

Networking Working Group
Internet-Draft
Intended status: Informational
Expires: March 29, 2009

P. Levis
Stanford University
A. Tavakoli
S. Dawson-Haggerty
UC Berkeley
September 25, 2008

Overview of Existing Routing Protocols for Low Power and Lossy Networks [draft-ietf-roll-protocols-survey-01](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on March 29, 2009.

Abstract

Networks of low power wireless devices introduce novel IP routing issues. Low-power wireless devices, such as sensors, actuators and smart objects, have difficult constraints: very limited memory, little processing power, and long sleep periods. As most of these devices are battery-powered, energy efficiency is critically important. Wireless link qualities can vary significantly over time, requiring protocols to make agile decisions yet minimize topology change energy costs. Routing over such low power and lossy networks has novel requirements that existing protocols may not address. This document provides a brief survey of the strengths and weaknesses of

existing protocols with respect to this class of networks. From this survey it examines whether existing protocols as described in RFCs and mature drafts could be used without modification in these networks, or whether further work is necessary.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Table of Contents

1.	Terminology	3
2.	Introduction	3
3.	Methodology	4
4.	Suitability Summary	5
4.1.	Formal Definitions	6
4.2.	Table Scalability	6
4.3.	Loss Response	7
4.4.	Control Cost	8
4.5.	Link and Node Cost	8
5.	Routing Protocol Taxonomy	9
6.	Link State Protocols	11
6.1.	OSPF	11
6.2.	OLSR	12
6.3.	TBRPF	12
7.	Distance Vector protocols	13
7.1.	RIP	13
7.2.	Ad-hoc On Demand Vector Routing (AODV)	13
7.3.	DYMO	14
7.4.	DSR	14
8.	Neighbor Discovery	14
8.1.	IPv6 Neighbor Discovery	14
8.2.	MANET-NHDP	15
9.	Security Issues	15
10.	Manageability Issues	15
11.	IANA Considerations	15
12.	Security Considerations	15
13.	Acknowledgements	15
14.	Annex A - Routing protocol scalability analysis	15
15.	References	19
15.1.	Normative References	19
15.2.	Informative References	19
	Authors' Addresses	20
	Intellectual Property and Copyright Statements	22

1. Terminology

AODV: Ad-hoc On Demand Vector Routing

DSR: Dynamic Source Routing

DYMO: Dynamic Mobile On-Demand

LLN: Low power and Lossy Network

LSA: Link State Advertisement

LSDB: Link State Database

MANET: Mobile Ad-hoc Networks

MAC: Medium Access Control

MPLS: Multiprotocol Label Switching

MPR: Multipoint Relays

MTU: Maximum Transmission Unit

OLSR: Optimized Link State Routing

ROLL: Routing in Low power and Lossy Networks

TDMA: Time Division Multiple Access

2. Introduction

Wireless is increasingly important to computer networking. As Moore's Law has reduced computer prices and form factors, networking includes not only servers and desktops, but laptops, palmtops, and cellphones. As computing device costs and sizes have shrunk, small wireless sensors, actuators, and smart objects have emerged as an important next step in inter-networking. The sheer number of the low-power networked devices means that they cannot depend on human intervention (e.g., adjusting position) for good networking: they must have routing protocols that enable them to self-organize into multihop networks.

Energy is a fundamental challenge in these devices. Convenience and ease of use requires they be wireless and therefore battery powered. Correspondingly, low power operation is a key concern for this class of networked device. Cost points and energy limitations cause these

devices to have very limited resources: a few kB of RAM and a few MHz of CPU is typical. As energy efficiency does not improve with Moore's Law, these limitations are not temporary. This trend towards smaller, lower power, and more numerous devices has led to new low-power wireless link layers to support them. In practice, wireless networks observe much higher loss rates than wired ones do, and low-power wireless is no exception. Furthermore, many of these networks will include powered as well as energy constrained nodes. Nevertheless, for cost and scaling reasons, many of these powered devices will still have limited resources.

These low power and lossy networks introduce constraints and requirements that other networks typically do not possess ([\[I-D.ietf-roll-home-routing-reqs\]](#) and [\[I-D.ietf-roll-indus-routing-reqs\]](#)). As they were not designed with these requirements in mind, existing protocols may or may not work well in LLNs. The first step to reaching consensus on a routing protocol for LLNs is to decide which of these two is true. If an existing protocol can meet LLN requirements without any changes, then barring extenuating circumstances, it behooves us to use an existing standard. However, if no current protocol can meet LLN's requirements, then further work will be needed to define and standardize with a protocol that can. Whether or not such a protocol involves modifications to an existing protocol or a new protocol entirely is outside the scope of this document: this document simply seeks to answer the question: do LLNs require a new protocol specification document at all?

[3.](#) Methodology

To answer the question of LLNs require new protocol specification work, this document examines existing routing protocols and how well they can be applied to low power and lossy networks. It provides a set of criteria with which to compare the costs and benefits of different protocol designs and examines existing protocols in terms of these criteria.

The five criteria this document uses are derived from a set of drafts that describe the requirements of a few major LLN application scenarios. The five criteria, presented in [Section 3](#), are neither exhaustive nor complete. Instead, they are one specific subset of high-level requirements shared across all of the application requirement drafts. Because every application requirement draft specifies these criteria, then a protocol which does not meet one of them cannot be used without modifications or extensions. However, because these criteria represent a subset of the intersection of the application requirements, any given application domain may impose

additional requirements which a particular protocol may not meet. For this reason, these criteria are "necessary but not sufficient." A protocol that does not meet the criteria cannot be used as specified, but it is possible that a protocol meets the criteria yet is not able to meet the requirements of a particular application domain. Nevertheless, a protocol that meets all of the criteria would be very promising, and deserve a closer look and consideration in light of LLN application domains.

This document considers "existing routing protocols" to be protocols that are specified in RFCs or, in the cases of DYMO [[I-D.ietf-manet-dymo](#)] or OLSRV2 [[I-D.ietf-manet-olsrv2](#)] , a very mature draft which will most likely become an RFC. This document does not seek to answer the question of whether there is any protocol anywhere which could meet LLN application requirements. Rather, it seeks to answer whether protocols, as specified in current IETF standards documents, can meet such requirements. If an existing protocol specification can be used unchanged, then writing additional protocol specifications is unnecessary. For example, there are many academic papers and experimental protocol implementations available; while one or more of these may meet LLN requirements, if they are not specified in an RFC then a working group will need to write a new RFC for them to be a standard. The question this document seeks to answer is not whether proposed, evaluated, theoretical or hypothetical protocol designs can satisfy LLN requirements: the question is whether existing IETF standards can.

Whether a protocol meets these criteria was judged by thinking through each specification and considering the best implementation possible. The judgement is based on what a specification allows, rather than any particular implementation of that specification. For example, while many DYMO implementations use hopcount as a routing metric, the DYMO specification allows a hop to add more than one to the routing metric, so DYMO as a specification can support some links or nodes being more costly than others.

4. Suitability Summary

In this section, we present five important requirements for routing in low power and lossy networks, and evaluate protocols against them. This evaluation attempts to take a complicated and interrelated set of design decisions and trade-offs and condense them to a simple "pass", "fail", or "?". As with any simplification, there is a risk of removing some necessary nuance. However, we believe that being forced to take a position on whether or not these protocols are acceptable according to binary criterion will be constructive.

We derive these criteria from existing documents that describe ROLL network application requirements. These metrics do not encompass all application requirements. Instead, they are a common set of routing protocol requirements that most applications domains share.

Considering this very general and common set of requirements sets a minimal bar for a protocol to be generally applicable. If a protocol cannot meet even these minimalist criteria, then it cannot be used in several major ROLL application domains and so is unlikely to be a good candidate for further analysis and examination. Satisfying these minimal criteria is necessary but not sufficient: they do not represent the complete intersection of application requirements and applications introduce additional, more stringent requirements. But this simplified view provides a first cut of the applicability of existing protocols, and those that do satisfy them might be reasonable candidates for further study.

The five criteria are "table scalability", "loss response", "control cost", "link cost", and "node cost". For each of these, the value "pass" indicates that a given protocol has satisfactory performance according to the metric. The value "fail" indicates that the protocol does not have acceptable performance according to the metric, and that the RFC defining the protocol does not, as written, contain sufficient flexibility to alter the protocol to do so. Finally, "?" indicates that an implementation could exhibit satisfactory performance while following the RFC, but that the implementation descisions necessary to do so are not specified and may require some exploration. In other words, a "fail" means a protocol would have to be modified so it is not compliant with its RFC in order to meet the criterion, while a "?" means a protocol would require a supplementary document further constraining and specifying how a protocol should behave.

4.1. Formal Definitions

To provide precise definitions of these metrics, we use formal big-O notation, where N refers to the number of nodes in the network, D refers to the number of unique destinations, and L refers to the size of a node's local, single-hop neighborhood (the network density). We explain the derivation of each metric from application requirements in its corresponding section.

4.2. Table Scalability

Scalability support for large networks of sensors is highlighted as a key requirement by all three application requirements documents. Network sizes range from a minimum of 250 nodes in the home routing requirements [[I-D.ietf-roll-home-routing-reqs](#)] to very large networks of "tens of thousands to millions" of devices noted of the urban

requirements [[I-D.ietf-roll-urban-routing-reqs](#)]. Networks are expected to have similar size in industrial settings, the requirements draft states that depths of up to 20 hops are to be expected [[I-D.ietf-roll-indus-routing-reqs](#)]. Given that network information maintained at each node is stored in routing and neighbor tables, along with the constrained memory of nodes, necessitates bounds on the size of these tables.

This metric examines whether routing tables scale within reasonable memory resources of low-power nodes. According to this metric, routing protocols that scale linearly with the size of the network or a node's neighborhood fail. Scaling with the size of the network prevents networks from growing to reasonable size, while scaling with the network density precludes dense deployments. However, as many low-power and lossy networks behave principally as data collection networks and principally communicate through routers to data collection points in the larger Internet, scaling with the number of such collection points is reasonable. Protocols whose state scales with the number of destinations pass.

More precisely, routing table size scaling with $O(N)$ or $O(L)$ fails. A table that scales $O(D)$ (assuming no N or L) passes.

[4.3.](#) Loss Response

In low power and lossy networks, links routinely come and go due to being close to the SINR threshold. It is important that link churn not trigger unnecessary responses by the routing protocol. This point is stressed in all the application requirement documents, pointing to the need to localize response to link failures with no triggering of global network re-optimization, whether for reducing traffic or for maintaining low route convergence times ([[I-D.ietf-roll-home-routing-reqs](#)], [[I-D.ietf-roll-urban-routing-reqs](#)], and [[I-D.ietf-roll-indus-routing-reqs](#)]). The industrial routing requirements draft states that protocols must be able to "recompute paths based on underlying link characteristics which may change dynamically", as well as reoptimize when the device set changes to maintain service requirements. The protocol should also "always be in the process of optimizing the system in response to changing link statistics." Protocols with these properties should take care not to require global updates.

A protocol which requires many link changes to propagate across the entire network fails. Protocols which constrain the scope of information propagation to only when they affect routes to active destinations, or to local neighborhoods, pass. Protocols which allow proactively path maintenance pass if the choice of which paths to

maintain is user-specified.

More precisely, loss responses that require $O(N)$ transmissions fail, while responses that can rely on $O(1)$ local broadcasts or $O(D)$ route updates pass.

[4.4.](#) Control Cost

Battery-operated devices are a critical component of all three application spectrums, and as such special emphasis is placed on minimizing power consumption to achieve long battery lifetime, [[I-D.ietf-roll-home-routing-reqs](#)], with multi-year deployments being a common case [[I-D.ietf-roll-indus-routing-reqs](#)]. In terms of routing structure, any proposed L2N routing protocol ought to support the autonomous organization and configuration of the network at the lowest possible energy cost [[I-D.ietf-roll-urban-routing-reqs](#)].

All routing protocols must transmit additional data to detect neighbors, build routes, transmit routing tables, or otherwise conduct routing. As low-power wireless networks can have very low data rates, protocols which require a minimum control packet rate can have unbounded control overhead. This is particularly true for event-driven networks, which only report data when certain conditions are met. Regions of a network which never meet the condition can be forced to send control traffic even when there is no data to send. For these use cases, hard-coded timing constants are unacceptable, because they imply a prior knowledge of the expected data rate.

If control traffic is unbounded by data traffic, a protocol fails according to Control Cost metric. Protocols which pass bound their control traffic rate to their data traffic. Protocols which pass do not use resources to maintain unused state. More specifically, any protocol which requires fixed-rate periodic control packets in the absence of data traffic fails.

[4.5.](#) Link and Node Cost

These two metrics specify how a protocol chooses routes for data packets to take through the network. Classical routing algorithms typically acknowledge the differing costs of paths and may use a shortest path algorithm to find paths. This is a requirement for low power networks, as links must be evaluated as part of an objective function across various metric types, such as minimizing latency and maximizing reliability [[I-D.ietf-roll-indus-routing-reqs](#)].

However, in low power networks it is also desirable to account for the cost of routing through particular routers. Applications require node or parameter constrained routing, which takes into account node

properties and attributes such as power, memory, and battery life that dictate a router's willingness or ability to route other packets. Home routing requirements note that devices will vary in their duty cycle, and that routing protocols should prefer nodes with permanent power [[I-D.ietf-roll-home-routing-reqs](#)]. The urban requirements note that routing protocols may wish to take advantage of differing data processing and management capabilities among network devices [[I-D.ietf-roll-urban-routing-reqs](#)]. Finally, industrial requirements cite differing lifetime requirements as an important factor to account for [[I-D.ietf-roll-indus-routing-reqs](#)]. Node cost refers to the ability for a protocol to incorporate router properties into routing metrics and use node attributes for constraint-based routing.

A "pass" indicates that the protocol contains a mechanism allowing these considerations to be considered when choosing routes.

5. Routing Protocol Taxonomy

Routing protocols broadly fall into two classes: link-state and distance-vector.

A router running a link-state protocol first establishes adjacency with its neighbors and then reliably floods the local topology information in the form of a Link State Advertisement packet. The collection of LSAs constitutes the Link State Database (LSDB) that represents the network topology, and routers synchronize their LSDBs. Thus each node in the network has a complete view of the network topology. Each router uses the LSDB to compute a routing table where each entry (reachable IP destination address) points to the next hop along the shortest path according to some metric. Link state protocols (such as OSPF and IS-IS) support the concept of area (called "level" for IS-IS) whereby all the routers in the same area share the same view (they have the same LSDB) and areas are interconnected by border routers according to specific rules that advertise IP prefix reachability between areas.

A distance vector protocol exchanges routing information rather than topological information. A router running a distance vector protocol exchanges information with its "neighbors" with which it has link layer connectivity. Tunneling and similar mechanisms can virtualize link layer connectivity to allow neighbors that are multiple layer 2 hops away. Rather than a map of the network topology from which each router can calculate routes, a distance vector protocol node has information on what routes its neighbors have. Each node's set of available routes is the union of its neighbors routes plus a route to itself. In a distance vector protocol, nodes may only advertise

routes which are in use, enabling on-demand discovery. In comparison to link state protocols, distance vector protocols have the advantage of only requiring neighbor routing information, but also have corresponding limitations which protocols must address, such as routing loops, count to infinity, split horizon, and slow convergence times. Furthermore, routing constraints are difficult to enforce with distance vector protocols.

Neighbor discovery is a critical component of any routing protocol. It enables a protocol to learn about which other nodes are nearby and which it can use as the next hop for routes. As neighbor discovery is a key component of many protocols, several general protocols and protocol mechanisms have been designed to support it. A protocol's neighbor set is defined by how many "hops" away the set reaches. For example, the 1-hop neighbor set of a node is all nodes it can directly communicate with at the link layer, while the 2-hop neighbor set is its own 1-hop neighbor set and the 1-hop neighbor sets of all of its 1-hop neighbors.

Because nodes often have very limited resources for storing routing state, protocols cannot assume that they can store complete neighbor information. For example, a node with 4kB of RAM cannot store full neighbor state when it has 1000 other nodes nearby. This means that ROLL protocols must have mechanisms to decide which of many possible neighbors they monitor as routable next hops. For elements such as 2-hop neighborhoods, these decisions can have a significant impact on the topology that other nodes observe, and therefore may require intelligent logic to prevent effects such as network partitions.

Protocols Today

Wired networks draw from both approaches. OSPF or IS-IS, for example, are link-state protocols, while RIP is a distance-vector protocol.

MANETs similarly draw from both approaches. OLSR is a link-state protocol, while AODV and DYMO are distance vector protocols. The general consensus in core networks is to use link state routing protocols as IGPs for a number of reasons: in many cases having a complete network topology view is required to adequately compute the shortest path according to some metrics. For some applications such as MPLS Traffic Engineering it is even required to have the knowledge of the Traffic Engineering Database for constraint based routing.

Furthermore link state protocols typically have superior convergence speeds (ability to find an alternate path in case of network element failure), are easier to debug and troubleshoot, and introduce less control packet overhead than distance vector protocols. In contrast,

distance vector protocols are simpler, require less computation, and have smaller storage requirements. Most of these tradeoffs are similar in wireless networks, with one exception. Because wireless links can suffer from significant temporal variation, link state protocols can have higher traffic loads as topology changes must propagate globally, while in a distance vector protocol a node can make local routing decisions with no effect on the global routing topology. One major protocol, DSR, does not easily fit into one of these two classes. Although it is a distance vector protocol, DSR has several properties that make it differ from most other protocols in this class. We examine these differences in our discussion of DSR.

The next two sections summarize several well established routing protocols. This table shows, based on the criteria described above, whether these protocols meet ROLL criteria. Annex A contains the reasoning behind each value in the table.

Protocol	Table	Loss	Control	Link Cost	Node Cost
OSPF	fail	fail	fail	pass	fail
OLSRv2	fail	fail	fail	pass	pass
TBRPF	fail	pass	fail	pass	?
RIP	pass	fail	fail	?	?
AODV	pass	fail	pass	fail	fail
DYMO[-low]	pass	fail	pass	?	fail
DSR	fail	pass	pass	fail	fail

6. Link State Protocols

6.1. OSPF

OSPF (specified in [[RFC2328](#)] for IPv4 and in [[RFC2740](#)] for IPv6)) is a link state protocol designed for routing within an Internet Autonomous System (AS). OSPF provides the ability to divide a network into areas, which can establish a routing hierarchy. The topology within an area is hidden from other areas and IP prefix reachability across areas (inter-area routing) is provided using summary LSAs. The hierarchy implies that there is a top-level routing area (the backbone area) which connects other areas. Areas may be connected to the back-bone area through a virtual link. OSPF maintains routing adjacencies by sending hello messages. OSPF calculates the shortest path to a node using link metrics (that may reflect the link bandwidth, propagation delay, ...). OSPF Traffic Engineering (OSPF-TE, [[RFC3630](#)]) extends OSPF to include information on reservable, unreserved, and available bandwidth.

6.2. OLSR

Optimized Link State Routing (OLSR) (see [[RFC3626](#)] and [[I-D.ietf-manet-olsrv2](#)]) is a link state routing protocol for wireless mesh networks. OLSR nodes flood route discovery packets throughout the entire network, such that each node has a map of the mesh topology. Because link variations can lead to heavy flooding traffic when using a link state approach, OLSR establishes a topology for minimizing this communication. Each node maintains a set of nodes called its Multipoint Relays (MPR), which is a subset of the one-hop neighbors whose connectivity covers the two-hop neighborhood. Each node that is an MPR maintains a set called its MPR selectors, which are nodes that have chosen it to be an MPR.

OLSR uses these two sets to apply three optimizations. First, only MPRs generate link state information. Second, nodes can use MPRs to limit the set of nodes that forward link state packets. Third, an MPR, rather than advertise all of its links, can advertise only links to its MPR selectors. Together, these three optimizations can greatly reduce the control traffic in dense networks, as the number of MPRs should not increase significantly as a network becomes denser.

OLSR selects routes based on hop counts, and assumes an underlying protocol that determines whether a link exists between two nodes. OLSR's constrained flooding allows it to quickly adapt to and propagate topology changes.

OLSR is closely related to clustering algorithms in the wireless sensor networking literature, in which cluster heads are elected such that routing occurs over links between cluster heads and all other nodes are leafs that communicate to a cluster head.

6.3. TBRPF

Topology Dissemination Based on Reverse Path Forwarding (see [[RFC3684](#)]) is another proactive link state protocol. TBRPF computes a source tree, which provides routes to all reachable nodes. It reduces control packet overhead by having nodes only transmit a subset of their source tree as well as by using differential updates.

The major difference between TBRPF and OLSR is the routing data that nodes advertise and who chooses to aggregate information. In OLSR, nodes select neighbors to be MPRs and advertise their link state for them; in TBRPF, nodes elect themselves to advertise relevant link state based on whether it acts as a next hop.

7. Distance Vector protocols

7.1. RIP

The Routing Information Protocol (RIP) (defined in [[RFC2453](#)]) predates OSPF. As it is a distance vector protocol, routing loops can occur and considerable work has been done to accelerate convergence since the initial RIP protocols were introduced. RIP measures route cost in terms of hops, and detects routing loops by observing a route cost approach infinity where "infinity" is referred to as a maximum number of hops. RIP is typically not appropriate for situations where routes need to be chosen based on real-time parameters such as measured delay, reliability, or load or when the network topology needs to be known for route computation.

"Triggered RIP" (defined in [[RFC2091](#)]) was originally designed to support "on-demand" circuits. The aim of triggered RIP is to avoid systematically sending the routing database on regular intervals. Instead, triggered RIP sends the database when there is a routing update or a next hop adjacency change: once neighbors have exchanged their routing database, only incremental updates need to be sent. Because incremental updates cannot depend on periodic traffic to overcome losses, triggered RIP uses acknowledgment based mechanisms for reliable delivery.

7.2. Ad-hoc On Demand Vector Routing (AODV)

AODV (specified in [[RFC3561](#)]) is a distance vector protocol intended for mobile ad-hoc networks. When one AODV node requires a route to another, it floods a request in the network to discover a route. A depth-scoped flooding process avoids discovery from expanding to the most distant regions of the network that are in the opposite direction of the destination. AODV chooses routes that have the minimum hop count.

If an AODV route request reaches a node that has a route to the destination (this includes the destination itself), that node sends a reply along the reverse route. All nodes along the reverse route can cache the route. When routes break due to topology changes, AODV floods error messages and issues a new request. Because AODV is on-demand it only maintains routes for active nodes. When a link breaks, AODV issues a Route Error (RERR) and a new route request message (RREQ), with a higher sequence number so nodes do not respond from their route caches. These packets can flood the entire network, giving loss response a fail.

7.3. DYMO

Dynamic Mobile On-Demand routing (DYMO) ([\[I-D.ietf-manet-dymo\]](#)) is an evolution of AODV. The basic functionality is the same, but it has different packet formats, handling rules, and supports path accumulation. Path accumulation allows a single DYMO route request to generate routes to all nodes along the route to that destination. Like AODV, DYMO uses hop counts as its routing metric, but links may have a cost ≥ 1 , allowing DYMO to represent link costs. Like AODV, on link breaks DYMO issues a new route request message (RREQ), with a higher sequence number so nodes do not respond from their route caches. Correspondingly, a route request can flood the entire network.

7.4. DSR

Dynamic Source Routing ([\[RFC4728\]](#)) is a distance vector protocol, but a DSR packet source explicitly specifies the route for each packet. Because the route is determined at a single place -- the source -- DSR does not require sequence numbers or other mechanisms to prevent routing loops, as there is no problem of inconsistent routing tables. Unlike AODV and DYMO, by pushing state into packet headers, DSR does not require per-destination routing state. Instead, a node originating packets only needs to store a spanning tree of the part of the network it is communicating with.

8. Neighbor Discovery

A limit on maintained routing state (light footprint) prevents ROLL protocols from assuming they know all 1-hop, 2-hop, or N-hop neighbors. For this reason, while protocols such as MANET-NHDP ([\[I-D.ietf-manet-nhdp\]](#)) and IPv6's neighbor discovery ([\[RFC4861\]](#)) provide basic mechanisms for discovering link-layer neighbors, not all of their features are relevant. This section describes these two protocols, their capabilities, and how ROLL protocols could leverage them.

8.1. IPv6 Neighbor Discovery

IPv6 neighbor discovery provides mechanisms for nodes to discover single-hop neighbors as well as routers that can forward packets past the local neighborhood. There is an implicit assumption that the delegation of whether a node is a router or not is static (e.g., based on a wired topology). The fact that all routers must respond to a Router Solicitation requires that the number of routers with a 1-hop neighborhood is small, or there will be a reply implosion. Furthermore, IPv6 neighbor discovery's support of address

autoconfiguration assumes address provisioning, in that addresses reflect the underlying communication topology. IPv6 neighbor discovery does not consider asymmetric links. Nevertheless, it may be possible to extend and adapt IPv6's mechanisms to wireless in order to avoid response storms and implosions.

8.2. MANET-NHDP

The MANET Neighborhood Discovery Protocol (MANET-NHDP) provides mechanisms for discovering a node's symmetric 2-hop neighborhood. It maintains information on discovered links, their interfaces, status, and neighbor sets. MANET-NHDP advertises a node's local link state; by listening to all of its 1-hop neighbor's advertisements, a node can compute its 2-hop neighborhood. MANET-NHDP link state advertisements can include a link quality metric. MANET-NHDP's node information base includes all interface addresses of each 1-hop neighbor: for low-power nodes, this state requirement can be difficult to support.

9. Security Issues

TBD

10. Manageability Issues

TBD

11. IANA Considerations

This document includes no request to IANA.

12. Security Considerations

TBD

13. Acknowledgements

14. Annex A - Routing protocol scalability analysis

This aim of this Annex is to provide the details for the analysis routing scalability analysis.

"OSPF"

OSPF floods link state through a network. Each router must receive this complete link set. OSPF fails the table size criterion because it requires each router to discover each link in the network, for a total routing table size which is $O(N * L)$. This also causes it to fail the control cost criterion, since this information must be propagated. Furthermore, changes in the link set require re-flooding the network link state even if the changed links were not being used. Since link state changes in wireless networks are often uncorrelated with data traffic and are instead caused by external (environmental) factors, this causes OSPF to fail both the control cost and loss response criteria. OSPF routers can impose policies on the use of links and can consider link properties (Type of Service), as the cost associated with an edge is configurable by the system administrator [[RFC2328](#)], so receive a pass for link cost. However, there is no way to associate metrics with routers (as costs are only applied to outgoing interfaces, i.e. edges) when computing paths, and so fails the node cost criteria. While [[RFC3630](#)] discusses paths that take into account node attributes, it specifically states that no known algorithm or mechanism currently exists for incorporating this into the OSPF RFC.

"OLSRv2"

OLSRv2 is a proactive link state protocol, flooding this information through a set of multipoint relays (MPRs). Routing state includes 1-hop neighbor information for each node in the network, 1-hop and 2-hop information for neighbors (for MPR selection), and a routing table (consisting of destination, and next hop), resulting in state proportional to network size and density ($O(N*L + L^2)$), and failing the table scalability criteria.

Unacceptable control traffic overhead arises from flooding and maintenance. HELLO messages are periodically broadcast local beacon messages, but TC messages spread topology information throughout the network (using MPRs). As such, control traffic is proportional to $O(N^2)$. MPRs reduce this load to $O(N^2 / L)$. As the number of MPRs is inversely proportional to the density of the network and L is bounded by N , this means control traffic is at best proportional to $O(N)$, and fails the control cost metric.

Furthermore, changes in the link set require immediately re-flooding the network link state even if those links were not being used by routing, which fails the loss response metric.

OLSR allows for specification of link quality, and also provides a 'Willingness' metric to symbolize node cost, giving it a pass for

both those metrics.

"TBRPF"

As a link state protocol where each node maintains a database of the entire network topology, TBRPF's routing table size scales with network size and density, leading to table sizes which are $O(N * L)$ when a node receives disjoint link sets from its neighbors. This causes the protocol to fail the table size criteria. The protocol's use of differential updates should allow both fast response time and incremental changes once the distributed database of links has been established. Differential updates are only used to reduce response time to changing network conditions, not to reduce the amount of topology information sent, since each node will periodically send their piece of the topology. As a result, TBRPF fails the control overhead criteria. However, its differential updates triggered by link failure do not immediately cause a global re-flooding of state (but only to affected routers) [[RFC3684](#)], leading to a pass for loss response.

TBRPF has a flexible neighbor management layer which enables it to incorporate various types of link metrics into its routing decision by enabling a USE_METRIC flag [[RFC3684](#)]. As a result, it receives a pass for link cost. It also provides a mechanism whereby routers can maintain multiple link metrics to a single neighbor, some of which can be advertised by the neighbor router [[RFC3684](#)]. Although the RFC does not specify a policy for using these values, developing one could allow TBRPF to satisfy this requirement, leading to a ? for the node cost requirement.

"RIP"

RIP is a distance vector protocol: all routers maintain a route to all other routers. Routing table size is therefore $O(N)$. However, if destinations are known apriori, table size can be reduced to $O(D)$, resulting in a pass for table scalability. Each node broadcasts a beacon per period, and updates must be propagated by affected nodes, irrespective of data rate, failing the control cost metric. Loss triggers updates, only propagating if part of a best route, but even if the route is not actively being used, resulting in a fail for loss response. The rate of triggered updates is throttled, and these are only differential updates, yet this still doesn't account for other control traffic (or tie it to data rate) or prevent the triggered updates from being flooded along non-active paths. [[RFC2453](#)]

RIP receives a ? for link cost because while current implementations focus on hop count and that is the metric used in [[RFC2453](#)], the RFC also mentions that more complex metrics such as differences in

bandwidth and reliability could be used. However, the RFC also states that real-time metrics such as link-quality would create instability and the concept of node cost only appears as metrics assigned to external networks. It also receives a ? because the concept of a network cost is introduced, which is added to link cost, but does not describe its use.

"AODV"

AODV table size is a function of the number of communicating pairs in the network, scaling with $O(D)$. This is acceptable and so AODV passes the table size criteria. As an on-demand protocol, AODV does not generate any traffic until data is sent, and so control traffic is correlated with the data and so it receives a pass for control traffic. When a broken link is detected, AODV will use a precursor list maintained for each destination to inform downstream routers (with a RERR) of the topology change. However, the RERR message is forwarded by all nodes that have a route that uses the broken link, even if the route is not currently active, leading to a fail for loss response [[RFC3561](#)].

AODV fails the link cost metric because the only metric used is hop count, and this is hardcoded in the route table entry, according to the RFC [[RFC3561](#)]. It fails the node cost requirement because there is no way for a router to indicate its [lack of] willingness to route while still adhering to the RFC.

"DYMO/DYMO-low"

The design of DYMO shares much with AODV, with some changes to remove precursor lists and compact various messages. It still passes the table size criteria because it only maintains routes requested by RREQ messages, resulting in $O(D)$ table size. Control traffic (RREQ, RREP, and RREQ) are still driven by data, and hence DYMO passes the control cost criterion. However, RERR messages are forwarded by any nodes that have a route using the link, even if inactive, leading to a fail of the loss response criteria [[I-D.ietf-manet-dymo](#)].

While DYMO does indicate that the metric used for a link can vary from 1-65535, it specifically refers to this as distance, which is incremented by at least one at each hop [[I-D.ietf-manet-dymo](#)], leading to a ? in link cost. While additional routing information can be added DYMO messages, there is no mention of node cost, leading to a fail in node cost.

"DSR"

DSR performs on-demand route discovery, and source routing of

packets. It maintains a source route for all destinations, and also a blacklist of all unidirectional neighbor links [[RFC4728](#)], leading to a total table size of $O(D + L)$, failing the table size criterion. Control traffic is completely data driven, and so DSR receives a pass for this criteria. Finally, a transmission failure only prompts an unreachable destination to be sent to the source of the message, passing the loss response criteria.

DSR fails the link cost criterion because its source routes are advertised only in terms of hops, such that all advertised links are considered equivalent. DSR also fails the node cost criterion because a node has no way of indicating its willingness to serve as a router and forward messages.

[15. References](#)

[15.1. Normative References](#)

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[15.2. Informative References](#)

[I-D.ietf-manet-dymo]

Chakeres, I. and C. Perkins, "Dynamic MANET On-demand (DYMO) Routing", [draft-ietf-manet-dymo-14](#) (work in progress), June 2008.

[I-D.ietf-manet-nhdp]

Clausen, T., Dearlove, C., and J. Dean, "MANET Neighborhood Discovery Protocol (NHDP)", [draft-ietf-manet-nhdp-07](#) (work in progress), July 2008.

[I-D.ietf-manet-olsrv2]

Clausen, T., Dearlove, C., and P. Jacquet, "The Optimized Link State Routing Protocol version 2", [draft-ietf-manet-olsrv2-07](#) (work in progress), July 2008.

[I-D.ietf-roll-home-routing-reqs]

Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirement in Low Power and Lossy Networks", [draft-ietf-roll-home-routing-reqs-03](#) (work in progress), September 2008.

[I-D.ietf-roll-indus-routing-reqs]

Networks, D., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low Power and Lossy

Networks", [draft-ietf-roll-indus-routing-reqs-01](#) (work in progress), July 2008.

[I-D.ietf-roll-urban-routing-reqs]

Dohler, M., Watteyne, T., Winter, T., Jacquenet, C., Madhusudan, G., Chegaray, G., and D. Barthel, "Urban WSNs Routing Requirements in Low Power and Lossy Networks", [draft-ietf-roll-urban-routing-reqs-01](#) (work in progress), July 2008.

[RFC2091] Meyer, G. and S. Sherry, "Triggered Extensions to RIP to Support Demand Circuits", [RFC 2091](#), January 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.

[RFC2453] Malkin, G., "RIP Version 2", STD 56, [RFC 2453](#), November 1998.

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", [RFC 2740](#), December 1999.

[RFC3561] Perkins, C., Belding-Royer, E., and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing", [RFC 3561](#), July 2003.

[RFC3626] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", [RFC 3626](#), October 2003.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.

[RFC3684] Ogier, R., Templin, F., and M. Lewis, "Topology Dissemination Based on Reverse-Path Forwarding (TBRPF)", [RFC 3684](#), February 2004.

[RFC4728] Johnson, D., Hu, Y., and D. Maltz, "The Dynamic Source Routing Protocol (DSR) for Mobile Ad Hoc Networks for IPv4", [RFC 4728](#), February 2007.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.

Authors' Addresses

Philip Levis
Stanford University
358 Gates Hall, Stanford University
Stanford, CA 94305-9030
USA

Email: pal@cs.stanford.edu

Arsalan Tavakoli
UC Berkeley
Soda Hall, UC Berkeley
Berkeley, CA 94707
USA

Email: arsalan@eecs.berkeley.edu

Stephen Dawson-Haggerty
UC Berkeley
Soda Hall, UC Berkeley
Berkeley, CA 94707
USA

Email: stevedh@cs.berkeley.edu

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

