

Delay Tolerant Networking with OLSRv2

Ulrich Herberg
Trusted Systems Innovation Group
Fujitsu Laboratories of America, USA
ulrich@herberg.name

Thomas Clausen
Hipercom@LIX
Ecole Polytechnique, France
Thomas@ThomasClausen.org

Abstract—This paper proposes a simple mechanism for enabling basic delay tolerant networking with off-the-shelf MANET routing protocols – with the objective being to enable trading off slightly longer data delivery delays against resilience to a temporary lack of connectivity between a router and the ultimate destination of an IP datagram. As part of testing the benefit of said mechanism, an extreme network mobility model is proposed, entitled the “PopUp model”: a router appears in the network, and operates normally – then may disable and disappear from the network to appear later elsewhere. Observed to cause severely degraded performance for MANET routing protocols, this model is used for testing the proposed mechanism in OLSRv2-routed MANETs. The proposed mechanism shows to vastly increase the data delivery ration, with reasonably low increases in delays and control traffic overhead incurred.

I. INTRODUCTION

IP datagram transmission is based on a delivery principle, commonly known as “best effort”: if the routing table contains a “valid route” towards the destination of the datagram, then the datagram is transmitted – otherwise, it is silently dropped. In ad hoc networks, such a “best effort” behavior may be undesired: a route may not be available, but will be shortly, *e.g.*, as the routing protocol converges so as to reflect a changed topology in routing tables. A wireless link may at one point have been considered sufficiently reliable, and thus advertised and used as such, yet over time degrade until being declared as lost and possibly advertised as lost also. Destinations reachable via paths containing this now lost link are unreachable, at least until the routing protocol message exchange and table calculations complete, providing different paths through the network.

In networks in which the topology is dynamic and such transient situations are frequent, it may be preferable to slightly alter the “best effort delivery principle” so as to, colloquially speaking, *hold on to an un-transmittable IP datagram for a while, hoping that a route will become available shortly* – using graph terminology, may choose to buffer an un-transmittable IP datagram until its routing table indicates that it is in the same connected component as the intended destination. Doing so may increase the data delivery ratio, at the expense of longer delivery delays (“delay tolerant networking”) as well as additional state requirements in routers for the buffering process.

This paper proposes:

- 1) A modification to the IP datagrams transmission process, allowing for buffering un-transmittable IP datagrams;

- 2) A network mobility model, entitled “PopUp Networks”, wherein devices appear in the network, disappear, move, then re-appear;
- 3) A comprehensive set of simulations, studying the behavior of an OLSRv2-routed ad hoc network, when exposed to the previous two items.

The proposed PopUp network mobility model is presenting an “extreme” type of network behavior, and is so chosen as it generally results in vastly degraded routing protocol performance, unless special measures are taken. The purpose of this paper is to present how, even when faced with such a harsh network mobility model, a relatively simple mechanism can vastly improve the performance an off-the-shelf ad-hoc routing protocol.

A. Paper Outline

The remainder of this paper is organized as follows: section II provides a basic overview of OLSRv2, and section III specifies the proposed modification to the IP datagram transmission mechanism, denoted “Link Buffering”, for allowing delay tolerant message delivery. The “PopUp” network mobility model is proposed in section IV. Section V presents a performance evaluation of the delay-tolerance mechanism in OLSRv2-routed networks simulated when exposed to the PopUp network mobility model. Section VI concludes the paper.

II. OLSRV2 OVERVIEW

The Optimized Link State Routing Protocol version 2 (OLSRv2) [?], [?], [?], [?] is a successor to the widely deployed OLSR [?] routing protocol for MANETs. OLSRv2 retains the same basic algorithms as its predecessor, however offers various improvements, *e.g.*, a modular and flexible architecture, and in particular a flexible message format [?] by way of TLVs, allowing extensions, such as for security, to be developed as add-ons to the basic protocol whilst retaining backwards and forwards compatibility. OLSRv2 contains three basic processes: Neighborhood Discovery, MPR Flooding and Link State Advertisements. The basic operation of OLSRv2 is detailed in section II-1 to II-3 below.

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, as well as detects its 2-hop neighbors. This,

by way of a periodic HELLO message exchange, as specified in [?]. NHDP enables routers to apply a hysteresis mechanism for determining when to admit the link to a given neighbor, as well as the ability to signal a link no longer satisfactory as “LOST”, triggering recalculation of MPRs and possibly renewed Link State Advertisement.

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.

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.

III. LINK BUFFERING

This section proposes a link buffer mechanism, based on [?], which allows delaying datagram transmission when the ultimate destination is (temporarily) unavailable in a router’s routing table – and then transmit that datagram later, once the destination re-appears.

A. IP Buffering

Traditional IP datagram transmission considers, essentially, two “states” for any destination: either a route exists (transmit) or no route exists (drop). The proposed mechanism changes the default behavior of “drop” to “buffer datagrams for expected later delivery”. Table I lists the different states, combined with the appropriate action to be taken.

Table I
LINK BUFFERING STATES AND ACTIONS

State	Description	Action
No Route	No routing information is available for this destination	Buffer IP Datagrams
Route Valid	A route entry exists for this destination, and the link to the next hop towards this destination is not disconnected	Transmit IP Datagrams

Whenever a routing protocol adds or modifies a route to a destination, all buffered IP datagrams to that destination are retransmitted.

B. L2 Buffering

A variant of the mechanism described in section III-A is to consider a third state in addition to “no route” and “route valid”, called “route invalid”. This state corresponds to the situation where a destination is believed to be present in the network, although the previously selected “next hop” is

currently unavailable. This variant requires the link layer (L2) to be aware of, and able to signal to the network layer and the routing table if transmissions to the next hop fail, *i.e.* the link to the next hop is disconnected. For example, 802.11 [?] applies such a mechanism using acknowledgments: if a frame is not acknowledged by the destination (of the frame, *i.e.*, the ultimate destination or the intended next-hop), the frame is dropped (after a bounded number of failed attempts with an exponential back-off mechanism). If such a L2 mechanism is in place, IP datagrams can also be buffered in the “route invalid” state.

C. Flow Chart

Figure 1 depicts the proposed mechanism, for both IP and L2 link buffering. When a datagram is to be forwarded by a router (state “Datagram to be forwarded”), the router verifies whether it has a route towards the destination. If yes, it hands the datagram off for transmission (states “Send datagram” and “Datagram arrives at L2”). If the datagram cannot be sent successfully (*e.g.* no ACK has been received), the datagram is buffered.

If no valid route exists (state “valid route available?”), the behavior depends on whether the router is the source of the datagram or not. If yes, the datagram is buffered, otherwise it is dropped. The rationale for this differentiation is provided in section III-D.




Figure 1. Link buffering process for L3 and L2

D. Buffering at Intermediate Hops

IP datagrams may traverse several hops (*i.e.* several routers) from the source to the destination. Invalid routes can appear at any of these routers – in part, as routing protocols do not converge instantaneously on all routers. Typically, the further a router is away from a destination (in terms of hops), the longer is the convergence time for that destination¹. Thus, it is possible that a router forwarding an IP datagram still has

¹Unless there are unstable links to a “close” destination in terms of hop count, and stable links to a destination far away

a valid route to the destination and forwards it, while another router further along the path towards the destination has a fresher information about that destination and is aware that the destination is no longer available.

Consequently, buffering could be performed at any of these “intermediate” routers, however such may lead to a large demand of memory on some routers that are bottlenecks in the network. Consider the example in figure 2, where the gray router (called ‘I’) buffers datagrams if it does not have a valid routing table entry for their ultimate destinations – requiring vast amount of memory in ‘I’, or provides a target for denial-of-service attacks.




Figure 2. The gray router buffers datagrams that it forwards from the left side to the right side of the network (and vice versa)

While these reasons may make it undesirable to allow buffering in intermediate routers, the performance of this mechanism is also evaluated in section V.

IV. MODEL DESCRIPTION

In order to test the link buffering mechanism presented in section III in extreme conditions (*i.e.* where the end-to-end delivery ratio of IP datagrams is otherwise low without such a mechanism), the “PopUp” network mobility model is proposed.

The PopUp mobility model is defined as follows: every router in a MANET has two conjunct states (as depicted in figure 3): *enabled* and *disabled*.

Whenever a router is “enabled”, it operates as usual, *i.e.*, it may exchange control messages of a routing protocol, send and receive data traffic etc. In this state, the router does not move. The router remains for a certain time (denoted *enabled_time*) in this state. It then switches to the “disabled” state which may involve a change in position: the router can move up to a certain upper bounded distance (denoted *max_distance*) with infinite speed.

If a router is “disabled”, it cannot send or receive any messages from the network (neither data nor control traffic). The router can be considered as being “switched off” and is thus completely unresponsive to any network events. The router stays a maximum time of *disabled_time* in that state until it returns to the “enabled” state. Note that both *enabled_time* and *disabled_time* can be infinite, *i.e.*, that a router can remain in either of these states forever.

Figure 4 depicts an exemplary MANET that behaves according to the above-mentioned PopUp mobility model. At time t_0 , routers 1 to 4 do not move and are all enabled which means they operate as expected, possibly exchanging control and data traffic. In the figure in the middle at time t_1 , router 4 becomes




Figure 3. State diagram of the PopUp model

disabled and moves with infinite speed up to a certain distance from its original position. At time t_2 , depicted in the right figure, the router is enabled again and participates normally in the network.




Figure 4. Example of the PopUp mobility model

V. PERFORMANCE EVALUATION

This section presents a performance evaluation of unicast data traffic in OLSRV2-routed MANETs, both *with* and *without* the link buffering mechanism described in section III. The performance evaluation is undertaken by way of NS2 simulations. Typical metrics such as delivery ratio of data traffic, control traffic overhead, average path length and delay of data traffic are considered.

A. God Routing Protocol

In order to provide an upper bound in terms of delivery ratio for the simulated scenarios, a “God routing protocol” (denoted “GodRP”) has been implemented. The routing protocol uses the “God” object of NS2, and calculates routes to all destination based on their position and radio range, instantaneously and without any signaling. The expected performance of the GodRP is close to the best possible routing protocol, which helps to understand how well a routing protocol could theoretically perform².

A variation of the GodRP is included in the simulation as well, which performs link buffering as described in section III, denoted “GodRP-LB”. Note that the variant with link buffering (LB), does not represent an upper bound in terms of data delivery ratio, illustrated in figure 5. Four routers run a link




Figure 5. GodRP with link buffering (“GodRP-LB”) does not represent an upper bound in terms of data delivery ratio for routing protocols with link buffering

state routing protocol and are arranged in a “strip” topology at time t_0 . Router 1 sends unicast data traffic to router 4. At time t_1 , the link between router 3 and 4 breaks. Due to the non-zero convergence time of link state routing protocols, router 1 may still have a valid route to 4, so it transmits the datagrams. When router 3 receives the datagrams (after being forwarded by router 2), it may have a fresher information about the disconnected link to 4. If no intermediate buffering is used, it will drop the datagram. However, when using intermediate buffering, the datagram is buffered, and delivered later (at time t_2). Using GodRP-LB, the datagram will never be delivered in this example from t_1 : at t_1 , router 1 will be immediately informed about the link disconnection between 3 and 4, so router 1 will buffer all datagrams. Since at t_2 , router 1 is not longer connected to 2, datagrams will not reach their final destination, router 4.

B. Graceful Shutdown

In the performance evaluation, a variation of OLSRv2 with link buffering as considered: if a router is gracefully shutdown when transiting from “enabled” to the “disabled” state, it sends

²In some cases, a perfect routing protocol would perform better by dropping some packets even if actually routable, e.g. to destinations further away in cases of massive load on the channel. This may lead to a higher delivery ratio for destinations close to the router (i.e. neighbors) if the transmissions contend for the channel.

a HELLO to inform neighbors that links to the router shutting down should be considered as unavailable. Figure 6 depicts the graceful shutdown process.

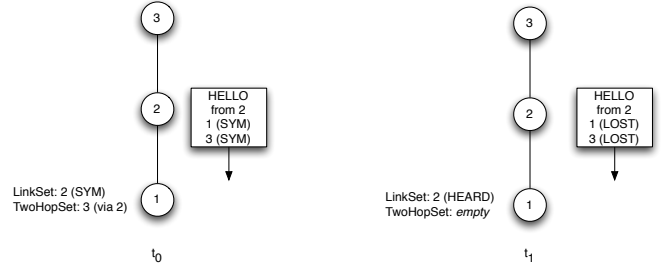


Figure 6. Graceful shutdown

At time t_0 , all three routers are in the “enabled” state. In the periodic HELLO messages from router 2, it will list both router 1 and 3 as symmetric (SYM) neighbors. Thus, router 1 will list 2 as symmetric neighbor, and 3 as two-hop neighbor reachable via 2. When router 2 is about to be shut down at time t_1 , it sends a HELLO message with all neighbors listed as LOST. Thus, router 1 will remove the 2-hop neighbor entry, and set the link status towards router 2 to HEARD. In addition, router 1 may trigger a HELLO message (and possibly a TC message, if it was selected as MPR by router 2), advertising the lost link.

Without the graceful shutdown mechanism, router 1 would have listed 2 as symmetric neighbor until the link expired. Thus, the graceful shutdown mechanism effectively reduces the convergence time of the routing protocol.

C. Simulation Settings

NS2 simulations settings are summarized in table II. Routers behave according to the PopUp mobility model as defined in section IV, in a square area of 1000x1000 meters. On average, the data traffic load is five concurrent CBR streams of 3.2 kB/s between random pairs of routers. The simulation time is 300 seconds, and each data point presented is an average over 30 runs of different scenarios, corresponding to the same abstract scenario description. A random seed was used for every simulation.

The following routing protocol variants are compared: OLSRv2, OLSRv2 with link buffering on L3 only (“OLSRv2-L3-LB”), OLSRv2 with link buffering (“OLSRv2-LB”), OLSRv2 with link buffering and graceful shutdown (“OLSRv2-LB-gf”), OLSRv2 with link buffering also at intermediate routers (“OLSRv2-IM-LB”), the God routing protocol (“GodRP”), God routing protocol with link buffering (“GodRP-LB”).

At the beginning of the simulation, all routers are in either “disabled” or “enabled” state randomly with a random `enabled_time` or `disabled_time` left before the next state transition.

D. Simulation Results

Figure 7 depicts the unicast data delivery ratio. The God routing protocol – not surprisingly – has a higher delivery rate

Table II
SIMULATION SETTINGS

Parameter	Value
NS2 version	2.34
Mobility model	PopUp model
Area	1000m x 1000m
Number of routers	20 - 70
Communication range	250m
Radio propagation model	Two-ray ground
Enabled time	25 - 30s
Disabled time	25 - 30s
Move distance	100m
Simulation time	300s
Avg. number of concurrent CBR streams	5
Interface type	802.11b
Radio frequency	2.4 GHz
OLSRv2 parameters	Proposed default values of [?]

than OLSRV2, which is due to the zero-convergence time and less collisions due to the lack of control traffic. However, even with the GodRP, no more than a 35% of delivery ratio with 70 routers is reached. This is because the network, by design, is not fully connected in the PopUp model. The link buffering variants result in significantly higher delivery ratios, with GodRP with link buffering producing the highest data delivery ratio. OLSRV2 with intermediate link buffering produces a data delivery ratio close to that of GodRP-LB. Due to the non-zero convergence time of OLSRV2, buffered datagrams will be delivered only a bit later than with the link buffering variant of the GodRP (which can be observed in figure 8). OLSRV2 without intermediate link buffering has a significantly lower delivery ratio, whereas the graceful shutdown mechanism has a positive effect on the delivery ratio due to the reduced convergence time.




Figure 7. Unicast data delivery ratio

Figure 8 depicts the average delay from the moment a datagram has been sent until it arrives at its destination. The delay largely depends on the selected `enabled_time` and `disabled_time` values. All buffering variants have significantly higher delays than standard OLSRV2 (at the benefit of a higher delivery ratio). The different buffering variants lead to similarly high delay, however, the delay of OLSRV2-IM-LB is significantly higher than the other protocols. This can be explained with the increased routing stretch – paths towards

the destination are longer, as observed in figure 9.




Figure 8. Average delay per datagram

Figure 9 presents the average path length of every data packet from its source to its destination. OLSRV2 has the shortest paths; this can be explained by the lowest delivery ratio amongst the compared protocols (as depicted in figure 7) – destinations further away from the source receive fewer of the transmitted datagrams. For the same reason, the other link buffering variants have a lower path length than GodRP-LB. Comparing GodRP-LB and OLSRV2-IM-LB (which both have similar delivery ratios as shown in figure 7), the path stretch of OLSRV2-IM-LB can be clearly observed.




Figure 9. Path length

Figure 10 shows the accumulated control traffic overhead over the simulation time (GodRP is not included, because the overhead is 0). Not surprisingly, the different link buffering variants of OLSRV2 have no influence on the control traffic. The graceful shutdown mechanism includes transmitting additional HELLOs before shutdown of a router, plus possibly triggered HELLOs and TCs of neighbors of that router.

Figure 11 shows the total number of dropped frames because of collisions during the simulation. For GodRP, almost no collisions appear, since no control traffic congests the network and only little unicast traffic is injected into the network. OLSRV2 without link buffering has the lowest number of collisions, increasing with a growing number of routers in the network – the increasing amount of control traffic leads to most of the collisions. All link buffering variants have a




Figure 10. Accumulated control traffic overhead

higher collision ratio, which can be explained by the fact that once a router reappears (*i.e.* changes its state from “disabled” to “enabled”), all datagrams from different sources that have been buffered while it was “disabled”, are retransmitted almost the same time, as soon as valid routes towards that destination router are added to the routers that had buffered the datagrams. Again, collisions due to increasing control traffic explain the difference between the link buffering variants of OLSRv2 and GodRP-LB with higher number of routers. The bottleneck effect, that has been described in section III-D, can be well observed due to the large number of collision when applying intermediate buffering (OLSRv2-IM-LB).




Figure 11. Collisions

The fact that despite the high number of collision with OLSRv2-IM-LB, the delivery ratio (as depicted in figure 7) is still the highest amongst the OLSRv2 variants, can be explained by the retransmission mechanism of 802.11 [?]. Even if more frames are lost due collisions on L2, eventually one may reach the destination if the frame (*i.e.* the next hop along the route). Figure 12 depicts the total accumulated overhead of

unicast data traffic, when counting each L2 (re-)transmission at every router along the path from source to destination.

It can be seen that OLSRv2-IM-LB produces the highest control traffic overhead: The higher delivery ratio despite the higher collision ratio comes at the expense of available bandwidth on the channel.




Figure 12. Accumulated unicast traffic overhead (counting each transmission at every router)

VI. CONCLUSION

This paper has proposed a “link buffering” mechanism, allowing a router to delay the transmission of an IP datagram if there is no valid route to the destination of that datagram at transmission time, and which is able to complement an off-the-shelf MANET routing protocol for increasing end-to-end delivery ratio of datagrams, at the expense of higher delay. In order to test that mechanism, a network mobility model, denoted the “PopUp model” is specified: routers can be either “enabled” or “disabled”. In the “disabled” state, routers cannot send or receive any data, but are allowed to change position within a bounded distance.

The off-the-shelf MANET routing protocol OLSRv2 is studied, by way of NS2 simulations, with different variations of the link buffering mechanism and when subject to the PopUp model. These simulations have shown that when using a “link buffering” mechanism, the data delivery ratio can be largely improved. Among the variations of the link buffering mechanism are a variant with buffering on layer 3 only (denoted “OLSRv2-L3-LB”), link buffering on layer 2 and 3 (denoted “OLSRv2-LB”), link buffering on intermediate routers (denoted “OLSRv2-IM-LB”), and a variant with a graceful shutdown mechanism (denoted “OLSRv2-LB-gf”). All link buffering variants yield a much higher end-to-end delivery ratio than without using such a mechanism, but also incur an increased end-to-end delay.