

A New Approach to On-Demand Loop-Free Routing in Ad Hoc Networks

J.J. Garcia-Luna-Aceves
Computer Engineering Dept.
Baskin School of Engineering
University of California
Santa Cruz, CA 95064
jj@soe.ucsc.edu

Marc Mosko^{*}
Computer Engineering Dept.
Baskin School of Engineering
University of California
Santa Cruz, CA 95064
mmosko@soe.ucsc.edu

Charles E. Perkins
Communications Systems
Laboratory
Nokia Research Center
Mountain View, CA 94303
charliep@iprg.nokia.com

ABSTRACT

A new protocol is presented for on-demand loop-free routing in ad hoc networks. The new protocol, called labeled distance routing (LDR) protocol, uses a distance invariant to establish an ordering criterion and per-destination sequence numbers to reset the invariant resulting in loop-freedom at every instant. The distance invariant allows nodes to change their next hops or distances to destinations without creating routing-table loops. The destination sequence number, which only the destination may increment, permits nodes to reset the values of their distance invariants. The performance of LDR is compared against the performance of three other protocols that are representative of the state-of-the-art, namely AODV, DSR and OLSR.

Keywords

routing, loop-free, ad hoc network, distance invariant

1. INTRODUCTION

An ad hoc wireless network is characterized by nodes functioning as routers, as well as sources and sinks of data traffic, with radio network interfaces and no fixed infrastructure to support communications. Wireless networks usually have limited bandwidth and battery power, so their routing protocols should have low control overhead. Reactive or on-demand routing protocols have been developed for this reason. In an on-demand routing protocol, a node only maintains routes for in-use destinations and does not pro-actively advertise routes. Rather, it queries for needed routes and offers routes in response to queries.

^{*}This work was supported in part by the Advanced Technology Office of the Defense Advanced Research Projects Agency (DARPA) under grant No. DAAD19-01-C-0026 and by the U.S. Air Force/OSR under grant No. F49620-00-1-0330.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODC'03, July 13–16, 2003, Boston, Massachusetts, USA.

Copyright 2003 ACM 1-58113-708-7/03/0007...\$5.00.

Many on-demand routing protocols have been proposed over the past few years, and all of them attempt to provide loop-free paths at every instant through various methods. Examples of such protocols include the Ad hoc On-demand Distance Vector (AODV) protocol [12], the Dynamic Source Routing (DSR) protocol [8, 9], the Neighborhood-aware Source Routing (NSR) protocol [16], the Temporally-Ordered Routing Algorithm (TORA) protocol [11], and the Routing On-demand Acyclic Multipath (ROAM) protocol [15]. There are also many examples of pro-active routing protocols that attain loop-free routing [7] as well as pro-active routing protocols that tolerate temporary routing loops [6, 2, 10].

Two factors that decrease routing efficiency in ad hoc networks are routing loops and the maintenance of complete reachability information for all nodes. Routing loops increase packet-delivery latencies and reduce the number of packets delivered to the intended destinations. Furthermore, maintaining routing information for all destinations becomes less efficient than maintaining routing information on-demand at each node as the number of nodes in the network increases and the average number of destinations contacted by each source becomes a smaller and smaller portion of the total number of nodes. Accordingly, this paper addresses the problem of providing efficient on-demand loop-free routing in ad hoc networks. Prior work on on-demand loop-free routing has been based on the following approaches: (a) using source routes in data packets, (b) coordinating nodes by way of the directed graph implied by the next hop entry for a given destination at each node, and (c) using sequence numbers to establish an ordering among nodes.

DSR and NSR are examples of using source routing to avoid routing loops. In DSR, each route request records its traversed path, and the complete route between source and the requested destination is specified in the route reply sent back to the source by any node with a valid route to the destination. The discovered route is stored in a route cache at the origin and the relay nodes. The header of every data packet specifies the source routes to their intended destinations. DSR incorporates a number of optimizations to shorten source routes, learn routes, and cache routes efficiently. NSR extends the source routing approach of DSR by having nodes communicate information regarding their two-hop neighborhood in route requests and route replies path in addition to path information regarding specific in-

use destinations.

TORA uses a link-reversal algorithm [4] to maintain loop-free multipaths that are created by a query-reply process similar to that used in DSR and AODV. TORA relies on synchronized clocks to create time stamps that maintain the relative ordering of events. The link-reversal algorithm is a form of synchronization among nodes spanning multiple hops.

The Diffusing Update Algorithm (DUAL) [7] attains loop freedom by means of nodal coordination. Based on routing advertisements, a node i may choose independently a successor j from its neighbors for a destination d as long as that choice cannot form a loop; such a choice is called feasible if it cannot form a loop. If the potential exists to form a loop, node i must coordinate with other nodes before changing its route to d . The feasibility of a change is calculated based on a feasibility condition. If the condition is satisfied, the change is feasible, and node i may proceed without coordinating with any other node. If the change is not feasible, node i starts a diffusing computation [3]. DUAL was developed for wire-line networks and the diffusing computation requires reliable communications between neighbors to enforce synchronization over potentially large segments of a network. The diffusing computation resets information in the network, ensuring that any change at i for d will not cause a loop. After the computation terminates, node i is free to make a change and then inform the network of that change. ROAM extends DUAL to provide loop-free routing on demand. The distances to a given destination are used to establish ordering among nodes. ROAM uses a route reply-request process similar to that used in AODV, DSR and other on-demand protocols. However, a node can change its next hop to a destination without notifying its neighbors as long as it has a neighbor with a distance that is shorter than the node's own feasible distance value to the destination, where such a distance is the smallest value attained by the node's distance since it obtained a route to the destination after sending a route request. If such an invariant condition is not satisfied, the node must reliably send a route request to its neighbors, which serves the same purpose of DUAL's resets. After sending a route request, the node cannot select a new next hop to a destination until it receives route replies from all its neighbors.

ROAM and TORA require reliable exchanges among neighbors and coordination among nodes over multiple hops. Routes are locked down at a node until the portion of the distributed calculation in which it participates is complete, which is signaled to the node when all its neighbors reply to its route request. This type of mechanisms incurs more control messages compared to AODV, DSR, and other on-demand protocols that work correctly even with unreliable transmissions of route requests and replies among neighbors.

AODV attains loop-free routing by using a sequence number for each destination as the means to establish an ordering invariant among nodes. AODV defines an active route as one that is fresh and likely to have a good successor path. Such a route was either recently learned through some advertisement or has been recently used without error. An inactive (or invalid) route is one that has expired its cache time without use or one for which the next hop is in an error state (e.g., lack of connectivity or downstream route failure). The sequence numbers of active routes for a given destination are non-increasing moving away from the destination.

When a node A needs to establish a route to a destination D , it broadcasts a route request to its neighbors. If A previously knew a route to D that became invalid, A increases the sequence number for D and includes it in the route request. A node receiving the request can send back a unicast route reply along its shortest path to node A only if it has an active route to D and the sequence number stored for D is no less than the sequence number in the route request. Otherwise, the recipient must forward the route request. When node A sends a route request for a destination, it increases the sequence number for itself as well, which is used by other nodes that learn about new routes to node A .

By increasing the stored sequence number for a destination when a route breaks, AODV ensures that no upstream nodes may reply to a route request from a node on their successor path. Unfortunately, this also inhibits responses from downstream nodes with the prior sequence number, even if they have a valid loop-free path to the destination. A key limitation to using only sequence numbers as the loop-free routing invariant is that it inhibits responses both moving away from the destination and moving toward the destination. In many cases, this causes the destination to be the only node able to satisfy the request, because it alone can increase its own sequence number for a response.

We present the labeled distance routing protocol (LDR), which is an on-demand routing protocol that uses distance labels rather than sequence numbers, source routing, or inter-nodal coordination to ensure loop freedom at every instant. Section 2 describes LDR through two examples and presents the feasibility conditions and protocol procedures. Like any other on-demand routing protocol, LDR discovers routes through the network only when there are data for a destination.

Section 3 shows that LDR works correctly and is loop free at every instant. LDR uses a loop-free invariant for each destination similar to that first introduced in DUAL [7] and employs sequence numbers that can be incremented only by the destinations themselves to permit nodes to reset the values of their distance invariants. The combined use of a distance invariant and destination-controlled sequence numbers eliminates the need for inter-nodal coordination used in DUAL and other loop-free routing protocols, and enables more efficient responses to route requests compared to AODV, which helps reduce network load.

Section 4 presents the results of simulation experiments comparing LDR with AODV, DSR, and OLSR. The other three protocols were used as points of comparison given that they are representatives of the state-of-the art in routing for ad hoc networks, and are being considered by the working group on mobile ad hoc networks (MANET) of the Internet Engineering Task Force (IETF). The simulation results clearly show that LDR attains higher packet delivery ratios, smaller packet latencies, and much lower signaling overhead than the other three protocols.

2. LABELED DISTANCE ROUTING PROTOCOL

In our description of LDR, we assume positive symmetric link costs; asymmetric costs would require extra packet fields to communicate the costs upstream, but would not change the basic operation of LDR.

LDR uses a route request (RREQ), route reply (RREP),

Notation	
sn_D^A	The sequence number of D as known at node A .
d_D^A	The measured distance from node A to D . If all link costs are 1, it is a hop count.
fd_D^A	The feasible distance from node A to D , being the minimum d_D^A for the current sn_D^A .
lc_B^A	The link cost from node A to neighbor B , assumed to be positive and equal to unity if using hop count metrics.
*	An advertisement, for example sn_D^* is the sequence number in an advertisement for destination D .
#	A solicitation, for example $sn_D^\#$ is the sequence number in a solicitation for destination D . Each issuer adds its own unique identifier $rreqid$.
$rr_D^\#$	Reset required bit (T bit) for solicitation # for destination D . Indicates that an invariant ordering violation could occur and the path must be reset.

and route error (RERR) messaging structure that is based on that of AODV. We use the term *advertisement* to denote the portion of a packet that proffers reachability to a destination, and the term *solicitation* to denote the portion of a packet that requests information for a destination. The RREQ in LDR and other on-demand protocols constitute both an advertisement of a route to the node issuing the RREQ and a solicitation for a route to another node. In general, we will discuss advertisements and solicitations as separate entities apart from their concrete realizations in RREQs or RREPs. Table 1 summarizes the notation used to describe LDR in the rest of this paper.

The RREQ is the tuple $\{dst, sn_{dst}, rreqid, src, sn_{src}, fd, dist, flags\}$, where src is the identifier of the source of the RREQ seeking a path to the destination with identifier dst . The sequence numbers for the destination and source are sn_{dst} and sn_{src} , respectively. The $rreqid$ field is a source-specific unique identifier to control the flooding of the RREQ. The source's feasible distance is fd , and the measured distance of the path traversed by the RREQ is $dist$. Control bits are contained in $flags$.

The RREP is the tuple $\{dst, sn_{dst}, src, rreqid, dist, lifetime, flags\}$. The field $lifetime$ is the milli-seconds of time remaining for the route to dst and reflects the maximum time to cache the route if it is not used.

For a given destination D for which node A has a route, it maintains the sequence number originated by D (sn_D^A), its distance to D (d_D^A), its next hop to D , and its feasible distance to D (fd_D^A). The feasible distance fd_D^A is the minimum distance d_D^A ever known to D for the current sequence number sn_D^A .

2.1 Sufficient Conditions for Loop Freedom

The key to loop-freedom in previous works using feasible distances (DUAL [7], LPA [5], PDA [17]) is an invariant condition and some form of inter-nodal synchronization used when such a condition is not satisfied. DUAL states three invariant conditions, with the *source node condition* (SNC) being similar to our numbered distance condition. SNC is a minimum-cost condition that renders loop-free shortest

paths. LPA uses a condition equivalent to SNC. PDA uses the *loop-free invariant condition* (LFI), which relaxes the shortest-path requirement of SNC to multiple loop-free successors. When nodes cannot satisfy the feasibility condition in these algorithms, they begin a co-ordinated computation to prevent loops while resetting their distances to higher values. In DUAL, a node enters an *active* state for a destination and requires that all up-stream nodes that potentially use it as a successor either accept the proposed larger distance or change their successor. The node then resets the feasible distance to infinity and chooses a new successor. When a node using LPA or PDA wishes to change its successor and detects that doing so could cause a loop, it co-ordinates with all its immediate neighbors to ensure that none can continue to use it as a successor to the destination. The reliable co-ordination with neighbors preempts any possible loops.

The key to loop-freedom in LDR is the dissemination of route requests over a tree, forcing route replies to follow paths in that tree, and enforcing a strict ordering of feasible distances along successor paths. The RREQ tree is formed by the conventional reverse-path flooding technique of AODV. The route reply paths are constructed by looking up the RREQ ID in the RREQ cache and sending route replies only along the reverse path of the flood. Strict ordering of feasible distances for a given destination is attained by ensuring that the following conditions are satisfied.

Numbered Distance Condition (NDC): Node A may accept a route advertisement from neighbor B for destination D and update its routing table independently of other nodes if A has no information about destination D or either one of the following two conditions is satisfied:

$$sn_D^B > sn_D^A \quad (1)$$

$$sn_D^B = sn_D^A \wedge d_D^B < fd_D^A. \quad (2)$$

Feasible Distance Condition (FDC): Node I must set $rr_D^\# = 1$ in a relayed solicitation for destination D that it forwards if $sn_D^I = sn_D^\#$ and $fd_D^I \geq fd_D^\#$.

Start Distance Condition (SDC): Node I may initiate an advertisement for a solicitation from node A for destination D if I has an active route to D , and either of the following conditions is satisfied:

$$sn_D^I > sn_D^\# \quad (3)$$

$$sn_D^I = sn_D^\# \wedge d_D^I < fd_D^\# \wedge \neg rr_D^\# \quad (4)$$

NDC is used to allow nodes to change successors without coordination among nodes. FDC is used to enforce ordering of the feasible distances of all the nodes along a path to a destination. SDC is used to allow a node that does not have a neighbor satisfying NDC to find a distant node that can provide a loop-free path to the destination. More specifically, a RREQ specifies the feasible distance of the node that originated the request and that nodes sequence number. Subject to TTL restrictions, nodes relay a RREQ, until it reaches a node that satisfies SDC, and that node issues a RREP. Any node along the path taken by the RREQ that does not satisfy FDC sets the T bit to prevent nodes closer to the destination and with a smaller feasible distance from sending a RREP. Requests with the T bit set require a path reset so a node with higher sequence number must reply.

THEOREM 1. *Using NDC at node A to update successors for destination D independently of other nodes is sufficient to ensure that no loop is created.*

Proof: Let neighbor I send an advertisement to A for destination D . If $sn_D^I > sn_D^A$, then in the absence of node failures, node A cannot be on node I 's path to destination D , for otherwise it would have known of the higher sequence number. If the sequence numbers are the same, NDC is equivalent to SNC from DUAL, which is shown to be loop-free [7, Theorem 1, pp. 132ff]. ■

PROPOSITION 1. *Using SDC at node I to initiate an advertisement for a solicitation from node A for destination D does not create loops.*

The proof of the above proposition is immediate from Theorem 1 and FDC. If a node I issues an advertisement that is not feasible, it would be rejected by node A according to NDC. Furthermore, a path from A to D can be discovered successfully only if each node in the path from A to D satisfy FDC.

THEOREM 2 (ORDERING CRITERIA). *FDC ensures that, along a successor path $P = \{n_k, \dots, n_1\}$ from node n_k to node n_1 , it is always true for $i \in [k, 2]$ that $(sn_{n_i}^{n_i} < sn_{n_i}^{n_{i-1}}) \vee (sn_{n_i}^{n_i} = sn_{n_i}^{n_{i-1}} \wedge fd_{n_i}^{n_i} > fd_{n_i}^{n_{i-1}})$.*

Proof: For any given sequence number for node n_1 , $fd_{n_i}^{n_i} \leq d_{n_i}^{n_1}$. According to NDC, for equal sequence numbers, node n_i can use node n_{i-1} as a successor only if $d_{n_i}^{n_{i-1}} < fd_{n_i}^{n_i}$, which implies that $fd_{n_i}^{n_{i-1}} < fd_{n_i}^{n_i}$. At some later time, node n_{i-1} could change its successor, but that can only increase the sequence number or decrease the feasible distance, or remain the same. ■

Note that because stable paths to destinations are preferable, if node A already has an active route to destination D , it should not change its successor after receiving an advertisement from node B if $sn_D^B = sn_D^A$, unless its distance to D through node B can be reduced, i.e., $d_D^B + lc_B^A < d_D^A$.

2.2 Route Discovery

Conditions NDC, FDC, and SDC avoid the creation of routing-table loops. However, they do not guarantee the establishment of a path from a node sending a solicitation to the intended destination, even if a physical path does exist. The reason is that, for the same sequence number for a given destination, no node is allowed to send an advertisement in response to a solicitation after the T bit is set according to conditions NDC, FDC and SDC.

To remedy this shortcoming, LDR augments FDC and provides a reset operation in the event that the T bit must be set. If the RREQ/RREP path would violate the ordering of feasible distances, i.e., FDC is not satisfied, the RREQ is unicast to the destination, so that it can increase its own sequence number and the RREP then can reset feasible distances along the path and thus maintaining ordering. The node that unicasts the RREQ to destination D is the first node along the path followed by the RREQ that satisfies SDC without consideration to the T bit. That node must ensure the RREQ's TTL is sufficient to reach the destination because in an expanding ring search, the broadcast RREQ might not have enough time-to-live left.

A given node A enters into a route computation for destination D when it issues a solicitation for D with identifier




Figure 1: Example using LDR

ID_A (the *rreqid*). Such a node is called *active* for D in computation (A, ID_A) . For a given destination, a node may have at most one active computation. The computation (A, ID_A) terminates when A receives any feasible advertisement for D or a timer expires. If A receives a feasible advertisement for A , the computation terminates in success, otherwise timer expiry indicates failure. The termination of computation (A, ID_A) is a local event at A and does not imply that all solicitations for (A, ID_A) are out of the network or that intermediate nodes participating in the computation have terminated their engagement.

If a node relays the solicitation, it participates in the computation (A, ID_A) and must cache certain data for a period of time. Such a node is said to be *engaged* in (A, ID_A) . A relay node must record the tuple $\{A, ID_A, lasthop\}$, where *lasthop* identifies the previous hop node participating in the computation. An engaged node terminates computation (A, ID_A) at the expiry of a timer. A node may be engaged in multiple computations, but may only enter the engaged state once per computation (A, ID_A) . For the same computation (A, ID_A) the node A may not be both active and engaged (i.e., it may not relay its own solicitation). Note that a relay node has no state about the destination D of a computation and in fact may go active for D while being engaged in other computations for D . A node that is neither active or engaged in a computation (A, ID_A) is said to be *passive* for (A, ID_A) . This is the default state of a node.

PROCEDURE 1 (INITIATE SOLICITATION). *A node A that requires a route for destination D first checks to see if it is active for D. If it is, A should queue the packet that requires the route. If A is not active for D, it becomes active and increments its rreqid. Let ID_A be the incremented identifier. A issues a solicitation for D identified by (A, ID_A) and starts a timer with expiry $t = 2 \cdot ttl \cdot latency$, where *ttl* is the time-to-live of the broadcast flood and *latency* is the estimated per-hop latency of the network. If the timer expires, A may retry the solicitation and increase the *ttl* based on the network policies. If after the final attempt, A does not find a route to D, A should inform the packet origins of the failure and drop the queued packets.*

PROCEDURE 2 (RELAY SOLICITATION). *A node B that receives a solicitation (A, ID_A) for destination D firsts checks to see if it is passive for (A, ID_A) . If it is not passive, it silently ignores the solicitation. If it is passive, it becomes engaged. If B satisfies SDC, it may issue an advertisement for D. Otherwise, B relays the solicitation. Let the last hop be node C (possibly equal to A) and let the new solicitation be denoted by \ddagger . Node B must cache the tuple $\{A, ID_A, C\}$ for a sufficient period of time such that all solicitation instances of (A, ID_A) have left the network and any advertisements in*

response to (A, ID_A) have had time to complete.

$$sn_D^\dagger \leftarrow \begin{cases} sn_D^B & \text{if } sn_D^B > sn_D^\# \\ sn_D^\# & \text{otherwise} \end{cases} \quad (5)$$

$$fd_D^\dagger \leftarrow \begin{cases} fd_D^B & \text{if } sn_D^B > sn_D^\# \\ \min\{fd_D^B, fd_D^\#\} & \text{if } sn_D^B = sn_D^\# \\ fd_D^\# & \text{otherwise} \end{cases} \quad (6)$$

$$d_S^\dagger \leftarrow d_S^\# + lc_C^B \quad (7)$$

$$rr_D^\dagger \leftarrow \begin{cases} 0 & \text{if } sn_D^B > sn_D^\# \\ rr_D^\# & \text{if } (fd_D^B < fd_D^\#) \wedge (sn_D^B = sn_D^\#) \\ 1 & \text{otherwise} \end{cases} \quad (8)$$

Eq. 8 controls the path-reset request mechanism. The first condition resets the T bit to zero when B 's sequence number exceeds the requested sequence number. This is because B has increased the requested sequence number by Eq. 5, so that any advertisement sent in response to the solicitation functions as a path reset. The second condition reflects that B matches the ordering criteria, and therefore it relays the path-reset value that already exist. The third condition sets the T bit when B violates the ordering criteria.

If the destination receives a solicitation with the T bit set, it must reset the path. If $sn_D^D > sn_D^\#$, it may use the current sequence number to reset the path. Otherwise, it must increment its stored sequence number before sending the advertisement. As with all solicitations, the destination will only send one advertisement per source and rreqid pair.

When a node relays a solicitation, it has to record the last hop of the solicitation, because it will use that last hop as part of a reverse path. When a node relays a RREP, it will look up the (originator, rreqid) pair and force the RREP to follow the RREQ reverse path, even if the node has an active route to the originator. This reverse path information may be kept in the RREQ cache.

As an advertisement progresses through the network, it may happen that a relay node has stronger invariants than those contained in the advertisement. In such a case, the relay node must discard the advertisement and issue a new advertisement, if possible. It may also be the case that the relay node has stronger invariants, but an invalid route. In such as case, the relay cannot issue a new advertisement. At other times, a node relaying a route request may need to update the invariants of the request, such that the relaying node may use any solicited reply.

In the following, when a node drops an advertisement and issues a new one with stronger invariants or strengthens the invariants in a solicitation, we call this "updating the invariants," which is an oxymoron but succinct.

PROCEDURE 3 (SET ROUTE). *When node A adds or updates a route to destination D via successor B, it updates its sequence number, distance, and feasible distance for destination D.*

$$sn_D^A \leftarrow sn_D^\dagger \quad (9)$$

$$d_D^A \leftarrow d_D^\dagger + lc_B^A \quad (10)$$

$$fd_D^A \leftarrow \begin{cases} d_D^A & \text{if } sn_D^A < sn_D^\dagger \\ \min\{fd_D^A, d_D^A\} & \text{if } sn_D^A = sn_D^\dagger \end{cases} \quad (11)$$

Procedure 3 guarantees that the feasible distance for a given sequence number is a non-increasing function over

time for a given node. The distance for a given sequence number and feasible distance may fluctuate, but is never less than the feasible distance. This behavior prevents loops.

If a node has an active route to the destination of an advertisement and is not itself the terminus of the advertisement, the node should issue a new advertisement for the route. If the node does not have an active route to the destination (because it could not update its routing table based on NDC), the node must not relay the advertisement.

It is possible for a RREQ to propagate through the network without creating the reverse path. If a node relays a RREQ without having an active reverse path to the RREQ origin, the relay node must set the (newly specified) N bit to indicate that the RREQ is no longer an advertisement for the RREQ origin. The N bit is not part of the current AODV specification. If the node replying to the RREQ does not have a reverse path, it sets the new corresponding N bit in the RREP indicating such. When the origin receives a RREP with the N bit set, it may send a unicast RREQ probe along its forward path with the D bit set. It should increase its sequence number to ensure that the reverse path is built. Nodes otherwise should not increase their sequence number when issuing a RREQ.

PROCEDURE 4 (RELAY ADVERTISEMENT). *If node A is not the terminus of the advertisement (e.g., the source address in a RREP), and it has an active route to destination D, node A should issue a new advertisement for D upon receipt of an advertisement for the destination. Node A may create or update its own routing table by Procedure 3 upon receiving an advertisement, and uses its RREQ cache to ensure that it does not forward more than one reply per (originator, rreqid) pair. Let the new advertisement be denoted by \dagger , then $sn_D^\dagger \leftarrow sn_D^A$, $d_D^\dagger \leftarrow d_D^A$.*

2.3 Example

Figure 1 shows the directed acyclic successor graph for destination T in a six node network. The numbers represent the stored distance and feasible distance to node T . Initially, node E does not have a route to T and issues a RREQ. Nodes $\{B, C, D\}$ respond with RREPs. Let node C respond first. It happens to have a measured distance of 3 and a feasible distance of 2. These numbers may occur due to mobility and changing successors. When node C issues a RREP, it sets the measured distance to 3. When node E receives this RREP, it sets its measured distance and feasible distance both to 4. Node B then issues a RREP with start distance 4. Node E ignores that RREP when it receives it, because it does not provide a shorter distance than E 's current feasible distance. Node D issues a RREP with measured distance 1. When E receives that RREP, it updates both its feasible distance and measured distance to 2 and set its successor to D .

Node E issues a new RREQ with a feasible distance of 2 if links e_2 and e_3 fail at some future time. When node B receives the RREQ, it must forward it because it cannot create a RREP for it, given that its measured distance of 4 does not satisfy the requesting feasible distance. Node B must also set the T bit to indicate the destination must reset the path. Node C must forward the RREQ because its measured distance of 3 is not sufficient either. Node D , finally, could issue a RREP because its measured distance satisfies the requesting feasible distance. However, the reset bit is set

so C must unicast the RREQ to T , which would then issue a RREP with larger sequence number and a distance of 0. D would relay it to C . When node C receives the RREP, it changes its measured distance to 2 and keeps its feasible distance at 2, then relays the RREP with a distance of 2. When node B receives the RREP, it sets both its measured distance and feasible distance to 3, then relays the RREP with distance 3. Finally, node E receives the RREP and sets its measured distance to 4 and resets its feasible distance to 4.

3. ANALYSIS

We first show that LDR is loop free at every instant. Then we demonstrate that a source that requests a route to a given destination is successful within a finite time, provided that there is a physical path between source and destination and the network is stable for a sufficiently long period of time after an arbitrary sequence of topology changes. Because all computations in LDR are bounded by finite timers, showing that LDR is live is trivial and is omitted for brevity.

LDR uses a sequence number consisting of a destination-specific time stamp taken from a node's real-time clock and an unsigned monotonically increasing counter. When the counter reaches its maximum value, the node places a new time stamp in its sequence number and resets the counter to zero. We assume that a node's real-time clock does not reset on reboot or adjust for daylight savings. This scheme is adopted because it does not require synchronized clocks or the explicit reset of sequence numbers throughout the network. Furthermore, it avoids using AODV's reboot-hold procedure, which requires that a node stay off-line long enough for the network to forget any cached information about it.

THEOREM 3. *Solicitations and advertisements in LDR do not loop.*

Proof: For a given calculation (A, ID_A) , a node may be passive, engaged, or active. A node enters any calculation at most once. Therefore, the propagation graph of the calculation forms a tree. By using the cached information at engaged nodes, advertisements for the calculation follow paths only in the calculation tree.

If a node unicasts a solicitation, it is guaranteed to not flow in a loop, even if the underlying routing table contains loops. This is because nodes enter the engaged or active states at most once per computation, regardless of the unicast or broadcast nature of the solicitation. Thus, the T bit does not affect the loop-freedom of control packets. ■

LDR treats each solicitation independently of each other by identifying each such computation by the identifier of its origin and a sequence number assigned by the origin. This independence enables LDR to guarantee successful termination of simultaneous calculations for the same destination by multiple active nodes.

LEMMA 1. *If a node updates its routing table by Procedure 3 and relays advertisements by Procedure 4, then NDC ensures that the network maintains the ordering criteria during the creation of a successor path assuming no node along the path changes successor once on the path.*

Proof: By induction on the number of hops to the source of an advertisement. According to Procedure 4, the relaying of

an advertisement is no different than initiating an advertisement. The relay node places its distance and sequence number in the packet before transmitting it. Consider the successor path from node n_k to node n_1 to be $P = \{n_k, \dots, n_1\}$. Let the time $t_{n_i}^s$ be when node n_i chooses node n_{i-1} as its successor. Node n_1 initiates an advertisement \star_1 at time t_0 with $sn_{n_1}^{\star_1} = sn_{n_1}^{\star_1}(t_0)$ and $d_{n_1}^{\star_1} = 0$. For node n_2 to choose node n_1 as successor, one of the following cases must be satisfied according to NDC.

Case I: Node n_2 has no information about node n_1 . By Procedure 4, $sn_{n_2}^{n_2}(t_{n_2}^s) = sn_{n_1}^{\star_1}$ and $fd_{n_2}^{n_2}(t_{n_2}^s) > d_{n_1}^{\star_1}$. Because the sequence number is a non-decreasing function with time, $sn_{n_2}^{n_2}(t_{n_2}^s) \leq sn_{n_1}^{\star_1}(t_{n_2}^s)$. The sequence number \star_1 is fixed at time t_0 but the sequence number variable stored in $sn_{n_1}^{\star_1}$ may increase. Node n_1 will never have a positive distance to itself and the feasible distance at node $n_i \neq n_1$ is always positive, so the ordering criteria is true.

Case II: $sn_{n_1}^{\star_1} > sn_{n_2}^{n_2}(t_{n_2}^s)$. Because the sequence number is non-decreasing, $sn_{n_1}^{\star_1}(t_{n_2}^s) \geq sn_{n_1}^{\star_1}(t_0)$. If $sn_{n_1}^{\star_1}(t_{n_2}^s) > sn_{n_1}^{\star_1}(t_0)$, the ordering is maintained. Otherwise, $sn_{n_1}^{\star_1}(t_{n_2}^s) = sn_{n_1}^{\star_1}(t_0)$. In this case by Procedure 3, $fd_{n_1}^{n_2}(t_{n_2}^s) > d_{n_1}^{\star_1}$ so the ordering is maintained.

Case III: $sn_{n_1}^{\star_1} = sn_{n_2}^{n_2}(t_{n_2}^s) \wedge d_{n_1}^{\star_1} < fd_{n_1}^{n_2}(t_{n_2}^s)$. As per Case II, $sn_{n_1}^{\star_1}(t_{n_2}^s)$ could have increased, in which case the ordering criteria is true. If the sequence number did not change, then by Case I, $fd_{n_1}^{n_2}$ is always greater than $fd_{n_1}^{\star_1}$.

By the inductive assumption, nodes $\{n_i, \dots, n_1\}$ have a path that obeys the ordering criteria. We show that when node n_{i+1} chooses n_i , it maintains the ordering criteria. At time $t_{n_i}^s < t < t_{n_{i+1}}^s$, node n_i emits the advertisement. Because n_i does not change successor during $(t_{n_i}^s, t_{n_{i+1}}^s)$, $sn_{n_i}^{\star_i} = sn_{n_i}^{\star_i}(t_{n_i}^s)$ and $d_{n_i}^{\star_i} = d_{n_i}^{\star_i}(t_{n_i}^s)$. Following the three cases above, we only need to validate in Case I $fd_{n_1}^{n_{i+1}}(t_{n_{i+1}}^s) > fd_{n_1}^{\star_i}(t_{n_{i+1}}^s)$, as the other statements do not depend on the identity of node n_1 .

Case I (revised): By Procedure 4, $sn_{n_1}^{n_{i+1}}(t_{n_{i+1}}^s) = sn_{n_i}^{\star_i}$ and $fd_{n_1}^{n_{i+1}}(t_{n_{i+1}}^s) > d_{n_1}^{\star_i}$. Because the sequence number is a non-decreasing function, $sn_{n_1}^{n_{i+1}}(t_{n_{i+1}}^s) \leq sn_{n_1}^{n_i}(t_{n_{i+1}}^s)$. If the sequence number at node n_i increases (which does not happen by the premise of the lemma), then the ordering criteria is maintained regardless of $fd_{n_1}^{n_i}(t_{n_{i+1}}^s)$. If the sequence number remains the same, then $fd_{n_1}^{n_i}(t_{n_{i+1}}^s) \leq fd_{n_1}^{\star_i}(t) \leq fd_{n_1}^{\star_i}$. Thus, the ordering is maintained. ■

LEMMA 2. *Given an established path that obeys the ordering criteria, any change of successor along that path by NDC and Procedure 4 maintains the ordering.*

Proof: Let time $t_{n_i}^c$ be the time at which node n_i changes its successor off an established path. Taking the path in Lemma 1 $\{n_k, \dots, n_1\}$ that obeys the ordering criteria, we show that if some node n_i changes its successor to n_1 to some other node m_j along the path $\{m_j, \dots, m_1, n_1\}$, which is in order, that the ordering criteria is obeyed. That is, $sn_{n_1}^{n_i}(t_{n_i}^c) < sn_{n_1}^{m_j}(t_{n_i}^c)$ or $sn_{n_1}^{n_i}(t_{n_i}^c) = sn_{n_1}^{m_j}(t_{n_i}^c) \wedge fd_{n_1}^{n_i}(t_{n_i}^c) > fd_{n_1}^{m_j}(t_{n_i}^c)$. For the change to occur, node m_j must issue an advertisement \star_{m_j} at time $t_{n_i}^s < t < t_{n_i}^c$. As per the discussion in Lemma 1, the invariants at m_j at time $t_{n_i}^c$ can only be stronger than at time t , so if \star_{m_j} is feasible at n_i , Procedure 4 maintains the ordering criteria. What we must show is that Procedure 4 does not violate the ordering criteria for node n_{i+1} , which then by the assumption that P is ordered, is sufficient to show that the change does not

violate the ordering criteria anywhere along the path.

If \star_{m_j} is feasible at n_i at time $t_{n_i}^c$, then either $sn_{n_1}^{\star_{m_j}} > sn_{n_1}^{n_i}(t_{n_i}^c)$ or $sn_{n_1}^{\star_{m_j}} = sn_{n_1}^{n_i}(t_{n_i}^c) \wedge d_{n_1}^{\star_{m_j}} < fd_{n_1}^{n_i}(t_{n_i}^c)$. In the first case, Procedure 4 ensures that n_i 's sequence number increases, which satisfies the ordering criteria. In the second case, Procedure 4 decreases the feasible distance at node n_j , which also satisfies the ordering criteria. ■

THEOREM 4. *LDR is loop-free at every instant, as long as nodes update their routing tables according to NDC and Procedure 3, and relay messages by Procedures 2 and 4.*

Proof: Let node I be the node issuing the routing advertisement for destination D and let the advertisement take the path $P = \{n_1, \dots, n_j\}$, where $n_1 = I$ and $n_j = A$. I may be equivalent to D . Let the path from I to D be $Q = \{m_1, \dots, m_k\}$, where $m_1 = D$ and $m_k = I$ and it may be null if $n_1 = D$.

For a loop to form, A must be on the path Q . The path P is loop-free by Theorem 3. If $n_1 = D$, this is trivially true. If $n_1 \neq D$, we show that it is impossible for A to be on I 's successor graph using a proof by contradiction.

At time t_0 , assume that path Q exists and is loop free and at time t_1 , A chooses a successor path based on an advertisement from I . Let node A be some node m_i , $1 < i < k$. By the ordering criteria, $sn_D^A(t_0) > sn_D^I(t_0)$ or $sn_D^A(t_0) = sn_D^I(t_0)$ and $fd_D^A(t_0) < fd_D^I(t_0)$. At this time, node I sends an advertisement along the loop-free path P to A , thus $sn_D^* = sn_D^I(t_0)$ and $d_D^* = d_D^I(t_0)$.

At time t_1 , node A receives the advertisement. The sequence number for a destination is a non-decreasing function with time, so $sn_D^A(t_1) \geq sn_D^A(t_0)$.

In the case where $sn_D^A(t_0) > sn_D^I(t_0)$, A cannot accept the advertisement because $sn_D^* < sn_D^A(t_1)$.

In the case where $sn_D^A(t_0) = sn_D^I(t_0) = sn_D^A(t_1)$, we know that $d_D^* \geq d_D^I(t_0) \geq fd_D^I(t_0) > fd_D^A(t_0)$. Because the feasible distance is a non-increasing function with time for the same sequence number, $fd_D^A(t_1) \leq fd_D^A(t_0)$, so $d_D^* \geq fd_D^A(t_1)$. By NDC, node A cannot accept the advertisement.

In the case where $sn_D^A(t_0) = sn_D^I(t_0) < sn_D^A(t_1)$, we have $sn_D^* < sn_D^A(t_1)$ and A cannot accept the advertisement. ■

We now consider several lemmas that show LDR successfully terminates route calculations in several event contexts that span the route calculation space. Theorem 5 combines the lemmas to show that LDR terminates a route discovery successfully in an error-free and stable connected network.

In our analysis of the route calculation process, we have to make certain assumptions about the network. No routing protocol can converge if there are certain errors or if the network topology changes frequently enough. A route discovery process will also fail if no node exists that can satisfy the route, such as in a partitioned network. We will thus impose three conditions on our analysis arguments when we consider a node A initiating a route discovery: (1) There exists a node B , perhaps equal to D , such that node A and B are connected and B 's route to D is feasible for every node along the path from A to B ; (2) node A sends the solicitation with large enough time-to-live that it may reach a node capable of sending a feasible advertisement back to A ; and (3) relay nodes follow Procedure 2.

LEMMA 3. *If a single node A initiates a route discovery for destination D identified by (A, ID_A) in an error-free stable connected network, LDR guarantees that each solicitation*

for a route is answered with a feasible advertisement for the destination.

Proof: We consider the case of the first advertisement \star_A to reach A response to (A, ID_A) . If there are multiple advertisement, it does not affect the fact that A terminates successfully based on the first advertisement; A may improve its route.

Let node A send a solicitation for destination D , and let that solicitation traverse the path $P = \{n_1 \dots n_{k-1}\}$ before arriving at node n_k (possibly equal to D) which satisfies SDC. Node n_k issues an advertisement for D with terminus A . We show that the way in which nodes relay advertisements ensures that any solicited advertisement is usable by the relaying nodes. It is usable by A by virtue of satisfying SDC. By Procedure 2, we know any advertisement sent in response to (A, ID) will follow the reverse path of the solicitation, so it will follow the path P from n_k to A .

Let us first consider the case in which no node along the solicitation path P is affected by another route discovery event for D during the computation period. In this case, no node along P can satisfy A 's invariants and has an active route; otherwise, such a node would have responded to the solicitation instead of node n_k . Each node $n \in P$ matches one of three cases: (i) n has no information about D , (ii) n 's information is invalid, or (iii) the invariants of A are stronger than n 's invariants.

In case (i), node n may use any advertisement sent by n_k . In case (iii), the advertisement sent by n_k will satisfy the invariants at n because it satisfies A .

For case (ii), we show by induction that the advertisement issued by n_k in response to node A 's solicitation will satisfy all nodes along P if they followed Procedure 2. Let us first consider node n_1 , being the immediate neighbor of node A . If $sn_D^* > sn_D^\#$, then node n_1 has an invalid route to destination D and placed its own sequence number and feasible distance in the advertisement. Because the sequence number increased, node A may use any solicited advertisement. If $sn_D^* = sn_D^\#$, then node n relayed the solicitation with the minimum feasible distance of nodes A and n_1 , which ensures that nodes n_1 and A may use the solicited advertisement. Otherwise, node A had a higher sequence number than node n_1 , and any solicited advertisement will be usable by node n_1 .

If by the inductive assumption, all nodes $A \dots n_{i-1}$ may use the advertisement, then we show that the actions at node n_i do not affect predecessors, and n_i will be satisfied. Node n_i may only increase the sequence number or, keeping the sequence number the same, decrease the feasible distance; n_i cannot invalidate the usability of the advertisement for any predecessor node. Node n_i will relay invariants to be the stronger of its own or those already in the solicitation and ensure that it may use the advertisement. ■

LEMMA 4. *Considering the case of Lemma 3, let there be one or more other nodes $m_i \notin P$ that go active for D during the calculation (A, ID_A) . LDR guarantees that each solicitation for a route is answered with a feasible advertisement for the destination.*

Proof: By symmetry, if we show A 's route calculation successfully terminates, the other calculations for set m successfully terminate. We thus restrict ourselves to considering A 's calculation.

Because the relaying of a solicitation does not change the invariants stored at a node or the information cached for calculation (A, ID_A) , we only need to consider the interaction of advertisements. Let \star_j be the set of advertisements generated by the calculations (m_i, ID_{m_i}) . For these advertisements to affect (A, ID_A) they must intersect $\{A, P\}$.

If one or more of \star_j intersect A before \star_A , then A 's calculation (A, ID_A) terminates in success. By Lemma 3, the first \star_j to reach A is feasible at A .

Let $\star \in \star_j$ intersect some node $n \in P$ before \star_A . By Lemma 3, \star is feasible at n and n may relay it toward the appropriate active node. The sequence number and feasible distance at n could only have remained the same or been strengthened by \star .

If they remained the same, \star_A will be usable because there was no change. It could be that n changed from an inactive to an active route, but that does not affect the feasibility of \star_A at n .

If the invariants at n were strengthened by \star , then either (1) some node in $\{A, P\}$ has stronger invariants, or (2) n 's invariants became stronger than the strongest in $\{A, P\}$. In case (1), \star_A is stronger than n , so it will supersede the route at n and continue as per Lemma 3. In case (2) n will follow Procedure 4. It will discard \star_A and issue a new advertisement \uparrow_A with the stronger invariants. ■

LEMMA 5. *Considering the case of Lemma 3, let there be one or more other nodes $m_i \in P$ that go active for D during the calculation (A, ID_A) . LDR guarantees that each solicitation for a route is answered with a feasible advertisement for the destination.*

Proof: A node $n \in P$ going active for destination D does not change the cached information from the engagement (A, ID_A) nor does it change the invariants at n . It follows from Lemma 4 that any advertisements sent in response to calculation (n, ID_n) cannot interfere with calculation (A, ID_A) . ■

THEOREM 5. *If node A initiates a route discovery for destination D identified by (A, ID_A) in an error-free stable connected network, LDR guarantees that A receives a feasible advertisement for D .*

Proof: Considering the set of possible events, we must have the situation described in Lemma 3, Lemma 4, Lemma 5 or both Lemmas 4 and 5. The first three situations follow from the Lemmas. The four situation, being nodes both on and off P go active for D during the calculation (A, ID_A) , also follows from the Lemmas. No matter how many other nodes go active for D during A 's calculation, the other calculations cannot affect the cached information at nodes engaged in (A, ID_A) and by Lemma 4, any advertisements generated by those other calculations must have invariants at least as strong as necessary to complete (A, ID_A) if they are to interfere with it. ■

4. PERFORMANCE

We present results that show LDR out-performs other routing protocols over varying loads and mobility. Simulations are run in GlomoSim [1] for LDR, AODV, DSR, and OLSR. They follow the parameters in [13]. There are packet jitter problems in the OLSR code from INRIA for Linux [14], and we introduce a new FIFO jitter queue to OLSR. The

FIFO jitter queue adds a uniformly chosen inter-packet jitter between 0 and 15ms and maintains FIFO packet order. The modified code performs substantially better than the base OLSR. We use the DSR implementation from GlomoSim, which implements DSR Draft 3.

In our simulations, the performance of DSR is poor with mobility and load. In most delivery ratio figures, the DSR plot is substantially below other protocols. We reprogrammed our simulations in Qualnet 3.5.2, which implements DSR draft 7, but we observed similar results. Fig. 6 for Qualnet has the same mobility and traffic load patterns as Fig. 3 in GlomoSim. The performance of DSR is slightly better, but still shows the same downward trend with increasing mobility. The AODV implementation in Qualnet also showed poor performance in the 120 packet-per-second scenario. Although we only present one graph, DSR performance in other Qualnet runs mirrored GlomoSim over all tested scenarios.

We use the following optimizations to LDR:

Multiple RREPs: A node may relay multiple RREPs for the same (originator, rreqid) pair as long as only RREPs with stronger invariants cross over time.

Request as error: If node A receives a solicitation for destination D from neighbor B and A has an active route for D with next hop B , it is likely that B no longer has a valid route to D . If $fd_D^\# > d_B^A - lc_B^A$, then B should have answered the query if it had an active route.

Reduced distance: Because of mobility and link failures, it is often desirable in an ad hoc network to use non-optimal bounds. A node may place in a RREQ an *answering distance extension*, which is any distance no greater than the node's feasible distance. Nodes use the answering distance to test SDC. We use a factor of 0.8, truncated to the lowest integer no less than 1.

Minimum lifetime: A node should not respond to a RREQ if the lifetime remaining in its active route is less than a threshold. We use 1/3 the *ACTIVE_ROUTE_TIMEOUT*, or 1 second using default values. If a node receives such a RREQ, it relays the RREQ.

Optimal TTL: The initial TTL of a RREQ should be set according to the known distance and RREQ feasible distance. Here, let FD be the value, possibly lowered by the *reduced distance* optimization, given in the RREQ. The initial TTL should be $TTL = D - FD + 1 + LOCAL_ADD_TTL$.

There are two main sets of simulations, one on a 50-node network over a 1500m x 300m terrain, and one on a 100-node network over a 2200m x 600m terrain with 10-flow and 30-flow traffic loads using 512 byte packets at 4 packets per second per flow. The simulations use the 802.11 MAC layer with a 275m transmission range. The simulations run for 900 seconds. Nodes move between 1 m/s and 20 m/s. Flows have a mean length of 100 seconds, chosen from an exponential variate. We repeat each configuration (nodes, number of sources, routing protocol, and pause time) for 10 trials using different random number seeds. The LDR results reflect using the suggested optimizations.

We present six metrics. A "transmitted" packet count includes all hop-wise transmissions. An "initiated" packet count only includes the first transmission of a packet. The *delivery ratio* is the fraction of CBR data packets received at destinations. The *network load* is the total number of control packets (RREQ, RREP, RERR, Hello, TC, etc.) transmitted divided by total number of received data packets.

The *RREQ Load* is the total number of RREQs transmitted per received data packet. The *data latency* is the mean latency of data packets. The *RREP Init* is the number of RREPs initiated per RREQ initiated. The *RREP Recv* is the number of hop-wise usable RREPs received per RREQ initiated. A RREP may be usable at multiple nodes along its path to the RREQ origin. In the graphs, the vertical error bars represent the 95% confidence interval for all measurements. Table 1 summarizes the results by averaging over all pause times and both 50-node and 100-node scenarios for a given number of flows. The columns show the mean value and the 95% confidence interval range.

Figs. 2, 3, 4, 5 show the delivery ratio. In practically all cases, LDR has a higher delivery ratio than the other protocols. The exception is at high load and high mobility, where AODV sometimes performs better, and at high load and low mobility, where DSR sometimes performs better. In the low-load, 40pps, scenarios, LDR maintains a very high delivery ratio at all mobility speeds. The minimum delivery ratio is 98.5% for the 200s pause time in Fig. 4 (100-node, 10-flow, 40pps). The average LDR delivery ratio for all pause times and all 10-flow scenarios is 0.992 ± 0.002 . The next-best protocol in terms of delivery ratio is AODV, but in many cases it is almost identical to OLSR. For 30-flows, LDR, AODV and OLSR are statistically identical on the average over all mobility pause times.

Fig. 7 shows the mean destination sequence number for LDR and AODV for low load and high load. LDR with 10-flows has a maximum mean sequence number of 0.8 and with 30-flows has a maximum of 15.8. AODV, on the other hand, has a 10-flow maximum of 104 and a 30-flow maximum of 108.

In terms of delay from Table 1, we see that OLSR and LDR have statistically identical latencies (overlapping confidence intervals), although at 10-flows they barely overlap. DSR and AODV have significantly higher latency. In terms of load from Table 1, LDR and AODV have statistically identical network load for 10-flows. LDR, AODV, DSR, and OLSR are equivalent for 30-flows. Looking at the RREQ load, LDR transmits about 1/3 fewer broadcast RREQs. LDR initiates about 39% fewer RREPs per RREQ packet for 10 flows and 8% fewer for 30 flows than AODV. Yet, LDR has a higher number of received RREPs per RREQ packet.

5. CONCLUSION

We have presented LDR, a novel on-demand loop-free routing protocol using a new distance label invariant based on DUAL and a messaging structure based on AODV. LDR removes the requirement of AODV for nodes to independently increase other node's sequence numbers, placing firm control of a sequence number with the owning node. LDR avoids loops by introducing a second routing invariant, the feasible distance. If the start distance of an advertisement is less than a node's feasible distance, there can be no loop, even if the measured distance increases. Destination sequence numbers play the role of resets for feasible distances and allow nodes to increase their feasible distance along paths. Through simulation, we showed that LDR exhibits better packet delivery ratios than AODV, DSR, and OLSR, except at 3 out of 32 data points. At times, LDR's delivery ratio is significantly higher than other protocols and is never significantly worse. Our data also shows that LDR's

network load is less than or comparable to the loads of other protocols. In terms of mean latency of data packets, LDR is second to OLSR at low load and comparable to OLSR at high load.

6. REFERENCES

- [1] L. Bajaj et al. GloMoSim: A scalable network simulation environment. Technical Report 990027, UCLA Computer Science Department, 1999.
- [2] T. Clausen et al. Optimized link state routing protocol. IETF Internet draft, draft-ietf-manet-olsr-06.txt, Sep 2001.
- [3] E. W. Dijkstra and C. S. Scholten. Termination detection for diffusing computations. *Information Processing Letters*, 11(1):1-4, Aug. 1980.
- [4] E. M. Gafni and D. P. Bertsekas. Distributed algorithms for generating loop-free routes in networks with frequently changing topology. *IEEE Trans. Comm.*, COM-29(1):11-18, Jan. 1981.
- [5] J. Garcia-Luna-Aceves and S. Murthy. A path finding algorithm for loop-free routing. *IEEE/ACM Trans. Networking*, 5(1):148-160, Feb. 1997.
- [6] J. J. Garcia-Luna-Aceves and M. Spohn. Source-tree routing in wireless networks. In Proc. *IEEE ICNP'99*, pp. 273-82, Oct. 1999.
- [7] J. J. Garcia-Luna-Aceves. Loop-free routing using diffusing computations. *IEEE/ACM Transactions on Networking*, 1(1):130-41, Feb. 1993.
- [8] D. Johnson et al. The dynamic source routing protocol for mobile ad hoc networks (DSR). IETF Internet draft, draft-ietf-manet-dsr-07.txt, Feb 2002.
- [9] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.
- [10] R. Ogier et al. Topology broadcast based on reverse-path forwarding (TBRPF). IETF Internet draft, draft-ietf-manet-tbrpf-05.txt, Mar 2001.
- [11] V. D. Park and M. S. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *IEEE INFOCOM*, pp. 1405-13 vol.3, Apr. 1997.
- [12] C. Perkins, E. Belding-Royer, and S. Das. Ad hoc on demand distance vector (AODV) routing. IETF Internet draft, draft-ietf-manet-aodv-10.txt, Mar 2002.
- [13] C. Perkins et al. Performance comparison of two on-demand routing protocols for ad hoc networks. *IEEE Personal Communications*, 8(1):16 - 28, Feb 2001.
- [14] A. Plakoo and A. Laouiti. INRIA OLSR draft 3 linux source code. <http://menetou.inria.fr/olsr/#code>, ported by Marco Spohn, 2001.
- [15] J. Raju and J. J. Garcia-Luna-Aceves. A new approach to on-demand loop-free multipath routing. In *IEEE IC3N'99*, pp. 522-7, Oct. 1999.
- [16] M. Spohn and J. J. Garcia-Luna-Aceves. Neighborhood aware source routing. In *ACM MOBIHOC 2001*, pp. 11-21, Oct. 2001.
- [17] S. Vutukury and J. Garcia-Luna-Aceves. A simple approximation to minimum-delay routing. In Proc. *AMC SIGCOMM 1999*, pp. 227 - 238, Sept. 1999.

Table 1: Performance average over all pause times

flows	protocol	delivery ratio	Net load	RREQ load	latency (sec)	RREP Init	RREP Recv
10	LDR	0.992 ± 0.003	0.464 ± 0.132	0.410 ± 0.121	0.066 ± 0.034	3.379 ± 0.129	1.335 ± 0.020
10	AODV	0.962 ± 0.017	0.777 ± 0.203	0.609 ± 0.136	0.731 ± 0.538	5.641 ± 1.356	1.303 ± 0.006
10	DSR	0.736 ± 0.087	8.337 ± 3.719	0.043 ± 0.015	2.111 ± 0.831		
10	OLSR	0.904 ± 0.026	5.773 ± 2.040		0.029 ± 0.004		
30	LDR	0.857 ± 0.044	2.063 ± 0.759	1.532 ± 0.556	0.656 ± 0.113	7.183 ± 0.998	1.642 ± 0.105
30	AODV	0.822 ± 0.028	2.353 ± 0.623	1.746 ± 0.455	1.030 ± 0.105	7.775 ± 0.774	1.533 ± 0.158
30	DSR	0.603 ± 0.094	3.270 ± 1.625	0.033 ± 0.011	3.837 ± 1.567		
30	OLSR	0.799 ± 0.044	2.680 ± 1.009		0.489 ± 0.146		




Figure 2: Delivery 50-nodes, 10-flow, 40pps




Figure 3: Delivery 50-nodes, 30-flow, 120pps




Figure 4: Delivery 100-nodes, 10-flow, 40pps




Figure 5: Delivery 100-nodes, 30-flow, 120pps




Figure 6: Qualnet delivery 50-node, 30-flow, 120pps




Figure 7: Average node sequence number