

# Towards AODV Version 2

Thomas Clausen, Jiazi Yi, Axel Colin de Verdiere  
Laboratoire d'Informatique - LIX, Ecole Polytechnique, France  
thomas@thomasclausen.org, jiazi@jiaziyi.com, axel@axelcdv.com

**Abstract**—The Ad hoc On-demand Distance-Vector routing protocol (AODV) was published in 2003 by the IETF, as experimental RFC 3561. This routing protocol was one of four routing protocols, developed by the IETF for use in mobile ad hoc networks (MANETs) – with the other being DSR, TBRPF and OLSR. As operational experiences with these protocols accumulated, the IETF set forth on standardization of OLSRv2, a successor to OLSR, and DYMO – with DYMO being the intended successor to DSR and AODV. Alas, while there was traction for and standardization of OLSRv2, interest in, development, standardization, and use of DYMO in MANETs slowly withered.

AODV did, however, attract interest for routing in Low-power Lossy Networks (LLNs) due to its limited state requirements. Since 2005, several proposals for simplifying and adapting AODV specifically for LLNs emerged, in 2011 and 2012 with the use of one such adaptation of AODV in the G3-PLC standard for power line communications in smart grids, and with efforts within the IETF emerging towards a single “AODVv2” specification.

This paper presents this development – from AODV, as specified in RFC3561 – to “AODVv2”. While the basic operation remains unchanged, AODVv2 presents simplifications, and additional features and flexibilities are introduced. This paper studies the impact of these changes “from AODV to AODVv2”, and observes that AODVv2 unites simplification, flexibility and performance improvements.

## I. INTRODUCTION

With the advancements of micro-controller and wireless technology, the concept of “being online” is no longer exclusively reserved for computers, but expected also for phones, vehicles, televisions, refrigerators, utility meters, etc. “The Internet of Things” assumes objects in our environment to be part of the Internet, communicating with users and with each other – and that these objects have communication as a commodity, rather than as their *raison d’être*.

Communication in “The Internet of Things” is a challenge, subject to resource constraints, fragile and low-capacity links, dynamic and arbitrary topologies. Among the challenges is routing, requiring efficient protocols, able to converge rapidly even in very large networks, while exchanging limited control traffic and requiring limited memory and processing power.

### A. Background and History

Since the late 90s, the IETF<sup>1</sup> has embarked upon a path of developing routing protocols for networks with increasingly more fragile and low-capacity links, with less pre-determined connectivity properties and with increasingly constrained router resources. In '97, by chartering the MANET working group, then subsequently in 2006 and 2008 by chartering the 6LowPAN and ROLL working groups.

1) *MANET Protocol Developments*: The MANET working group converged on the development of two protocol families: reactive protocols, including AODV [1], and proactive protocols, including Optimized Link State Routing (OLSR) [2]. A distance vector protocol, AODV operates in an *on-demand* fashion, acquiring and maintaining routes only while needed for carrying data, by way of a *Route Request-Route Reply* exchange. A link state protocol, OLSR is based on periodic control messages exchanges, and each router proactively maintaining a routing table with entries for all destinations in the network, which provides low delays but constant control overhead. A sizable body of work exists, including [3], studying the performance of these protocols in different scenarios, and justifying their complementarity [4]. For the purpose of this paper, it suffices to observe that OLSR provides low delays and predictable, constant control overhead – at expense of requiring memory in each router for maintaining complete network topology. AODV limits the memory required for routing state to that for actively used routes – at the expense of delays for the *Route Request-Route Reply* exchange to take place, and control overhead dependent on data flows.

After acquiring operational experiences, the MANET working group commenced developing successors to OLSR and AODV, denoted OLSRv2 and DYMO. Whereas a relatively large and active community around OLSR thus standardized OLSRv2 [5], [6], [7], [8] and [9], the momentum behind DYMO withered in the MANET working group<sup>2</sup>.

2) *6LowPAN and ROLL Protocol Developments*: The 6LowPAN working group was chartered for adapting IPv6 for operation over IEEE 802.15.4, accommodating characteristics of that MAC layer, and with a careful eye on resource constrained devices (memory, CPU, energy, ...). Part of the original charter for this working group was to develop protocols for routing in multi-hop topologies under such constrained conditions, and over this particular MAC. Two initial philosophies to such routing were explored: *mesh-under* and *route-over*. The former, mesh-under, would, as part of an adaptation layer between 802.15.4 and IP, provide L2.5 multi-hop routing, presenting an underlying mesh-routed multi-hop topology as a single IP link. The latter, route-over, would expose the underlying multi-hop topology to the IP layer, whereupon IP routing would build multi-hop connectivity. Several proposals for routing were presented in 6LowPAN, for each of these philosophies, including LOAD [10]. LOAD was a derivative of AODV, but adapted for L2-addresses and mesh-under routing,

<sup>1</sup><http://www.ietf.org/>

<sup>2</sup><http://tools.ietf.org/wg/manet/minutes?item=minutes81.html>

and with some simplifications over AODV (*e.g.*, removal of intermediate node replies and sequence numbers). However, 6LowPAN was addressing other issues regarding adapting IPv6 for IEEE 802.15.4, such as IP packet header compression, and solving the routing issues was suspended, delegated to a working group ROLL, created in 2008 for this purpose. ROLL produced a routing protocol denoted “Routing Protocol for Low-power lossy networks” (RPL) [11] in 2011.

3) *Finally, Towards AODVv2*: While LOAD [10] development was suspended by the 6LowPAN working group, pending the results from ROLL and experiences with RPL, AODV derivatives live on: IEEE 802.11s [12] is based on AODV, and the G3-PLC standard [13], published in 2011, specifies the use of [10] at the MAC layer, for providing mesh-under routing for utility (electricity) metering networks. Justifications for using an AODV derivative in preference to RPL include that the former better supports bi-directional data flows such as a request/reply of a meter reading [14], as well as algorithmic and code complexity reasons [15]. The emergence of LLNs thus triggered a renewed interest in AODV-derived protocols for specific scenarios, resulting in work within the IETF [16] and [17] for the purpose of standardization of “AODVv2”, incorporating the experiences from deploying AODV – including, but not only, in LLNs.

### B. Statement of Purpose

A successor to AODV, in this paper denoted AODVv2, must support constrained environments, both in terms of computational power and memory, as well as in energy. An example platform is the ESB (Embedded Sensor Board [18]), with a Texas Instruments MSP430 low-power micro-controller, an 1MHz CPU, 2kB RAM and 60kB flash ROM. The link layers typically used in LLNs impose strict limitations on packet sizes: in IEEE 802.15.4, the maximum physical layer packet size is 127 bytes, the resulting maximum frame size at the mac-layer is 102 bytes. If link-layer security is used, this may consume up to a further 21 bytes, which leaves just 81 bytes for upper layer protocols.

Flexibility is non-optional: IPv4 and IPv6 addresses are of different lengths, 6LowPAN [19] proposes “short” addresses of 16 bits, AODV derivatives are used for routing at L2, where address lengths may be different again. Also, extensibility is required, to support *e.g.*, different link metrics, or future optimizations.

These constraints impose challenges to a routing protocol design, and can be summarized as follows: flexible and extensible, small code footprint and low code complexity, low memory and CPU usage – and, of course, low control traffic overhead.

This paper presents AODVv2, designed with the above in mind: a slim core specification satisfying the listed constraints, designed to accommodate extensions. AODVv2 removes many elements from AODV [1] – including intermediate RREPs, gratuitous RREPs, precursor lists, etc., the rationale behind removing these being detailed in this paper. These elements were originally proposed by [1] as performance optimizations

– their removal therefore being justified by a performance comparison between AODV as specified in [1], and AODVv2.

### C. Paper Outline

The remainder of this paper is organized as follows. In section II, the AODVv2 core specification is introduced, including its new packet format and main operations. The simulation study is performed in section III, in which AODV and AODVv2 are compared. Section IV describes the modular approach to be used with AODVv2, and we conclude the paper in section V.

## II. AODVv2 CORE SPECIFICATION

A reactive distance-vector protocol, AODVv2 inherits the basic properties and operations of AODV, including generation of Route Requests (RREQs) by a router (originator) for discovering a route to a destination, forwarding of such RREQs until they reach the destination router, generation of Route Replies (RREPs) upon receipt of a RREQ by the indicated destination, and unicast hop-by-hop forwarding of these RREPs towards the originator. If a route is detected broken, *i.e.*, if forwarding of a data packet to the recorded next hop on the route to the destination is detected to fail, local route repair can be attempted, or a Route Error (RERR) message can be returned to the originator of that data packet. Compared to AODV, AODVv2 contains both extensions and simplification.

Extensions in AODVv2:

- AODVv2 is modular. The core specification defines the simple and light-weight core functions of the protocol. AODVv2 is extensible, by way of a flexible packet format permitting addition of arbitrary attributes and information via new message types and/or TLVs.
- Optimized Flooding is explicitly supported, reducing the overhead incurred by RREQ forwarding. Jitter is employed, to reduce the probability of losses due to collisions on lower layers [5].
- Address lengths from 1-16 octets are supported<sup>3</sup>. The only requirement is, that within a given routing domain, all addresses are of the same address length.
- Different metrics are supported, to make better use of link information from different layers.

Simplifications in AODVv2:

- Only the destination is permitted to respond to an RREQ; intermediate routers are explicitly prohibited from responding to RREQs, even if they may have active routes to the sought destination. All messages (RREQ or RREPs) generated by a given router share a single unique, monotonically increasing sequence number. This also eliminates Gratuitous RREPs while ensuring loop freedom. The rationale for this simplification is reduced complexity of protocol operation and reduced message sizes – found to be without significant influence in the performance in section III.

<sup>3</sup>*i.e.*, IPv6, IPv4, 6LowPAN short addresses, L2 MAC addresses etc. are all supported by AODVv2

- An AODVv2 router does not maintain a precursor list, thus when forwarding of a data packet to the recorded next hop on the route to the destination fails, a RERR is sent only to the originator of that data packet. The rationale for this simplification is an assumption that few overlapping routes are in use concurrently, and delay is not a critical issue in a given network.

#### A. Packet Format

AODVv2 defines four types of routing packets:

##### Route Request (RREQ)

Generated by a router, when presented with a data packet to a destination, for which it has no valid route, and containing the address of the destination for that data packet.

##### Route Reply (RREP)

Generated by a router, when it receives and processes a RREQ containing an address for which the router is responsible<sup>4</sup> as a response to a RREQ.

##### Route Reply Acknowledgement (RREP-ACK)

Generated by an AODVv2 router as a response to a RREP, in order to signal to the neighbor which transmitted the RREP that the RREP was successfully received.

##### Route Error (RERR)

Generated by a router when a link on an active route to a destination is detected as broken, by way of inability to forward a data packet towards that destination.

Figure 1(a) shows the general format of all types of packets, using the notation from [6]:




Figure 1. AODVv2 packet and message format examples

Where <tlv-block> contains zero or more Type-Length-Value elements (TLVs), and <message> field is an RREQ, RREP, RREP-ACK or RERR message. An example of RREQ/RREP message is illustrated in figure 1(b). The <metric> field specifies the function, by way of which two messages (say, RREQs) are compared so as to identify which has come through the “best” path. Such a function may be specified so as to use any information available, e.g., <route-cost><sup>5</sup>, <weak-links><sup>6</sup> and various TLVs, and must define how a router compares two received messages so as to determine which of them represent the “best” path.

<sup>4</sup>i.e., an address of a destination, local to that router

<sup>5</sup>In its simplest form, a basic hop count metric.

<sup>6</sup>Links deemed, e.g., by a lower layer, to be marginally useful and which therefore should be avoided if at all possible.

AODVv2 allows protocol extensions to define additional message types, as well as additional TLV types, for inclusion in the message types defined in the core specification. Combined with the ability to define a <metric> using these TLVs, AODVv2 is thus accommodating for both functional extensions as well as refinements of core protocol operations for deployments with specific requirements.

#### B. Main Protocol Operations

AODVv2 retains the basic protocol operations from AODV, including *Route Discovery* and *Route Maintenance*, albeit in a greatly simplified form.

1) *Route Discovery*: During *Route Discovery*, RREQ messages are flooded through the network. In AODVv2, only the router, which is responsible for the address listed in the RREQ will respond with an RREP, sent in unicast to the source of the RREQ.

This is a departure from AODV, in which an intermediate router, having a (valid) route to the address listed in the RREQ, would send a RREP to the source of the RREQ, as well as a gratuitous RREP to the intended destination (both as unicast). This simplification in AODVv2 permits eliminating intermediate RREPs, gratuitous RREPs, Destination Sequence Number and Originator Sequence Number in RREQ messages, reducing the per-packet overhead – and without significant influence in the performance.

2) *Route Maintenance*: *Route Maintenance* is performed when an actively used route fails. Route failure is detected by way of a data packet not being deliverable to the next hop towards the intended destination. In AODVv2, when a route failure is detected, an RERR message is generated, sent as unicast along the route to the source of data packet. On receiving the RERR at the source of data packet, a new *Route Discovery* should be performed.

This is a departure from AODV, in which each router maintains a list of “precursors”, containing the IP address for each of its neighbors that are likely to use it for a next hop. When a router detects that data delivery over a link fails, it informs all neighbors, listed in this “precursor list”. The simplification in AODVv2 of employing end-to-end signaling only permits eliminating the precursor list.

### III. SIMULATIONS AND RESULTS

In order to understand, in particular, the performance impact of the extensions and simplifications, present in AODVv2, ns2-simulation results are presented in the below. Simulations were made of a field of 1000 × 1000 meters, with varying numbers of routers placed randomly. For the purpose of this study, router mobility was not considered. The network scenarios were subject to two different traffic patterns: point-to-point traffic (P2P), where (source, destination) pairs were chosen at random in the network, and multipoint-to-point (MP2P), where all routers generate traffic, for which the destination always is a single, fixed router in the network. The latter, e.g., simulating a data collection network.

Each data source transmits five 512-byte data packets per second, in bursts lasting for 80 seconds each, for a total simulation time of 100s. For each data-point in the figures below, 50 randomly generated scenarios, representing the same abstract parameters (density, number of traffic flows) and the results presented are the average hereof.

### A. Point-to-point traffic

Figure 2 shows the packet delivery ratios, the average end-to-end delays, the number of collisions occurring in the network, and the routing overhead (in number of bytes), incurring from respectively AODV and AODVv2, when the networks were subject to 30 concurrent (source, destination) pairs. Packet delivery ratios and delays are almost identical between the two version of the protocols, whereas AODVv2 incurs both fewer collisions (due to the use of jitter on flooding) and less control traffic overhead (due to smaller control packets) than does AODV.

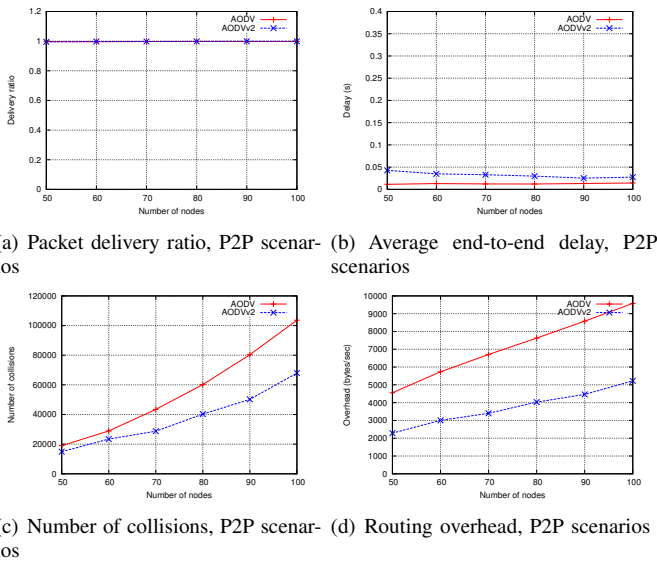


Figure 2. P2P Simulation Results - AODV and AODVv2

### B. Multipoint-to-point traffic

Figure 3 shows the packet delivery ratios, the average end-to-end delays, and the routing overhead (in number of bytes), incurring from respectively AODV and AODVv2, in the MP2P scenarios, *i.e.*, where the network is subject to a heavier data traffic load:  $n - 1$  routers, all sending traffic to a single other router in the network. AODVv2 yields at least 10% higher data delivery ratios than does AODV, while incurring a dramatically lower control traffic overhead.

### C. General Observations

The simulations presented accord some general observations: AODV, by way of using “intermediate RREPs” and expanding ring flooding for seeking to shorten delays, end up generating a few fewer control packets than does AODVv2. However as AODVv2’s control packets are significantly

smaller than those of AODV, the control traffic overhead of AODVv2 is substantially lower than that of AODV.

AODV does yield a shorter average delay than does AODVv2 - as shown in 3(b) for 80 routers or more. Note that in calculating the average delays, only data packets, successfully being delivered to their intended destination, are accounted - and in those high-density scenarios, AODVv2 delivers ~50% more packets than does AODV.

## IV. AODV COMPANION MODULES

As mentioned in section I, AODVv2 is designed to be modular. The aim is that the AODVv2 “core” specification be as slim as possible: small code footprint, low state and packet overhead, low processing - not to mention easy to implement and test. This, while remaining as general as possible: providing the ability to discover and maintain usable, short paths between any (source, destination) pair.

Specific scenarios (*e.g.*, traffic patterns, density) and deployments (*e.g.*, public or private, link characteristics, other protocols running on the routers) may benefit from specific optimizations. AODVv2 is designed so as to be able to accommodate these as plug-in modules, either by way of introducing new packet types, or by way of introducing additional TLVs to AODVv2 control packets (and specifying their processing) - or simply by specifying additional in-router data structures and processing. This section indicates a few such plug-in modules, without intending to be exhaustive.

1) “Precursor Lists”: was one feature, removed from AODV in creating AODVv2. If many concurrently active routes are expected to transit routers at the same time, and that therefore a link breakage may disrupt many concurrent data flows, it may be justifiable for a router to sacrifice state and complexity to track these flows. In case of a link breakage detected, such a router generate and forward multiple RERR, each towards one of the the sources of these flows. A plug-in module, introducing “Precursor Lists”, introduces additional state and processing in a router - which will be generating RERR messages, processed and forwarded by other routers as per AODVv2. As this entails no additional signaling, does not impact interoperability, and is beneficial only in specific scenarios, it is not part of the “core” specification.

2) “Path Accumulation”: for enabling source routing, an idea extracted from [20] and also used in [11] relieving intermediate routers from maintaining routing tables for the duration of the data traffic flow<sup>7</sup>. This functionality can be introduced in AODVv2 by way of introducing a TLV to a RREQ, to be inserted by the source of a RREQ and specifying that it be updated by each intermediate router such that the value of said TLV when received by the destination contains the succession of intermediate routers that have forwarded the RREQ. The <flags> field (figure 1(a)) allows that the RREQ source instructs intermediate routers to “drop TLVs you don’t understand before forwarding them”, and TLV processing may

<sup>7</sup>Trading off ‘bytes over the media’ by way of source routing headers on each packet, for ‘bytes in memory’....




Figure 3. MP2P Simulation Results - AODV and AODVv2

specify that the routing table entries, recorded by RREQ/RREP be expired after a short time even if the routes they represent are being used.

3) *Multi-path routing, Multicast tree construction, etc.:* can likewise be supported by way of adding TLVs and specifying TLV processing, to provide backup route or building multicast tree in the network.

## V. CONCLUSION

This paper has presented AODVv2, a proposed ultra-light-weight, extensible and flexible successor to AODV, applicable also in very constrained environments: low overhead, low computational and memory complexity. A preliminary performance evaluation has been conducted, which has shown AODVv2 yields comparative performance to that of AODV – except that AODVv2 incurs a substantially lower control traffic overhead.

## ACKNOWLEDGMENTS

The authors would like to thank Ulrich Herberg (Fujitsu Labs of America), Cedric Lavenu (EDF R&D), Thierry Lys (ERDF), Yuichi Igarashi and Satoh Hiroki (Hitachi YRL), as well as Charles E. Perkins (Tellabs) and Afshin Niktash (Maxim Integrated Products) for stimulating discussions and feedback on the work presented in this paper.

## REFERENCES

- [1] C. Perkins, E. Belding-Royer, and S. Das, “Ad hoc On-Demand Distance Vector (AODV) Routing,” Experimental RFC 3561, July 2003.
- [2] T. Clausen and P. Jacquet, “Optimized Link State Routing Protocol (OLSR),” Experimental RFC 3626, October 2003.
- [3] L. V. T. Clausen, P. Jacquet, “Comparative study of routing protocols for mobile ad-hoc networks.” Proceedings of the IFIP MedHocNet, September, Sardinia, Italy, 2002.
- [4] T. Clausen, P. Jacquet, and L. Viennot, “Analyzing control traffic overhead versus mobility and data traffic activity in mobile ad hoc network protocols,” *ACM Journal on Wireless Networks*, vol. 10 no. 4, 2004.
- [5] T. Clausen, C. Dearlove, and B. Adamson, “Jitter Considerations in MANETs,” IETF Inf. RFC 5148, February 2008.
- [6] T. Clausen, C. Dearlove, J. Dean, and C. Adjih, “Generalized MANET Packet/Message Format,” Std. Track RFC 5444, February 2009.
- [7] T. Clausen and C. Dearlove, “Representing Multi-Value Time in MANETs,” IETF Std. Track RFC 5497, February 2009.
- [8] T. Clausen, C. Dearlove, and J. Dean, “Mobile Ad Hoc Network Neighborhood Discovery Protocol,” Std. Track RFC 6130, April 2010.
- [9] T. Clausen, C. Dearlove, and P. Jacquet, “The Optimized Link State Routing Protocol version 2,” Internet Draft, draft-ietf-manet-olsrv2-11, work in progress, April 2010.
- [10] K. Kim, S. D. Park, G. Montenegro, S. Yoo, and N. Kushalnagar, “6LoWPAN Ad Hoc On-Demand Distance Vector Routing,” June 2007, Internet Draft, work in progress, draft-daniel-6lowpan-load-adhoc-routing-03.
- [11] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, and J. Vasseur, “RPL: IPv6 Routing Protocol for Low power and Lossy Networks,” March 2011, Internet Draft, work in progress, draft-ietf-roll-rpl-19.
- [12] G. Hiertz, S. Max, R. Zhao, D. Denteneer, and L. Berlemann, “Principles of IEEE 802.11s,” in *Proceedings of WiMAN in conjunction with the 16th ICCCN*, Honolulu, Hawaii, USA, Aug 2007, p. 6.
- [13] “ITU-T G.9956: Narrow-Band OFDM power line communication transceivers - Data link layer specification,” November 2011.
- [14] U. Herberg and T. Clausen, “A comparative performance study of the routing protocols load and rpl with bi-directional traffic in low-power and lossy networks (lln).” Proceedings of the 8th ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN), October 2011.
- [15] T. Clausen, U. Herberg, and M. Philipp, “A critical evaluation of the “IPv6 routing protocol for low power and lossy networks.”” Proceedings of the 5th IEEE International Conference on Wireless & Mobile Computing, Networking & Communication (WiMob), October 2011.
- [16] T. Clausen, A. C. de Verdiere, J. Yi, A. Niktash, Y. Igarashi, H. Satoh, and U. Herberg, “The lln on-demand ad hoc distance-vector routing protocol - next generation,” The Internet Engineering Task Force, October 2011, internet Draft, work in progress, draft-clausen-lln-loadng.
- [17] T. Clausen, A. Camacho, J. Yi, A. C. de Verdiere, Y. Igarashi, H. Satoh, and Y. Morii, “Experience with the loadng routing protocol for llns,” The Internet Engineering Task Force, October 2011, internet Draft, work in progress, draft-lavenu-lln-loadng-interoperability-report.
- [18] “The ESB Embedded Sensor Board,” <http://www.sics.se/~adam/contiki/contiki-2.0-doc/a00435.html>, 2011.
- [19] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler, “Transmission of IPv6 Packets over IEEE 802.15.4 Networks,” September 2007, Standards Track RFC 4944.
- [20] D. Johnson, Y. Hu, and D. Maltz, “The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks for IPv4,” IETF Std. RFC 4728, February 2007.