



HAL
open science

Performance Analysis of SNMP in OLSRv2-routed MANETs

Robert Cole, Ulrich Herberg

► **To cite this version:**

Robert Cole, Ulrich Herberg. Performance Analysis of SNMP in OLSRv2-routed MANETs. [Research Report] RR-7407, INRIA. 2011, pp.22. inria-00523607v2

HAL Id: inria-00523607

<https://hal.inria.fr/inria-00523607v2>

Submitted on 12 Jul 2011

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

*Performance Analysis of SNMP in
OLSRv2-routed MANETs*

Robert Cole, Ulrich Herberg

N° 7407

July 2011



*R*apport
de recherche

Performance Analysis of SNMP in OLSRv2-routed MANETs

Robert Cole*, Ulrich Herberg†

Thème : COM – Systèmes communicants
Équipe-Projet Hipercom

Rapport de recherche n° 7407 — July 2011 — 22 pages

Abstract: Mobile Ad Hoc NETWORKs (MANETs) are generally thought of as infrastructure-less and largely “un-managed”, capable of accommodating highly dynamic network topologies. Yet, while the network may be un-managed, monitoring performance and setting configuration parameters post-deployment, remains important in order to ensure proper “tuning” and maintenance of a MANET. While SNMP is sometimes considered too “heavy” for MANETs – a too chatty a protocol with too large protocol messages – it remains the predominant management and monitoring protocol in the Internet, and many implementations exist. This memorandum analyzes SNMP in an OLSRv2-routed MANET, with the purpose of investigating performance metrics, such as delivery ratio, delay, management overhead, collisions and performance monitoring accuracy. In order to address concerns both regarding SNMP being “heavy”, as well as regarding the accuracy of performance reports obtained via SNMP polling in MANETs, where path delays can be highly variable, the utility of performance reporting proxies, *i.e.* the REPORT-MIB, is studied. The obtained results show that a significant benefit can be obtained by so deploying performance reporting proxies in an SNMP managed MANET. The investigations are supported by way of network simulations (NS2).

Key-words: OLSRv2, MANET, management, performance management, control, SNMP, performance study, simulation, NS2

* Space and Terrestrial Communications - US Army CERDEC - Aberdeen Proving Ground, MD, USA 21005 - robert.g.cole@us.army.mil

† LIX - Ecole Polytechnique, Ulrich@Herberg.name

Analyse de Performance de SNMP dans des réseaux MANETs basés sur OLSRv2

Résumé : Lorsqu'on parle de réseaux mobiles ad-hoc (MANETs), on pense généralement à des réseaux sans infrastructure et à des déploiements en réseaux largement non-gérés, pouvant s'adapter à des topologies de réseau très changeantes. Néanmoins, bien que l'infrastructure du réseau est de nature non-gérée, la surveillance des performances du réseau et le choix des paramètres de configuration une fois le réseau déployé demeurent primordiaux pour la maintenance et le réglage d'un réseau MANET. Alors que SNMP est parfois considéré trop "lourd" pour des MANETs, il demeure le protocole prédominant de management et monitoring d'Internet, et beaucoup d'implémentations du protocole existent. Ce rapport analyse SNMP dans des MANETs basés sur OLSRv2, avec l'intention de déterminer des métriques de performance, comme le taux de remise, délai, overhead et collisions dans le simulateur de réseaux NS2.

Mots-clés : OLSRv2, MANET, management, performance management, control, SNMP, performance study, simulation, NS2

1 Introduction

Mobile Ad Hoc Network (MANET) routing protocols are commonly assumed to be entirely self-managing: routers perceive the topology of a MANET by way of control message exchanges, with changes to the topology being reflected in routing tables of each router after a bounded convergence time. Usually, no operator intervention is required: variable parameters for the routing protocol are either negotiated in the control traffic exchange, or are of only local importance to each router (*i.e.* do not influence interoperability). Still, external management and monitoring of a MANET routing protocol may be desired, for optimizing routing protocol operation, *e.g.* to attain a more stable perceived topology, a lower control traffic overhead, and ultimately a higher data delivery ratio, a lower end-to-end delay, and less bandwidth and energy usage.

This memorandum analyzes the performance of the Simple Network Management Protocol (SNMP), the prevailing management and monitoring protocol in the Internet, in the context of an OLSRv2 routed MANET. OLSRv2 is currently in the process of being standardized by the MANET working group of the IETF¹. Further, this memorandum analyzes the benefits of performance reporting proxies for reducing network management overhead, and for improving the accuracy of performance reports in MANETs.

Surveys of performance aspects of SNMP exist, *e.g.* [1], yet – to the best of the authors’ knowledge – none consider performance in MANETs. [2] proposes an extension to SNMP that allows aggregation, and presents a study of that extension applied in airborne tactical networks, with a static network of nodes arranged in a grid. However [2] presents no general performance analysis.

Reasons for the lack of research in this area may be twofold: (i) SNMP may be considered too “heavy” for MANETs, yet as no alternative “light-weight” management protocol has been standardized, SNMP remains *the* (Internet) management protocol². (ii) Despite the ‘S’ in SNMP meaning “simple”, SNMP is composed by a large corpus of RFCs, rendering a fully compliant implementation of SNMP for network simulators a daunting undertaking.

[4] presents a tool, AgentJ, which allows plugging unmodified Java protocols into NS2 for simulation studies. This memorandum uses AgentJ to plug “SNMP4J” [5] and JOLSRv2 [6] (Java implementations of OLSRv2 and SNMP) into NS2, as a basis for the studies undertaken.

1.1 Memorandum Outline

The remainder of this memorandum is organized as follows: Section 2 provides a brief overview of OLSRv2 and SNMP. Section 3 describes the motivation for monitoring and controlling OLSRv2 routed MANETs. Section 4 presents a management architecture for OLSRv2, including the role of performance reporting proxies, *e.g.* the REPORT-MIB [20]. Section 5 describes the simulation settings for the performance analysis of SNMP and the REPORT-MIB in OLSRv2-based MANETs, and details the results. This memorandum is concluded in section 6.

¹The Internet Engineering Task Force: <http://www.ietf.org>

²The IETF has standardized NETCONF [3] in 2006, but not with the focus on constrained devices such as MANET routers

2 Overview of OLSRv2 and SNMP

This section outlines OLSRv2 and SNMP.

2.1 OLSRv2 Overview

The Optimized Link State Routing Protocol version 2 (OLSRv2) [8, 9, 10, 11] is a successor to the widely deployed OLSR [12] routing protocol for MANETs. 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. OLSRv2 contains three basic processes: Neighborhood Discovery, MPR Flooding and Link State Advertisements. The basic operation of OLSRv2 is detailed in section 2.1.1 to 2.1.3 below, followed by the flexible message format used by OLSRv2, in section 2.1.4, and a discussion of the configuration and operation of OLSRv2 routers in section 2.1.5.

2.1.1 Neighborhood Discovery (NHDP)

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 – called SYM). 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. NHDP enables each router interface to apply a *hysteresis function* which, in addition to the message exchange, may constrain when a link is considered as “usable” or not: for example, a router may elect to not consider, and thus not advertise, a link as SYM or HEARD unless a certain ratio of HELLOs are received, unless the SNR reaches a given threshold, etc. Symmetrically, a router may decide to stop advertising a link as SYM or HEARD, subject to similar such constraints.

2.1.2 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.

2.1.3 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, 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 TCs are also bi-directional. TC messages are sent periodically, however certain events may trigger non-periodic TCs.

2.1.4 Flexible Message Format

OLSRv2 employs the format specified in [8], for all protocol messages, thereby enabling scope-limited message flooding, compact (aggregated) address representation, also of non-contiguous network addresses, and the ability to associate any number of arbitrary attributes to either of control messages or addresses, by way of inclusion of Type-Length-Value objects (TLVs). The TLV structure permits any given message to be parsed correctly by allowing an implementation to “skip over” TLVs not recognized, thus enabling extensions to be developed that embed information into existing OLSRv2 control messages.

2.1.5 OLSRv2 Router Configuration

The configuration of an OLSRv2 router consists of the set of prefixes “owned”, and thus advertised, by the router, as well as interfaces of that router, participating in the OLSRv2 routing protocol. For each such interface, a set of parameters (other than the IP address(es)) apply: control message emission intervals, hysteresis values and link quality estimations. Agreement between OLSRv2 routers on these values is *not* required for interoperability: link quality and hysteresis affect only which links a given router permits to become SYM or HEARD. Control message emission intervals and message content validity are encoded in outgoing control messages, by way of TLVs, such that a recipient router can correctly process these, regardless of its own configuration.

2.2 SNMP Overview

The Structure of Management Information (SMI) standardizes a way of exposing management data (system configuration, performance measurements, etc.) by way of defining a set of *objects* on the *managed devices*. These objects may then be read and, if appropriate, set in a standardized manner via SNMP. This, by way of a *Network Management System* communicating with an *agent* on the managed device – in this case, an OLSRv2 router. SNMP and SMI do not mandate a specific set of objects to read or set, on a device, but defines a standardized way for a device to present such objects – a Management Information Base (MIB). A SMI defines modules of related management objects within such a MIB.

Three versions of SNMP exists. SNMPv1 [13] specified a set of basic network management capabilities, including a basic security model. SNMPv2 [14] extended the functionality of SNMP, notably for retrieving more voluminous data through a single request. SNMPv3 [15] saw improvements to the security model, and otherwise left the protocol as in SNMPv2. The Structure of Management Information version 2 (SMIv2) [16] is the current version of SMI, and allows designing and describing the management model for the system, protocol or device being managed. SMIv2 allows for the definition of fairly complex management models, yet allows for simplicity of chosen implementations through the definition of *Compliance statements* within the MIB.

3 Problem Statement

As indicated in section 2.1.5, OLSRv2 imposes few constraints on valid router configuration parameters. Fundamentally, the only parameter upon which agreement is required for interoperability is C – a constant, used to fix the scale and granularity of the validity and interval time values, included in protocol control messages. [9] proposes a value for this constant. As control messages carry validity time and interval time values, a recipient OLSRv2 router can behave appropriately, even if it uses vastly different values itself, so long as the recipient and sender use the same value for C .

Link admittance, by way of the hysteresis values and link quality estimation are used for an individual router to determine a suitable threshold for “considering that a link *could* be a candidate for being advertised as usable”, and thus do not need agreement.

Still, external monitoring and management may be desirable in an OLSRv2 network. A network may benefit from having its control message emission tuned according to the network dynamics: in a mostly static network, *i.e.* a network in which the topology remains stable over long durations, the control message emission frequency could be decreased in order to consume less bandwidth or less energy. Conversely, of course, in a highly dynamic network, the emission frequency could be increased for improved responsiveness.

This example requires a more “global view” of the network, than that of a single OLSRv2 router – *i.e.* entails that a *Network Management System* is able to inquire as to various performance values of the network (*e.g.* to discover the network dynamics), and to set various router parameters (*e.g.* tune up or down emission intervals). Thus, a first-order task is to identify suitable management data for an OLSRv2 routed MANET, and to describe these by way of MIBs for use by an SNMP Network Management System. A second-order task is to develop a proxy in order to (i) provide highly accurate performance measurements in delay variable MANETs and (ii) minimize SNMP overhead in the MANET.

4 OLSRv2 Management Architecture

The OLSRv2 management system architecture consists of three MIB modules: NHDP-MIB [18], OLSRv2-MIB [19], and the REPORT-MIB [20]. Both the NHDP-MIB and the OLSRv2-MIB consist of different groups, allowing (i) changing protocol parameters such as message intervals (*e.g.* for HELLOs) and information validity times (*e.g.* hold times), and (ii) monitoring the router state (*e.g.* the neighbor set).

As is standard for SNMP management architectures, a Network Management System interacts with the various components of the device models directly over the network. However, frequent polling for object values in such a system involves a frequent and bandwidth-consuming message exchange – prohibitive in a MANET where connectivity often is wireless. Further, because of highly variable network delays in such MANETs, it is not possible for a management application to determine the time associated with object values obtained via polling. In order to specifically address these issues of performance management over low bandwidth and high latency networks, the proposed OLSRv2 management sys-

tem architecture includes a *proxy capability*, denoted REPORT-MIB [20]. This proxy is located in close proximity to the managed devices, and offers remote generation of performance reports established via the management application using Remote Monitoring (RMON) [17] style control and reporting. The proxy then polls (locally) for the current values of the relevant objects necessary for the generation of the performance reporting. Hence, the bulk of the SNMP traffic is removed from the MANET and is isolated to local interaction.

[21] provides further details regarding the MIBs modules and how they allow monitoring performance of NHDP and OLSRv2.

5 Performance Study of SNMP for OLSRv2

In order to understand the implications when running SNMP in an OLSRv2 routed MANET, this section presents a performance study of SNMP in the NS2 simulator. Typical performance metrics – such as delivery ratio, delay, overhead, collision ratio and performance measurement accuracy – are evaluated.

5.1 Simulation Settings

Simulations have been conducted with JOLSRv2 [6], a fully-compliant Java implementation of OLSRv2, as routing protocol, and SNMP4J [5], a Java implementation of SNMP, hooked into NS2 using AgentJ [4]. According to [5],

“SNMP4J is an enterprise class free open source and state-of-the-art SNMP implementation for Java 2 SE 1.4 or later. SNMP4J supports command generation (managers) as well as command responding (agents). Its clean object oriented design is inspired by SNMP++, which is a well-known SNMPv1/v2c/v3 API for C++ [...]”

Simulations have been performed using the scenario parameters in table 1. Each presented data point represents an average over 10 simulation runs of randomly generated scenarios, each corresponding to these parameters.

Table 1: NS2 parameters

Parameter	Value
NS2 version	2.34
Mobility scenario	Random walk
Grid size	1000m x 1000m
Number of routers	10 - 50
Communication range	250m
Radio propagation model	Two-ray ground
Simulation time	270 secs
Interface type	802.11b
Radio frequency	2.4 GHz
OLSRv2 parameters	Proposed default values of [11]

In all scenarios, one router (with ID of 0) is positioned at exactly the center of the simulated area, and does not move. This router runs an SNMP manager.

All other routers run an SNMP agent, providing the NHDP-MIB [18] and the REPORT-MIB [20].

For the first set of simulations, the SNMP manager continuously sends requests (“get-next-request”) for the NHDP parameter `N_HOLD_TIME` to all other routers, one by one. The manager starts sending these requests after 10s, in order to allow routing tables to converge. UDP is used as transport protocol. Each request has a 500ms timeout, *i.e.* the manager aborts the request if no response has been received after 500ms, and proceeds to send a request to the next router. 25 seconds after the first request is sent, all routers have been interrogated and either responded or timed out (50 routers · 500ms timeout). The manager, then, restarts interrogating the first router again – resulting in each router being interrogated 25 times during the simulation.

Simulations are run using SNMPv2c, SNMPv3 without authentication or privacy (“SNMPv3”), SNMPv3 with SHA authentication only (“SNMPv3 (SHA)”), SNMPv3 with authentication and privacy (denoted “SNMPv3 (SHADES)” [22] and “SNMPv3 (SHAAES128)” [23]³.

For the second set of simulations, the impact of performance reporting proxies, *i.e.* the REPORT-MIB, is investigated. For these simulations, the manager polls each router 20 times over a 10 second window to collect counter values for performance reports, corresponds to standard SNMP operation for data collection for performance monitoring. With the REPORT-MIB implemented locally on each router, the SNMP manager needs only to interact with the routers twice during this period: first, to set up the report control, and, second to collect the performance report from the local REPORT-MIB instance.

The goal of these simulation studies is to measure the reduction of SNMP overhead when using the REPORT-MIB – as well as to estimate the accuracy of the performance reports generated in MANETs, where path delays may be highly variable. Specifically, SNMP management applications typically poll periodically for a common set of objects for the purpose of computing performance statistics related to state of performance objects on managed devices.

Let V be the value associated with an SNMP object of interest. Typically, management applications are interested in the rate of change of the value of this object, *i.e.*:

$$a(t) = \frac{\partial V(t)}{\partial t} \quad (1)$$

where $a(t)$ represents the actual value of the quantity of interest. SNMP management applications estimate this derivative as:

$$e(t) = \frac{V(t_2) - V(t_1)}{t_2 - t_1} \quad (2)$$

where $V(t_i)$ is the value of the SNMP object at time t_i and $e(t)$ represents the estimate of $a(t)$ at time t between t_1 and t_2 .

Because of the distance between the management application and the managed device, there typically is some variation in the round trip delay between these devices. This causes an error in the derived estimate, termed $m(t)$ for measured.

³Some implementations, *e.g.* SNMP4J and Cisco SNMP, provide other ciphers such as SHAAES192, SHAAES256 and SHA3DES, however these have not been standardized, and have therefore not been considered.

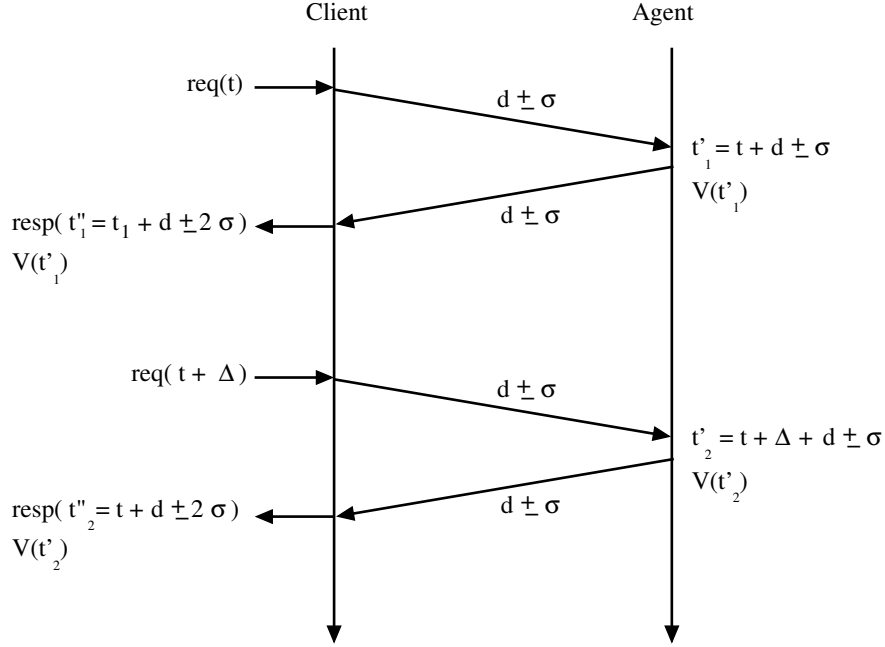


Figure 1: Impact of delay variation on SNMP performance data collection.

Referring to Figure 1 the measured result can be written, using $\Delta = t_2 - t_1$, as:

$$\begin{aligned}
 m(t) &= \frac{V(t'_2) - V(t'_1)}{t''_2 - t''_1} \\
 &\sim \frac{V(t_2) - V(t_1) \pm 2\sigma(\frac{\partial}{\partial t} V)}{\Delta[1 \pm 4(\frac{\sigma}{\Delta})]} \quad (3)
 \end{aligned}$$

which reduces to (given the definitions of the actual, measured and estimate):

$$m(t) \sim e(t) \left[1 \pm 4\left(\frac{\sigma}{\Delta}\right) \pm 2\left(\frac{\sigma}{\Delta}\right)\frac{a(t)}{e(t)} \right] \quad (4)$$

Assuming that the actual and the estimate are identical, then:

$$m(t) \sim e(t) \left[1 \pm 6\left(\frac{\sigma}{\Delta}\right) \right] \quad (5)$$

where $m(t)$ is the result as measured by the SNMP network management application, $e(t)$ is the estimate available if computed locally on the managed device, Δ is the time difference between the SNMP requests (*e.g.* polls) sent from the SNMP manager (the inverse of the polling frequency) and δ is the standard deviation of the one-way delay between the SNMP manager and the managed device, assumed to be identical for both directions across the network.

In the following, the results of the NS2-simulations will be used to develop estimates of the error between $m(t)$ and $e(t)$, illustrating the value of the REPORT-MIB for MANETs.

5.2 Simulation Results

This section presents the obtained simulation results. Figure 2 depicts the accumulated transmitted control traffic of OLSRv2 during the simulation, counting each retransmission of forwarded messages. The control traffic overhead, unsurprisingly, increases with the number of routers in the network.




Figure 2: OLSRv2 accumulated control traffic throughout the simulation

Figure 3 depicts the control traffic in a network with 50 routers, for different router velocities (except for router 0, which is static). The control traffic overhead is higher than in a static scenario, but otherwise mostly constant from 5 m/s. The reason for the difference between a static and a mobile scenario is that JOLSRv2 supports “triggered” HELLO and TC messages, generated when links break or appear between routers.

Figure 4 depicts the accumulated SNMP traffic for the different SNMP versions and security mechanisms. Again unsurprisingly, traffic grows linearly with the number of routers (and, thus, SNMP agents) in the network. SNMPv2 exhibits a far lower overhead than SNMPv3. SNMPv3 with authentication only (SHA) exhibits a higher overhead than SNMPv3 without authentication, but less than both tested encrypted SNMPv3 variants (which have an almost equal overhead).

SNMP messages for the different versions tested contain different amount of security related parameters, accounting for the differences in overhead incurred. Table 2 depicts the message sizes of the get-next-request message that is sent from the manager to the agents, measured with Wireshark between two physical routers using the same SNMP implementation, SNMP4J, as in the simulations.

SHADES and SHAAES128 have similar SNMP message sizes, confirming the almost equal plots in figure 4. The payload (the PDU) is of equal size in




Figure 3: OLSRv2 accumulated control traffic throughout the simulation in a network with 50 routers and variable router velocity.




Figure 4: Accumulated SNMP traffic overhead

Table 2: SNMP message sizes

Variant	Frame size	SNMP message size	PDU size
SNMPv3	140	103	48
SNMPv3-SHA	146	109	48
SNMPv3-SHADES	159	122	48
SNMPv3-SHAAES128	163	125	48

both, due to the mode of operation of the block cipher (see [24]). SHADES applies CBC (Cipher-block Chaining), which splits the plaintext in multiples of 8 bytes with possible padding. As the payload happens to be a multiple of 8 bytes, the cypher text has the same length as the plaintext. SHAAES128 uses a CFB (Cipher Feedback) operation mode, which always outputs the same length as the input plaintext.

Another reason for the different total SNMP traffic is the number of transmitted messages. Figure 5 compares SNMPv2c with the SHAAES128 variant of SNMPv3⁴. With SHAAES128, for each pair of routers exchanging SNMP messages, an additional initial message exchange has to be performed in order to provide replay protection, illustrated in Figure 6.




Figure 5: Number of transmitted SNMP messages

For SNMPv2, the get-next-request is directly sent and answered, whereas SNMPv3 exchanges the Authoritative EngineID and a counter of how often the agent has been rebooted, in order to provide replay protection. For the simulations presented in this memorandum, this initial exchange of parameters is only performed for the first request from the manager to an agent, not in any subsequent one – which explains why the plot in figure 5 for SNMPv3 show only slightly more frames set than SNMPv2.

Figure 7 depicts the MAC frame collision ratio. As the amount of OLSRv2 control traffic and SNMP unicast traffic increases with the number of routers in the network, so does the collision ratio. There is no significant difference between the different SNMP variants as the SNMP traffic makes up only a small fraction of the total traffic in the network. Note that this is no general observation: in the simulated scenarios, no concurrent SNMP message exchanges take place, and no other unicast data traffic is present in the network.

Figure 8 depicts the collisions for different router velocities. As OLSRv2 control traffic accounts for the majority of the traffic in the simulation, and as that control traffic does not considerably increase with velocity, as depicted in figure 3, the collision ratio remains stable, at about 12%.

⁴For the other encrypted variants the results are similar

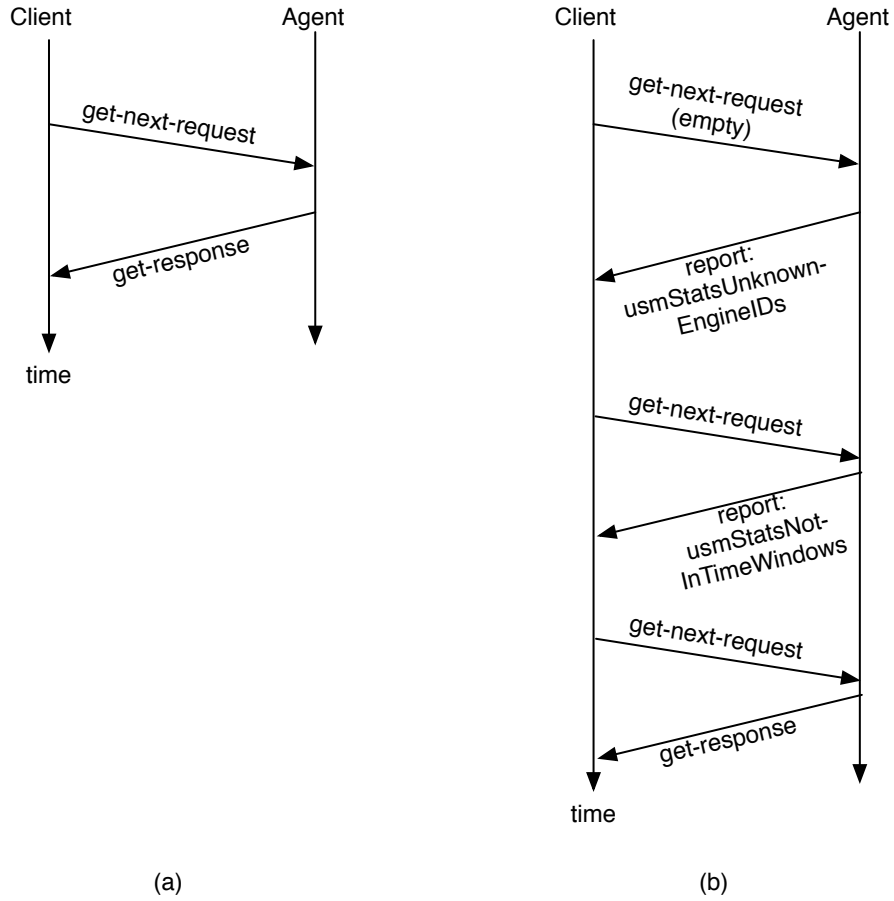


Figure 6: SNMP message exchange: (a) in SNMPv2 (b) in SNMPv3 with privacy




Figure 7: MAC collision ratio




Figure 8: MAC collision ratio in a network with 50 routers and variable router velocity.

Figure 9 depicts the message exchange delay between transmission of the get-next-request and the reception of a response by the manager. As the number of routers in the network increases, so does the message exchange delay across all SNMP variants. SNMPv3 and SNMPv3 with privacy incurs higher delays because of the initial message exchange for replay protection, depicted in figure 6.




Figure 9: Message exchange delay

Figure 10 depicts the delivery ratio for SNMP messages. With a low collision ratio (figure 7), the delivery ratio is relatively high, increasing network density. There is no significant difference observed between the different SNMP variants.




Figure 10: Delivery ratio of SNMP messages

Figure 11 depicts the delivery ratio for SNMP messages when routers are mobile. The delivery ratio decreases as the velocity increases, but remains at a relatively high level – because of the low collision ration, combined with the relatively low number of hops from the manager to all agents.




Figure 11: Delivery ratio in network a with 50 routers and variable router velocity.

Figure 12 depicts the average path length, measured in number of hops, between the SNMP manager and the agents. There is no significant difference between the different SNMP variants.

Beyond this basic understanding of the behavior of SNMP in an OLSRv2-network, the impact of the REPORT-MIB as an SNMP performance management proxy in MANET environments is of interest. Specifically, the subsequent




Figure 12: Average path length

simulations seek to quantify the possible reduction of the SNMP polling overhead in the MANET because of the placement of the REPORT-MIB proxy on the managed device, as well as to quantify the reduction in the potential error in the performance reports because of larger and varying path delays occurring in MANETs.

Figure 13 depicts number of frames sent when polling (as in standard SNMP) is used, as well as when the REPORT-MIB proxy is used. While the reduction in overhead when using the REPORT-MIB is substantial, note that this results in reports being generated only after the equivalent of 20 polling intervals. The SNMP manager interacts with each router only twice per report (configure report collection, collect performance report). Of course, the results will depend in general on the relative relationship of the report duration to the polling intervals, as well as other aspects of the network.

Equation 5 approximates the error in performance reports generated through the standard SNMP polling method, where the network manager is responsible for generating performance reports based upon collected SNMP counter. The error in these reports is related to the uncertainty on the measured times on the managed devices due to the possibility for highly variable path delays in MANETs.

Figure 14 depicts the average, standard deviation and maximum delays, experienced by these SNMP polls. These results show that the standard deviation of the round trip delays of the polls can be significant. These results can be considered as extremely conservative, as there is no other data traffic in the MANET.

Figure 15 depicts estimates of the accuracy of the performance reports, as generated through standard SNMP polling methods and based upon equation 5 and on the simulation results. The error in the performance reports can become quite significant. The middle curves in the figure are the estimates based on the simulation results, where the maximum reporting error is roughly 6% for the 50 node results. Polling at a lower frequency can improve the error estimate




Figure 13: The reduction in polling overhead in terms of frames due to the deployment of the Report-MIB.

– illustrated in the two lower curves, which assume a five times longer polling interval.

Care must be taken when increasing the polling interval, as the assumption that the difference between the actual performance statistic $a(t)$ and the sampled estimate $e(t)$ was small can break, if the polling interval becomes too long.

Finally, the upper two curves present the expected error if the standard deviation of the delay increased 5 times, such as would be the case if there was other data traffic (“background traffic”) in the network. This illustrates the potential for very large errors in the performance reports that would have been generated through standard SNMP polling over a MANET. Deployment of the REPORT-MIB proxy on the MANET routers would eliminate the existence of these performance reporting errors.

6 Conclusion

The MANET routing protocol OLSRv2 does not require any operator intervention once deployed: routers are able to accommodate frequently changing network topologies in a self-organizing manner, and the protocol is designed so as to enable a network to accommodate OLSRv2 routers with heterogeneous configurations. However, it may still be desirable to monitor the performance of a deployed network, and to tweak parameters for improving the performance of the routing protocol, *e.g.* if the conditions of the deployment change over time.




Figure 14: The average, standard deviation and maximum round trip delay for the SNMP polling over the MANET.

This memorandum analyzes the performance of SNMP, the predominant management and monitoring protocol for routers in the Internet, in OLSRv2-routed MANETs. Different evaluation metrics are considered, such as delivery ratio, delay, overhead, collisions, both in static and in mobile networks. Different variants of SNMP, notably SNMPv2c, SNMPv3 without authentication or privacy, SNMPv3 with SHA authentication only, and SNMPv3 with authentication and privacy (AES128 and DES) are studied.

This memorandum also analyzes the impact of performance reports collected through standard SNMP polling methods in the presence of large path delays and path delay variations, such as may occur in MANETs. These standard SNMP performance reporting methods generate a relatively large overhead in the network. Further, the accuracy of the resulting reports can be greatly diminished because of these delay variations. Colloquially speaking, SNMP polling across a MANET gives the worst possible result: a large traffic load on the network, resulting in erroneous, or at least inaccurate, results. This memorandum demonstrates both of these effects, as well as presents the benefits of a reporting proxy, *i.e.* the REPORT-MIB, for reducing the management overhead and in improving the accuracy of performance reports in these challenging MANET environments.




Figure 15: The typical and maximum error in standard SNMP performance reporting.

References

- [1] L. Andrey, O. Festor, A. Lahmadi, A. Pras, J. Schönwälder, “Survey of SNMP performance analysis studies”, *International Journal of Network Management*, 19: 527-548, 2009
- [2] G. Kuthethoor *et. al.*, “Performance analysis of SNMP in airborne tactical networks”, *Proceedings of the IEEE Military Communications Conference (MILCOM)*, 2008
- [3] R. Enns, “RFC4741: NETCONF Configuration Protocol”, *Std. Track*, <http://www.ietf.org/rfc/rfc4741.txt>
- [4] U. Herberg, I. Taylor, “Development Framework for Supporting Java NS2 Routing Protocols”, *Proceedings of the 2010 International Workshop on Future Engineering, Applications and Services (FEAS)*, May 2010
- [5] SNMP4J Website, <http://www.snmp4j.org>
- [6] U. Herberg, “JOLSRv2 – An OLSRv2 implementation in Java”, *Proceedings of the 4th OLSR Interop workshop*, October 2008
- [7] T. Clausen, C. Dearlove, B. Adamson, “RFC5148: Jitter Considerations in Mobile Ad Hoc Networks (MANETs)”, *Informational*, <http://www.ietf.org/rfc/rfc5148.txt>
- [8] 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>
- [9] T. Clausen, C. Dearlove, “RFC5497: Representing Multi-Value Time in Mobile Ad Hoc Networks (MANETs)”, *Std. Track*, <http://www.ietf.org/rfc/rfc5497.txt>
- [10] T. Clausen, C. Dearlove, J. Dean, “RFC6130: MANET Neighborhood Discovery Protocol (NHDP)”, *Std. Track*, <http://www.ietf.org/rfc/rfc6130.txt>
- [11] 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-olsrv2-11.txt>
- [12] T. Clausen, P. Jaquet, “RFC3626: Optimized Link State Routing Protocol (OLSR)”, *Experimental*, <http://www.ietf.org/rfc/rfc3626.txt>
- [13] J. Case, K. McCloghrie, M. Rose, S. Waldbusser, “RFC1157: A Simple Network Management Protocol (SNMP)”, <http://www.ietf.org/rfc/rfc1157.txt>
- [14] J. Case, K. McCloghrie, M. Rose, S. Waldbusser, “RFC1441: Introduction to version 2 of the Internet-standard Network Management Framework”, <http://www.ietf.org/rfc/rfc1441.txt>
- [15] R. Presuhn *et. al.*, “RFC3416: Version 3 of the Protocol Operations for the Simple Network Management Protocol (SNMP)”, *Std. Track*, <http://www.ietf.org/rfc/rfc3416.txt>

-
- [16] K. McCloghrie *et. al.*, “RFC2578: Structure of Management Information version 2 (SMIv2)”, Std. Track, <http://www.ietf.org/rfc/rfc2578.txt>
 - [17] Waldbusser, S., Cole, R.G., Kalbfleisch, C. and D. Romascanu, “RFC3577: Introduction to the RMON Family of MIB Modules”, Informational, <http://www.ietf.org/rfc/rfc3577.txt>
 - [18] U. Herberg, R. Cole, I. Chakeres, “I-D: Definition of Managed Objects for the Neighborhood Discovery Protocol”, Work In Progress, <http://tools.ietf.org/id/draft-ietf-manet-nhdp-mib-07.txt>
 - [19] U. Herberg, R. Cole, T. Clausen, “I-D: Definition of Managed Objects for the Optimized Link State Routing Protocol version 2”, Work In Progress, <http://tools.ietf.org/id/draft-ietf-manet-olsrv2-mib-03.txt>
 - [20] R. G. Cole, J. Macker, A. Morton, “I-D: Definition of Managed Objects for Performance Reporting”, Work in Progress, <http://tools.ietf.org/id/draft-ietf-manet-report-mib-01.txt>
 - [21] U. Herberg, T. Clausen, R. Cole, “MANET Network Management and Performance Monitoring for NHDP and OLSRv2”, Proceedings of the 6th International Conference on Network and Services Management (CNSM), October 2010
 - [22] U. Blumenthal, B. Wijnen, “RFC3414: User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)”, Std. Track, <http://www.ietf.org/rfc/rfc3414.txt>
 - [23] U. Blumenthal, F. Maino, K. McCloghrie, “RFC3826: The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model”, Std. Track, <http://www.ietf.org/rfc/rfc3826.txt>
 - [24] A. Menezes, P. van Oorschot, S. Vanstone, “Handbook of Applied Cryptography”, CRC Press, ISBN 0-8493-8523-7, 1996

Contents

1	Introduction	3
1.1	Memorandum Outline	3
2	Overview of OLSRv2 and SNMP	4
2.1	OLSRv2 Overview	4
2.1.1	Neighborhood Discovery (NHDP)	4
2.1.2	MPR Flooding	4
2.1.3	Link State Advertisement	4
2.1.4	Flexible Message Format	5
2.1.5	OLSRv2 Router Configuration	5
2.2	SNMP Overview	5
3	Problem Statement	6
4	OLSRv2 Management Architecture	6
5	Performance Study of SNMP for OLSRv2	7
5.1	Simulation Settings	7
5.2	Simulation Results	10
6	Conclusion	17



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