANET face à des utilisateurs imprudents ou malicieux se pose.

Un premier pas vers la protection de MANET est d'analyser les vulnérabilités du protocole de routage qui gère la connectivité. En comprenant en profondeur comment les algorithmes du protocole de routage opèrent et comment ils peuvent être exploités par des utilisateurs indécidés, des contre-mesures peuvent être développées afin de rendre MANET prêt à être déployé à plus grande échelle.

Ce rapport examine de manière conceptuelle les algorithmes qui constituent le protocole OLSRV2 (Optimized Link State Routing Protocol version 2) et pour chaque élément du protocole identifie les éventuelles vulnérabilités et attaques possibles. En quelque sorte, le rapport procure un manuel sur la meilleure façon d'attaquer un réseau OLSRV2 opérationnel, mais aussi sur les méthodes pour développer les contre-mesures pour se protéger de ces attaques.

Mots-clés : OLSRV2, MANET, analyse de vulnérabilité, sécurité

1 Introduction

OLSRv2 (the Optimized Link State Routing Protocol version 2) [10], [9], [11], [12], [13] is a successor to the widely deployed OLSR [7] routing protocol for MANETs (Mobile Ad hoc NETWORKS). OLSRv2 retains the same basic algorithms as its predecessor, however offers various improvements, *e.g.* a modular and flexible architecture allowing extensions, such as for security, to be developed as add-ons to the basic protocol.

The developments reflected in OLSRv2 have been motivated by increased real-world deployment experiences, *e.g.* from networks such as FunkFeuer [14], and the requirements presented for continued successful operation of these networks. With participation in such networks increasing (the FunkFeuer community network has, *e.g.*, roughly 400 individual participants), operating with the assumption, that participants can be “trusted” to behave in a non-destructive way, is utopia. Taking the Internet as an example, as participation in the network increases and becomes more diverse, more efforts are required to preserve the integrity and operation of the network. Most SMTP-servers were, *e.g.*, initially available for use by all and sundry on the Internet – with an increased populace on the Internet, the recommended practice is to require authentication and accounting for users of such SMTP servers [8].

A first step towards hardening against attacks disrupting the connectivity of a network, is to understand the vulnerabilities of routing protocol, managing the connectivity. This paper therefore analyzes OLSRv2, to understand its inherent vulnerabilities and resiliences. The authors do not claim completeness of the analysis, but hope that the identified attacks as presented form a meaningful starting-point for OLSRv2 security.

1.1 OLSRv2 Overview

OLSRv2 contains three basic processes: Neighborhood Discovery, MPR Flooding and Link State Advertisements.

1.1.1 Neighborhood Discovery

The process, whereby each router discovers the routers which are in direct communication range of itself (1-hop neighbors), and detects with which of these it can establish bi-directional communication. Each router sends HELLOs, listing the identifiers of all the routers from which it has recently received a HELLO, as well as the “status” of the link (heard, verified bi-directional). A router *a* receiving a HELLO from a neighbor *b* in which *b* indicates to have recently received a HELLO from *a* considers the link *a-b* to be bi-directional. As *b* lists identifiers of all its neighbors in its HELLO, *a* learns the “neighbors of its neighbors” (2-hop neighbors) through this process. HELLOs are sent periodically, however certain events may trigger non-periodic HELLOs.

1.1.2 MPR Flooding

The process whereby each router is able to, efficiently, conduct network-wide broadcasts. Each router designates, from among its bi-directional neighbors, a subset (MPR set) such that a message transmitted by the router and relayed by

the MPR set is received by all its 2-hop neighbors. MPR selection is encoded in outgoing HELLOs. The set of routers having selected a given router as MPR is the MPR-selector-set of that router. A study of the MPR flooding algorithm can be found in [3].

1.1.3 Link State Advertisement

The process whereby routers are determining which link state information to advertise through the network. Each router must advertise links between itself and its MPR-selector-set, in order to allow all routers to calculate shortest paths. Such link state advertisements are carried in TC messages, are broadcast through the network using the MPR Flooding process. As a router selects MPRs only from among bi-directional neighbors, links advertised in TC are also bi-directional. TC messages are sent periodically, however certain events may trigger non-periodic TCs.

1.2 Link State Vulnerability Taxonomy

Proper functioning of OLSRv2 assumes that (i) each router can acquire and maintain a topology map, accurately reflecting the effective network topology; and (ii) that the network converges, *i.e.* that all routers in the network will have sufficiently identical topology maps. An OLSRv2 network can be disrupted by breaking either of these assumptions, specifically (a) routers may be prevented from acquiring a topology map of the network; (b) routers may acquire a topology map, which does not reflect the effective network topology; and (c) two or more routers may acquire inconsistent topology maps.

1.3 OLSRv2 Attack Vectors

Besides “radio jamming”, attacks on OLSRv2 consist of a malicious router injecting “correctly looking, but invalid, control traffic” (TCs, HELLOs) into the network. A malicious router can either (a) lie about itself (its ID, its willingness to serve as MPR), henceforth *Identity Spoofing* or (b) lie about its relationship to other routers (pretend existence of links to other routers), henceforth *Link Spoofing*. Such attacks will *in-fine* cause disruption in the Link State Advertisement process, through targeting the MPR Flooding mechanism, or by causing incorrect link state information to be included in TCs, causing routers to have incomplete, inaccurate or inconsistent topology maps. In a different class of attacks, a malicious router injects control traffic, tuned to cause an *in-router resource exhaustion*, *e.g.* by causing the algorithms calculating routing tables or MPR sets to be invoked continuously, preventing the internal state of the router from converging.

1.4 Paper Outline

The remainder of this paper is organized as follows: section 2, 3, and 4 each represents a class of disruptive attacks against OLSRv2, detailing a number of attacks in each class. Section 5 summarizes the identified vulnerabilities in an OLSRv2 network, and section 6 summarizes the ways OLSRv2 has inherent resilience to. The paper is concluded in section 7.

2 Topology Map Acquisition

Topology Map Acquisition relates to the ability for a – any – given router in the network to acquire a representation of the network connectivity. A router, unable to acquire a topology map, is incapable of calculating routing paths and participating in forwarding data. Topology map acquisition can be hindered by (a) TC messages to not being delivered to (all) routers in the network, such as what happens in case of Flooding Disruption, or (b) in case of “jamming” of the communication channel.

2.1 Flooding Disruption

MPR selection (section 1.1.2) uses information about a router’s 1-hop and 2-hop neighborhood, assuming that (i) this information is accurate, and (ii) all 1-hop neighbors are equally apt as MPR. Thus, a malicious router will seek to manipulate the 1-hop and 2-hop neighborhood information in a router such as to cause the MPR selection to fail.

2.1.1 Flooding Disruption due to Identity Spoofing

In figure 1, a malicious router X spoofs the identity of b . The link between X and c is correctly detected and listed in X ’s HELLOs. a will receive HELLOs indicating a links, respectively $b:\{ b-e \}$, $X:\{ X-c, X-e \}$, and $d:\{ d-e, d-c \}$. For a , X and d are equal candidates for MPR selection.




Figure 1: *Identity Spoofing*: The malicious router spoofs address of router b .

If b and X (i) accept MPR selection and (ii) forward flooded traffic *as-if* they were both b , identity spoofing by X is harmless. If X does not forward flooded traffic (*i.e.* does not accept MPR selection), its presence entails flooding disruption: selecting b over d renders c unreachable by flooded traffic.




Figure 2: *Identity Spoofing*: flooding attach: 2-hop address duplication.

In figure 2, X (gray) spoofs the identity of x (white), *i.e.* a and c both receive HELLOs from a router identifying as x . For b , a and c present the same neighbor sets, and are equal candidates for MPR selection. If b selects only a as MPR, c will not relay flooded traffic from or transiting via b , and the (white) router x (and routers to the “right” of it) will not receive flooded traffic.

2.1.2 Flooding Disruption due to Link Spoofing

In the network in figure 3, the malicious router X spoofs links to the existing router c , as well as to a fictitious w . a receives HELLOs from X and b , reporting $X:\{X-c, X-w\}$, $b:\{b-c\}$. All else being equal, X appears a better choice as MPR than b , as X appears to cover all neighbors of b , plus w .




Figure 3: Link Spoofing: Flooding Disruption

As a will not select b as MPR, b will not relay flooded messages received from a . The routers left of b (starting with c) will, thus, not receive any flooded messages from or transiting a (*e.g.* a message originating from s).

2.2 Radio Jamming

Radio Jamming is an attack, in which access to the communication channel between routers is hindered by, *e.g.*, a powerful transmitter is generating “white noise” on the channel. Due to the ease of access to the channel, this is particularly possible in wireless networks. Jamming affects reception, thus interfaces on a “jammed” channel are unable to receive HELLO and TCs. Depending on lower layers, this may not affect transmissions: HELLOs and TCs from a router with “jammed” interfaces may be received by other routers. As the Neighborhood Discovery process of OLSRv2 identifies and uses only bi-directional links for the Link State Advertisement process, a link from a jammed router to a non-jammed router would not be considered, and the jammed router appear simply as “disconnected” for the un-jammed part of the network – which is able to maintain accurate topology maps.

3 Effective Topology

Link-state protocols assume that each router can acquire an accurate topology map, reflecting the *effective network topology*. This implies that the routing protocol, through its message exchange, identifies a path from a source to a destination, and this path is valid for forwarding data traffic.

3.1 Incorrect Forwarding

In OLSRv2, routers send TCs and HELLOs using link-local transmissions; the routing process in each router retransmits received messages, destined for network-wide diffusion. If the router is not configured to enable forwarding, this will not affect acquisition of a topology map by the routing protocol – but will cause a discrepancy between the effective topology and the topology map.

3.2 Wormholes

A wormhole, depicted in the example in figure 4, may be established between two collaborating *devices*, connected by an *out-of-band* channel; these devices send traffic through the “tunnel” to their alter-ego, which “replays” the traffic. Thus, router *d* and router *s* appear as-if direct neighbors and reachable from each other in 1 hop through the tunnel, with the path through the MANET being 100 hops long.




Figure 4: Wormholing between two collaborating devices not participating in the routing protocol, “tunneling” traffic between *s* and *d* over an “out-of-band” channel.

The impact of a wormhole depends on its detailed behavior. If the wormhole relays control traffic, but not data traffic, the considerations in section 3.1 applies. If it relays control and data traffic alike, it is identical to a usable link: the routing protocol will generate a topology map reflecting this as the effective network topology. The efficiency of the topology so obtained depends on (i) the wormhole characteristics, (ii) how the wormhole presents itself and (iii) how paths are calculated. If the cost of the wormhole “link” represents the actual cost of transit, then the wormhole may in the worst case cause no degradation in performance, in the best case improve performance by offering a better path. If the wormhole “misrepresents” the cost of transit, then the presence of the wormhole results in a degradation in performance as compared to using the non-wormhole path. Conversely, if the “link” presented by the wormhole has better characteristics, the wormhole results in improved performance.

An additional consideration with regards to wormholes is, that it may be undesirable to have data traffic transit such a path: an attacker could, by introducing a wormhole, acquire the ability to record and inspect transiting data traffic.

3.3 Sequence Number Attacks

OLSRv2 uses two different sequence numbers in TCs, to (i) avoid processing and forwarding the same message more than once (Message Sequence Number), and (ii) to ensure that old information, arriving late due to *e.g.* long paths other delays, is not allowed to overwrite fresher information (Advertised Neighbor Sequence Number – ANSN).

For (i), an attack may consist of a malicious router spoofing the identity of another router in the network, and transmitting a large number of TCs, each with different Message Sequence Numbers. Subsequent TCs with the same sequence numbers, originating from the router whose identity was spoofed, would thence be ignored, until eventually information concerning these “spoofed” TC messages expires.

For (ii), an attack may consist of a malicious router spoofing the identity of another router in the network, and transmitting a single TC, with an ANSN significantly larger than that which was last used by the legitimate router. Routers will retain this larger ANSN as “the most fresh information” and discard subsequent TCs with lower sequence numbers as being “old”.

3.4 Message Timing Attacks

In OLSRv2, each control message may contain “validity time” and “interval time” fields, identifying the time for which information in that control message should be considered valid until discarded, and the time until the next control message of the same type should be expected [11].

3.4.1 Interval Time Attack

A use of the expected interval between two successive HELLO messages is for determining the link quality in Neighbor Discovery process, as described in [7]: if messages are not received with the expected intervals (*e.g.* a certain fraction of messages are missing), then this may be used to exclude a link from being considered as useful, even if (some) bi-directional communication has been verified. If a malicious router X spoofs the identity of an existing router a , and sends HELLOs indicating a very low interval time, a router b receiving this HELLO will expect the following HELLO to arrive within the interval time indicated – or otherwise, decrease the link quality for the link $a-b$. Thus, X may cause b 's estimate of the link quality for the link $a-b$ to fall below the limit, where it is no longer considered as useful and, thus, not used.

3.4.2 Validity Time Attack

A malicious router, X , can spoof the identity of a router a and send a HELLO using a very low validity time (*e.g.* 1 ms). A receiving router b will discard the information upon expiration of that interval, *i.e.* a link between router a and b will be “torn down” by X .

3.5 Indirect Jamming

Indirect Jamming is when a malicious router X by its actions causes legitimate routers to generate inordinate amounts of control traffic. This increases channel

occupation, and the overhead in each receiving router processing this control traffic. With this traffic originating from legitimate routers, the malicious device may remain undetected to the wider network.

3.5.1 Indirect Jamming: Neighborhood Discovery

Figure 5 illustrates indirect jamming of the Neighborhood Discovery process. A malicious router X advertises a symmetric spoofed link to the non-existing router b (at time t_0). a selects X as MPR upon reception of the HELLO, and will trigger a HELLO at t_1 . Overhearing this triggered HELLO, the attacker sends another HELLO at t_2 , advertising the link to b as lost, which leads to router a deselecting the attacker as MPR, and another triggered message at t_4 . The cycle may be repeated, alternating advertising the link X - b as LOST and SYM.




Figure 5: Indirect Jamming in Neighborhood Discovery

Indirect Jamming of the Neighborhood Discovery process will cause additional MPR set calculations; are “triggered HELLOs” enabled, an increased HELLO frequency occurs.

3.5.2 Indirect Jamming: Link State Advertisement

Similar to 3.5.1, figure 6 illustrates indirect jamming of the Link State Advertisement process. A malicious router X may “flip” between selecting a as MPR, and between advertising the link a - X as lost. This leads a to update its set of advertised neighbors (as the MPR Selector Set of a changes), increase the corresponding ANSN, and advertise this in a subsequent TC – which X uses to trigger another “status flip”.

Each such TC with an updated ANSN causes all routers in the network to recalculate their routing tables; are “triggered TCs” enabled, an increased TC frequency occurs.

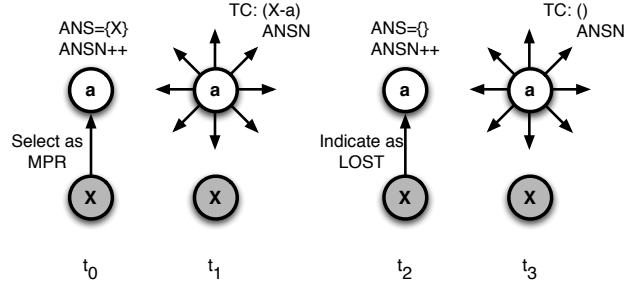


Figure 6: Indirect Jamming in Link State Advertisement

4 Inconsistent Topology

Inconsistent topology maps can occur by a malicious router employing either of identity spoofing or link spoofing for conducting an attack against an OLSRv2 network.

4.1 Identity spoofing

Identity spoofing can be employed by a malicious router via the Neighborhood Discovery process and via the Link State Advertisement process; either of which causing inconsistent topology maps in routers in the network.

4.1.1 Inconsistent Topology Maps due to Neighborhood Discovery

In order to minimize the risk of detection, the malicious router (gray circle) in figure 7 elects to not participate in the Link State Advertisement procedure, thus it does not select any MPRs and does not accept being elected as MPR (by advertising a willingness of zero). By not participating in the Link State Advertisement process, its presence is known only to c , d and e . X elects to spoof the identity of a , b , f and g , *i.e.* no routers whose identity it spoofs will not receive control messages allowing them to detect that these identities are *also* advertised elsewhere in the network. Traffic transiting d , from either side, to destination a , b , f and g will, rather than being forwarded to the intended destination, be delivered to the malicious router. Traffic transiting c with b as destination, will be delivered to the intended b . Traffic transiting c with a as destination may be delivered to the intended a via b or to the malicious router via d – as the paths are of equal length.

In figure 7, c is the only router receiving control traffic indicating two topologic locations of the identities a , b . However this is not an unusual situation: a valid link might indeed exist between routers a and d as well as between routers b and d , *e.g.* through another channel. This creates a situation wherein two or more routers have inconsistent topology maps: traffic for an identified destination is, depending on where in the network it appears, delivered to different routers.




Figure 7: *Identity Spoofing*: maximizing disruptive impact while minimizing risk of detection.

4.1.2 Inconsistent Topology Maps due to Link State Advertisements

An inconsistent topology map may also occur when the malicious router takes part in the Link State Advertisement process: spoofing an identity and selecting MPRs causes a link to the spoofed identities of the malicious router to be advertised through the network.




Figure 8: *Identity Spoofing due to Link State Advertisements*

In figure 8, the malicious router X spoofs the address of a . If X selects f as MPR, all routers in the network will be informed about the link f - a by way the TCs originating from f . Assuming that (the real) a selects b as MPR, the link b - a will also be advertised through the network.

b and c will calculate paths to a via b . e and f will calculate paths to a via f – i.e. through the malicious router X . e and f are thus disconnected from the real a . d will have the choice of selecting a path in to a in either direction.

In general, the following observations can be made: (i) the network will be split in two, with those routers closer to b than to X reaching a , and those routers closer to X than to b will be unable to reach a ; (ii) routers beyond b , i.e. routers beyond one hop away from a will be unable to detect this identity spoofing.

The impact of combining identity spoofing with Link State Advertisements is greater than the impact of section 4.1.1, as it causes alterations to the topology maps of all routers in the network. The attack is also easier to detect: with the malicious router advertised through the network, routers whose identities spoofed can detect this. When a receives a TC message from f advertising the link f - a , it can deduce that “something is wrong” as a does not have f recorded as a direct neighbor.

4.2 Link Spoofing

Link spoofing can be employed by a malicious router via the Neighborhood Discovery process and via the Link State Advertisement process; either of which causing inconsistent topology maps in routers in the network.

4.2.1 Inconsistent Topology Maps due to Neighborhood Discovery

The malicious router X in figure 3 spoofs two links to c and w . Consequently, a selects X as its sole MPR – and therefore router X is the sole router expected to advertise links to a . s selects a as MPR, thus a is expected to advertise the link $a-s$ through the network, *i.e.* using the MPR flooding process.

The topology maps acquired by the various routers in this example are:

- **a and b :** accurate topology map due to the Neighborhood Discovery process providing topological information up to 2 hops away.
- **c :** as in figure 9(a). Link state advertisements from a are not forwarded by b . Existence of s and the link $a-s$ is not known beyond b . Existence of a and the link $b-a$ and is known to b through the Neighborhood Discovery process.
- **d and beyond:** as illustrated in figure 9(b).
- **s :** accurate Topology Map corresponding to the network in figure 3. This may contain the dotted routers c and w , only if X participates in the Link State Advertisement process (section 4.2.2).




Figure 9: Perceived Topology Maps with malicious router X performing Link Spoofing in the Neighborhood Discovery Process.

4.2.2 Inconsistent Topology Maps due to Link State Advertisements

The malicious router X in figure 10 spoofs links to the existing a , by participating in the Link State Advertisement process and including the link $X-a$ in its advertisements.




Figure 10: **Link Spoofing:** The malicious router X advertises a spoofed link to router a in its TC messages, thus all routers will record the links $X-a$ and $b-a$.

As TC messages are flooded through the network, all routers will receive and record information describing the link $X-a$. If a has selected b as MPR, a will likewise flood this link state information through the network, and all routers will receive and record information describing a link $b-a$.

Routers b , c and d will calculate a shortest path via a different router than routers f and g , thus leading to network split in two. This is similar to the

impact of section 4.1.2, and when a receives a TC message from X advertising the link $X-a$, it can likewise deduce that “something is wrong” as a does not have X recorded as a direct neighbor.

5 Vulnerability Summary

Table 1 summarizes the vulnerabilities in OLSRv2, as presented in this paper. For each, a note indicates if OLSRv2 provides some inherent resilience; further discussion of the resilience of OLSRv2 is given in section 6.

| Attack | Section | OLSRv2 Resilience? |
|---------------------------------|---------|---|
| <i>Topology Map Acquisition</i> | | |
| Flooding disruption | 2.1 | None |
| Radio Jamming | 2.2 | Partly: Considers only bi-directional links |
| <i>Effective Topology</i> | | |
| Incorrect Forwarding | 3.1 | None |
| Wormholes | 3.2 | none |
| Sequence Numbers | 3.3 | Partly: rejecting “old” messages |
| Message Timing | 3.4 | None |
| Indirect Jamming | 3.5 | Yes: minimum and maximum intervals |
| <i>Inconsistent Topology</i> | | |
| Identity Spoofing | 4.1 | None |
| Link Spoofing | 4.2 | None |

Table 1: Vulnerability Summary of OLSRv2

6 Inherent OLSRv2 Resilience

While OLSRv2 does not specifically include security features (such as encryption), it has some inherent resilience against part of the attacks described in this paper. In particular, it provides the following resilience:

- *Sequence numbers*: OLSRv2 employs message sequence numbers, specific per router identity and message type. Routers keep an “information freshness” number (ANSN), incremented each time the content of a Link State Advertisement from a router changes. This allows rejecting “old” information and duplicate messages, and provides some protection against “message replay”. This also presents an attack vector (section 3.3).
- *Ignoring uni-directional links*: The Neighborhood Discovery process detects and admits only bi-directional links for use in MPR selection and Link State Advertisement. Jamming attacks (section 2.2) may affect only *reception* of control traffic, however OLSRv2 will correctly recognize, and ignore, such a link as not bi-directional.

- *Message interval bounds:* The frequency of control messages, with minimum intervals imposed for HELLO and TCs. This may limit the impact from an indirect jamming attack (section 3.5).
- *Additional reasons for rejecting control messages:* The OLSRv2 specification includes a list of reasons, for which an incoming control message should be rejected as malformed – and allows that a protocol extension may recognize additional reasons for OLSRv2 to consider a message malformed. This allows – together with the flexible message format [9] – addition of security mechanisms, such as digital signatures, while remaining compliant with the OLSRv2 standard specification.

7 Conclusion

This paper has presented a detailed analysis of security threats to the Optimized Link State Routing Protocol version 2 (OLSRv2), by taking an abstract look at the algorithms and message exchanges that constitute the protocol, and for each protocol element identifying the possible vulnerabilities and how these can be exploited. In particular, as link-state protocol, OLSRv2 assumes that (i) each router can acquire and maintain a topology map, accurately reflecting the effective network topology; and (ii) that the network converges, i.e. that all routers in the network will have sufficiently identical (consistent) topology maps. An OLSRv2 network can be effectively disrupted by breaking either of these assumptions, specifically (a) routers may be prevented from acquiring a topology map of the network; (b) routers may acquire a topology map, which does not reflect the effective network topology; and (c) two or more routers may acquire substantially inconsistent topology maps.

The disruptive attacks to OLSRv2, presented in this paper, are classified in either of these categories. For each, it is demonstrated, whether OLSRv2 has an inherent protection against the attack.

References

- [1] T. Clausen, G. Hansen, L. Christensen, G. Behrmann, "The Optimized Link State Routing Protocol, Evaluation through Experiments and Simulation", Proceedings of the IEEE conference on Wireless Personal Multimedia Communications (WPWC), October 2001, Aalborg, Denmark
- [2] U. Herberg, "Performance Evaluation of using a Dynamic Shortest Path Algorithm in OLSRv2", (To Appear) Proceedings of the 8th Annual Conference on Communication Networks and Services Research, May 2010, Montreal, Canada
- [3] A. Qayyum, L. Viennot, A. Laouiti, "Multipoint relaying: An efficient technique for flooding in mobile wireless networks", 35th Annual Hawaii International Conference on System Sciences (HICSS'2001)
- [4] L. Viennot, "Complexity results on election of multipoint relays in wireless networks", INRIA, Tech. Rep. RR-3584, 2007
- [5] Y. Li, C. Wang, W. Zhao, X. You, "The Jamming problem in IEEE 802.11-based mobile ad hoc networks with hidden terminals: Performance analysis and enhancement", International Journal of Communication Systems, Volume 22 Issue 8, Pages 937 - 958, April 2009
- [6] C. Adjih, E. Baccelli, T. Clausen, P. Jacquet, G. Rodolakis, "Fish eye OLSR scaling properties", IEEE Journal of Communication and Networks (JCN), Special Issue on Mobile Ad Hoc Wireless Networks, 2004
- [7] T. Clausen, P. Jacquet, "RFC3626: Optimized Link State Routing Protocol (OLSR)", Experimental, <http://www.ietf.org/rfc/rfc3626.txt>
- [8] C. Hutzler, D. Crocker, P. Resnick, E. Allman, T. Finch "Email Submission Operations: Access and Accountability Requirements", Best Current Practice, <http://www.ietf.org/rfc/rfc5068.txt>
- [9] T. Clausen, C. Dearlove, J. Dean, C. Adjih, "RFC5444: Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", Std. Track, <http://www.ietf.org/rfc/rfc5444.txt>
- [10] T. Clausen, C. Dearlove, B. Adamson, "RFC5148: Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", Informational, <http://www.ietf.org/rfc/rfc5148.txt>
- [11] T. Clausen, C. Dearlove, "RFC5497: Representing Multi-Value Time in Mobile Ad Hoc Networks (MANETs)", Std. Track, <http://www.ietf.org/rfc/rfc5497.txt>
- [12] T. Clausen, C. Dearlove, J. Dean, "I-D: MANET Neighborhood Discovery Protocol (NHDP)", Work In Progress, <http://tools.ietf.org/id/draft-ietf-manet-nhdp>
- [13] T. Clausen, C. Dearlove, P. Jaquet, "I-D: The Optimized Link State Routing Protocol version 2 (OLSRv2)", Work In Progress, <http://tools.ietf.org/id/draft-ietf-manet-olsrv2>

[14] <http://www.funkfeuer.at/>

[15] IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, June 2007

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 3 |
| 1.1 | OLSRv2 Overview | 3 |
| 1.1.1 | Neighborhood Discovery | 3 |
| 1.1.2 | MPR Flooding | 3 |
| 1.1.3 | Link State Advertisement | 4 |
| 1.2 | Link State Vulnerability Taxonomy | 4 |
| 1.3 | OLSRv2 Attack Vectors | 4 |
| 1.4 | Paper Outline | 4 |
| 2 | Topology Map Acquisition | 5 |
| 2.1 | Flooding Disruption | 5 |
| 2.1.1 | Flooding Disruption due to Identity Spoofing | 5 |
| 2.1.2 | Flooding Disruption due to Link Spoofing | 6 |
| 2.2 | Radio Jamming | 6 |
| 3 | Effective Topology | 6 |
| 3.1 | Incorrect Forwarding | 7 |
| 3.2 | Wormholes | 7 |
| 3.3 | Sequence Number Attacks | 8 |
| 3.4 | Message Timing Attacks | 8 |
| 3.4.1 | Interval Time Attack | 8 |
| 3.4.2 | Validity Time Attack | 8 |
| 3.5 | Indirect Jamming | 8 |
| 3.5.1 | Indirect Jamming: Neighborhood Discovery | 9 |
| 3.5.2 | Indirect Jamming: Link State Advertisement | 9 |
| 4 | Inconsistent Topology | 10 |
| 4.1 | Identity spoofing | 10 |
| 4.1.1 | Inconsistent Topology Maps due to Neighborhood Discovery | 10 |
| 4.1.2 | Inconsistent Topology Maps due to Link State Advertisements | 11 |
| 4.2 | Link Spoofing | 11 |
| 4.2.1 | Inconsistent Topology Maps due to Neighborhood Discovery | 12 |
| 4.2.2 | Inconsistent Topology Maps due to Link State Advertisements | 12 |
| 5 | Vulnerability Summary | 13 |
| 6 | Inherent OLSRv2 Resilience | 13 |
| 7 | Conclusion | 14 |



Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399