

Analysis of the Zone Routing Protocol

Jan Schaumann

December 8, 2002

Abstract

This document describes and analyzes the Zone Routing Protocol (ZRP), a hybrid mobile ad-hoc protocol which divides the network into overlapping routing zones, allowing for the use of independent protocols within and between the zones.

Contents

1	Overview	4
2	MANET in general	4
2.1	Pro-active vs. reactive	5
3	Introduction to ZRP	6
3.1	Intrazone Routing Protocol (IARP)	8
3.2	Interzone Routing Protocol (IERP)	9
3.3	Bordercast Resolution Protocol (BRP)	10
4	Examples	12
4.1	BRP	12
4.2	Stationary nodes in a dense network	13
4.3	Mobile nodes without a stationary fixpoint	14
5	Performance	15
5.1	Performance gain on the example of Bordercasting	16
6	Summary	17

List of Figures

1	Routing Zone of node A with $\rho = 2$	6
2	ZRP components	8
3	QD1 and QD2.	11
4	Selective Broadcasting.	11
5	A BRP Example.	12
6	A stationary network with $\rho = 2$	13
7	Node H's zone.	14
8	Overlapping zones.	14
9	Mobile nodes (1).	15
10	Mobile nodes (2).	15
11	Query Flooding	16
12	Selective Bordercasting	17
13	IERP Traffic per Route Discovery	18

1 Overview

This document describes and analyzes the Zone Routing Protocol or ZRP for Mobile Ad-Hoc Networks (MANETs). The basics of MANET and the implications on routing in particular are briefly covered in Section 2 to provide an introduction to the problems resulting from a rapidly changing topology without a fixed router.

As we will see in Section 3, ZRP, in contrast to other MANET routing protocols, utilizes a *hybrid* pro-active/reactive approach to maintain valid routing tables without too much overhead. Furthermore, ZRP does not provide a single protocol, but rather outlines a routing *framework* suitable for inclusion and extension of other existing protocols.

In describing the protocol, we will consider some specific examples in order to visualize how ZRP's characteristics influence its performance (Section 4).

Section 5 analyzes the performance of the ZRP as well as discusses various scenarios, and in Section 6 the information provided in this document is briefly summarized.

2 MANET in general

A *Mobile Ad-Hoc Network* (MANET) is a decentralized network of autonomous mobile nodes able to communicate with

each other over wireless links. Due to the mobility of the nodes, the topology of the network may rapidly be changing, making it impossible to use conventional routing tables maintained at fixed points (routers). Instead, each node is required to determine the best route to a given destination node by itself.

Given their dynamic nature, route discovery in a MANET differs significantly from the more or less static routes in wired networks: Not all nodes in a MANET necessarily have the same capabilities. Two nodes, even if they are direct neighbors, may differ with respect to signal strength, available power, reliability etc.

These differences require much more complicated and particularly more *active* distributed algorithms in order to maintain an accurate picture of the networks topology, while at the same time providing scalability for potentially large (and ever-growing) networks. At the same time, route discovery must not use up the majority of the often limited bandwidth available to today's mobile devices.

Furthermore, it is important to point out an important difference to conventional routing approaches: In wired networks, each link is bi-directional. If a node A can send packets to a node B, we know that node B can send packets back to node A, and a reverse path can be entered. This is not necessarily the case in a wireless network,

where the physical location and the individual power resources have great influence upon a nodes transmission capacity and signal strength.

MANET routing protocols are IP based and may use unicast, multicast or hybrid approaches and should allow for interaction with standard wired IP services rather than being regarded as a completely separate entity.

A detailed yet not overly complex overview of the various aspects of Mobile Ad-Hoc Networking is given in [7].

2.1 Pro-active vs. reactive

The IETF MANET Working Group has researched and developed a number of protocols for mobile ad-hoc networks, which have been described in [8], [9], [10], [11] and [12]. These protocols can generally be categorized into two groups: *pro*-active and *re*-active protocols.

Pro-active protocols follow an approach similar to the one used in wired routing protocols. By continuously evaluating the known and attempting to discover new routes, they try to maintain the most up-to-date map of the network. This allows them to efficiently forward packets, as the route is known at the time when the packet arrives at the node.

Pro-active or table-driven protocols, in order to maintain the constantly changing network graph due to new,

moving or failing nodes, require continuous updates, which may consume large amounts of bandwidth – clearly a disadvantage in the wireless world, where bandwidth is often sparse. Even worse so, much of the accumulated routing information is never used, since routes may exist only for very limited periods of time.

The family of Distance-Vector protocols, including Destination-Sequenced Distance-Vector Routing ([13]), fall into the category of pro-active protocols.

In contrast, reactive protocols determine the proper route only when required, that is, when a packet needs to be forwarded. In this instance, the node floods the network with a route-request and builds the route on demand from the responses it receives. This technique does not require constant broadcasts and discovery, but on the other hand causes delays since the routes are not already available. Additionally, the flooding of the network may lead to additional control traffic, again putting strain on the limited bandwidth.

These reactive (or on-demand) protocols include Dynamic Source Routing (DSR) [9] and Ad-hoc On demand Distance Vector Routing (AODV) [8], as well as the classical flooding algorithms. [1]

3 Introduction to ZRP

As explained above, both a purely proactive or purely reactive approach to implement a routing protocol for a MANET have their disadvantages. The *Zone Routing Protocol*, or ZRP, as described in this document combines the advantages of both into a *hybrid* scheme, taking advantage of pro-active discovery within a node's local neighborhood, and using a reactive protocol for communication between these neighborhoods.

In a MANET, it can safely be assumed that the most communication takes place between nodes close to each other. Changes in the topology are most important in the vicinity of a node – the addition or the removal of a node on the other side of the network has only limited impact on the local neighborhoods.

As mentioned earlier, the ZRP is not so much a distinct *protocol* as it provides a *framework* for other protocols. The separation of a node's local neighborhood from the global topology of the entire network allows for applying different approaches – and thus taking advantage of each technique's features for a given situation. These local neighborhoods are called *zones* (hence the name); each node may be within multiple overlapping zones, and each zone may be of a different size. The “size” of a zone is not determined by geographical mea-

surement, as one might expect, but is given by a radius of length ρ , where ρ is the number of hops to the perimeter of the zone.

By dividing the network into overlapping, variable-size zones, ZRP avoids a hierarchical map of the network and the overhead involved in maintaining this map. Instead, the network may be regarded as flat, and route optimization is possible if overlapping zones are detected.

While the idea of zones often seems to imply similarities with cellular phone services, it is important to point out that each node has its own zone, and does not rely on fixed nodes (which would be impossible in MANETs). Figure 1 shows an example routing zone with $\rho = 2$.




Figure 1: Routing Zone of node A with $\rho = 2$.

Note that in this example node A has multiple routes to node F, including one that has a hopcount of $c > \rho$. Since it also has a route with $c \leq \rho$, F still belongs to A's zone. Node G is out of A's zone,

The nodes on the perimeter of the zone (i.e. with a hopcount $hc = \rho$) are referred to as peripheral nodes (marked gray), nodes with $hc < \rho$ are interior nodes.

Obviously a node needs to first know about its neighbors before it can construct a routing zone and determine its peripheral nodes. In order to learn about its direct neighbors, a node may use the media access control (MAC) protocols directly. Alternatively, it may require a *Neighbor Discovery Protocol* (NDP). Again, we see that ZRP, as a framework, does not strictly specify the protocol used but allows for local independent implementations.

Such a Neighbor Discovery Protocol typically relies on the transmission of "hello" beacons by each node. If a node receives a response to such a message, it may note that it has a direct point-to-point connection with this neighbor. The NDP is free to select nodes on various criteria, such as signal strength or frequency/delay of beacons etc. Once the local routing information has been collected, the node periodically broadcasts discovery messages in order to keep its

map of neighbors up to date. In doing so, it is assumed that these "link-layer (neighbor) unicasts are delivered reliably and in-sequence." [1]

If the MAC layer of the nodes does not allow for such a NDP, the Intra-zone Routing Protocol must provide the possibility of direct neighbor discovery. This protocol is responsible for determining the routes to the peripheral nodes and is commonly a proactive protocol. The Intrazone Routing Protocol, or IARP, is described in more detail in Section 3.1.

Communication between the different zones is guarded by the Interzone Routing Protocol, or IERP, and provides routing capabilities among peripheral nodes only. That is, if a node encounters a packet with a destination outside its own zone – i.e. it does not have a valid route for this packet – it forwards it to its peripheral nodes, which maintain routing information for the neighboring zones, so that they can make a decision of where to forward the packet to. Through the use of a bordercast algorithm rather than flooding all peripheral nodes, these queries become more efficient. The Interzone Routing Protocol and the Bordercast Resolution Protocol are presented in Sections 3.2 and 3.3.

As we can see, the Zone Routing Protocol consists of several components, which only together provide the full routing benefit to ZRP. Each

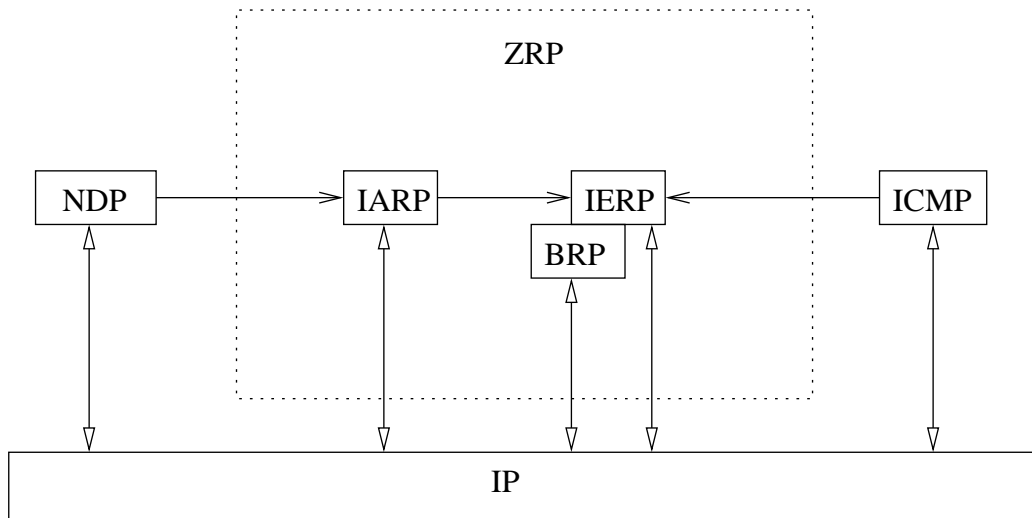


Figure 2: ZRP components

component works independently of the other and they may use different technologies in order to maximize efficiency in their particular area. For example, a reactive protocol such as AODV might be used as the IARP, while the IERP is most commonly a pro-active protocol such as OLSR [14].

Figure 2, as adapted from [1], illustrates the different protocols and their interactions.

Even though the hybrid nature of the ZRP seems to indicate that it is a hierarchical protocol, it is important to point out that the ZRP is in fact a *flat* protocol. In a hierarchical network architecture, two different protocols are maintained for communication among (a) each individ-

ual cluster's nodes and (b) the different clusters. The main difference here is that in the ZRP there is a one-to-one correspondence between nodes and routing zones, causing overlapping zones maintained by each individual nodes (see [1] for details).

3.1 Intrazone Routing Protocol (IARP)

Since ZRP assumes that local neighbor discovery is implemented on the link-layer and is provided by the NDP, the first protocol to be part of ZRP is the *Intrazone Routing Protocol*, or IARP. This protocol is used by a node to communicate with the interior nodes of its zone and as such is limited by the zones radius ρ (the number of hops

from the node to its peripheral nodes).

Since the local neighborhood of a node may rapidly be changing, and since changes in the local topology are likely to have a bigger impact on a node's routing behavior than a change on the other end of the network, the IARP is a pro-active, table-driven protocol.

The node continuously needs to update the routing information in order to determine the peripheral nodes as well as maintain a map of which nodes can be reached locally. The IARP allows for local route optimization through the removal of redundant routes and the shortening of routes if a route with fewer hops has been detected, as well as bypassing link-failures through multiple (local) hops, thus leveraging global propagation.

As mentioned earlier, it is possible that a node A can broadcast messages to a node B, but that node B, due to limitations in its signal-strength (caused by interference, for example) or low transmission power, can not reach node A. Therefore, it is important for the IARP to provide support for unidirectional links among the local nodes.

Due to its pro-active nature, local route discovery is very efficient and routes to local destinations are immediately available. In order to not overutilize the available bandwidth resources, the IARP – as the name sug-

gests – is restricted to routing within the zone, which is why it is referred to as a “limited scope pro-active routing protocol”[3].

Global route discovery, communication with nodes in a different zone, is done by guiding the route queries to the peripheral nodes instead of flooding all local nodes. In order to adopt a traditional pro-active link state protocol for use as the IARP in the ZRP, the scope of the protocol needs to be limited to the size of the zone ρ . This may be implemented by adding a Time To Live (TTL) to the route discovery requests, initialized to $\rho - 1$, and decremented by each node until it reaches 0 (when it is discarded).

In Figure 1, uninterrupted lines indicate the areas where the IARP is used to provide routing between the nodes. The Intrazone Routing Protocol including an example implementation (a Timer Based Link State IARP) is described in more detail in [3].

3.2 Interzone Routing Protocol (IERP)

As the global reactive routing component of the ZRP, the Interzone Routing Protocol, or IERP, takes advantage of the known local topology of a node's zone and, using a reactive approach enables communication with nodes in other zones.

Route queries within the IERP are issued on demand, that is only when a request for a route is made. The delay caused by the route discovery (in contrast to IARP, where the route is immediately available) is minimized through the use of bordercasting, an approach in which the node does not submit the query to all local nodes, but only to its peripheral nodes. Furthermore, a node does *not* send a query back to the nodes the request came from, even if they are peripheral nodes (as explained in Section 3.3, Figure 5).

In order to convert an existing reactive routing protocol for use as the IERP in the ZRP, it is necessary to disable pro-active updates for local routes, since this functionality is provided by the IARP. Furthermore, the IERP needs to be able to take advantage of the local routing information provided by the IARP, as well as change the way route discovery is handled: Instead of flooding a route request to all nodes, it should instead use the Bordercast Resolution Protocol (BRP) to only initiate route requests with peripheral nodes.

In Figure 1, dotted lines indicate the areas where the IERP is used to provide routing between the zones. The Interzone Routing Protocol including an example implementation (Reactive Source Routing) is described in more detail in [4].

3.3 Bordercast Resolution Protocol (BRP)

The Bordercast Resolution Protocol, or BRP, is used in the ZRP to direct the route requests initiated by the global reactive IERP to the peripheral nodes, thus removing redundant queries and maximizing efficiency. In doing so, it utilizes the map provided by the local pro-active IARP to construct a bordercast tree. Unlike IARP and IERP, it is not so much a routing protocol, as it is packet delivery service.

The BRP keeps track of which nodes a query has been delivered to, so that it can prune the bordercast tree of nodes that have already received (and relayed) the query. When a node receives a query packet for a node that does not lie within its local routing zone, it constructs a bordercast tree so that it can forward the packet to its neighbors. These nodes, upon receiving the packet, reconstruct the bordercast tree so that they can determine whether or not it belongs to the tree of the sending node. If it does not, it continues to process the request and determines if the destination lies within its routing zone and taking the appropriate action, upon which the nodes within this zone are marked as covered.

In order to detect when a routing zone they belong to has been queried, two levels of *Query Detection* are provided by BPR. As they relay the queries

to the peripheral nodes, the nodes detect the query and notes which zones have been covered. This is referred to as the first level of Query Detection, or QD1.

Secondly, in networks that use a single broadcast channel, a node can determine this information by listening to the traffic broadcast among other nodes. This approach is referred to as QD2. Figure 3 shows node A bordercasting a query to the peripheral nodes D and F. Nodes B and C, as they relay the query, note that node A's zone has been queried (QD1). In single-channel networks, node E can listen to the traffic and come to the same conclusion using QD2.




Figure 3: QD1 and QD2.

Simply detecting that a given node has already been covered, however, is not enough. The protocol needs to drop packets that would be sent to al-

ready covered nodes. This is done using *Early Termination* or ET, which obviously relies on Query Detection, as well as *Loopback Termination* or TL (in which routes that loop back into the querying nodes zone are eliminated, since these nodes can be reached locally using IARP).

In order to further eliminate unnecessary broadcasting, the BRP may implement *Selective Bordercasting*. In this approach, a node needs to know network topology information for an extended zone of size $2\rho - 1$.

Given this knowledge, a node can further eliminate peripheral nodes from its list of bordercast recipients, if the outer peripheral nodes overlap. Figure 4 shows an example of how node A is able to remove node C from its bordercasting spanning tree, since nodes G and H can be reached through nodes B and D respectively.




Figure 4: Selective Broadcasting.

In the context of ZRP, the BRP can be seen as the glue which ties together the IARP and the IERP in order to take full advantage of the proactive and reactive components where they are best used.

The Bordercast Resolution Protocol, including the implementation is described in more detail in [5].

4 Examples

In order to better understand how the different components of the ZRP work together and how routing is done using this approach, we will consider a few examples. Several scenarios are possible, among them stationary nodes in a dense network, mobile nodes moving in different directions with and without stable and/or stationary fix-points and stationary or moving networks with instable or frequently changing nodes. In these examples, a node is considered *stationary*, if it does not move relative to the other nodes (even though it may be moving into the same direction). A *stable* node is one that broadcasts with a constant signal and does not undergo power fluctuations.

4.1 BRP

Before we go into examples of how routes between nodes from different zones are determined, let us consider

an example of how BRP performs Query Detection.




Figure 5: A BRP Example.

In Figure 5, node A needs to send a packet to node H. Since node H is not within its zone, it constructs a bordercast tree spanning its peripheral nodes C, E and F, and sends the query to its local neighbors B and D. Both B and D note that H is not within their zone they in turn construct a bordercast tree. Note that they only include peripheral nodes of theirs that have not been covered previously.

That is, node B will not include node D, even though it is a peripheral node for B's zone. Similarly, node D's tree will not include node B. Node D then forwards the request to node

F, which determines that node H lies within its routing zone, and replies with the correct route.

4.2 Stationary nodes in a dense network

Let us extend the network from Figure 1 to span a few more nodes. An example situation of such a stationary network, where the number of nodes and their position does not change frequently might be a number of people attending a conference, communicating with each other. This example lends itself to show the basics of ZRP, so we will investigate it in more detail.

Figure 6 shows the network graph; the radius of each node's zone is still $\rho = 2$. Again, as before, the peripheral nodes for node A are marked gray. In this example, node A has to send a packet to node U.

Node A uses the IARP to determine if node U is within its zone. Since IARP is pro-active, the information that U is in this routing zone is readily available, and node A initiates a route request using IERP. As explained in Section 3.3, IERP now utilizes BRP to bordercast the request: It is not flooded to each of the nodes in A's zone, but only to B, D, E, F, H and J, the peripheral nodes, which in turn search their routing tables for the destination.




Figure 6: A stationary network with $\rho = 2$.

Node H does not find U in its routing table and thus bordercasts the request to its peripheral nodes. Figure 7 shows node H's routing zone.

Through the use of BPR, the bordercast tree of node H does not contain nodes F or A: these branches have been pruned, since these nodes have already been covered. Node H broadcasts the route query to nodes M and N. Some of the peripheral nodes of node N's zone are also peripheral nodes of node U's zone – an example of how zones can (and frequently do) over-




Figure 7: Node H's zone.




Figure 8: Overlapping zones.

lap. In Figure 8, the peripheral nodes for both node U and node N are shaded.

Now node N bordercasts the query to nodes R and S (nodes M and H are already marked as covered), which in turn reply with the correct route, as they both know U to be within their local zone.

4.3 Mobile nodes without a stationary fixpoint

In this example, we will see how ZRP deals with link-failures and link-optimization, since moving nodes will constantly have to update their network map, as nodes that used to be within

their zone move out of transmission range. For simplicity, let us consider a much smaller initial network. In Figure 9, the arrows indicate the direction in which the nodes are traveling. As we observe node A, we mark its peripheral nodes gray.

As nodes D and B move further away from each other, they lose their connection, just as nodes C and E. However, node E moves closer and is able to establish a point-to-point connection with nodes A and D. Similarly, node F moves closer and is able to establish connections with nodes E and A. These changes are reflected in




Figure 9: Mobile nodes (1).

Figure 10.




Figure 10: Mobile nodes (2).

In this example, it is interesting to point out that node D remains a peripheral node for A's zone, even though the local route from A to D changes. This change shows the need for a proactive IERP: the changes in the networks topology like will have taken

place within a short period of time. Nodes connected to E only need to add an additional peripheral node (F), as E maintains its connection to node D. However, were a node connected to E request a route to C, the IERP would need to initiate a new route request, as C no longer is within the querying node's zone and the route needs to go through D.

5 Performance

In order to maximize performance of the ZRP, we need to minimize the amount of control traffic that is sent. Thus, we wish to maintain an overview of the networks topology that is as accurate as possible (at any given time – thus minimizing delays caused by route discovery requests), while at the same time requiring sending as little packages as possible.

Given the hybrid nature of the ZRP, this goal can be reduced to finding the correct – i.e. optimal – size of the routing zone radius ρ for the given network – which may vary from case to case, depending on the circumstances.

For example, in a stationary network as considered in Section 4.2, it would be possible to increase ρ to a larger number, without too much of a penalty: in this situation, the position or the number of available nodes changes infrequently, so that, given a larger routing zone radius, the nodes

could take advantage of the comparably static and immediately available, since pro-actively maintained, routes.

The example in Section 4.3 on the other hand would not benefit from large zones: the cost of maintaining the ever-changing local routes is too high, particularly since most of the routes are so short-lived that they are never used. Instead, a zone radius of $\rho \leq 3$ would be beneficial, to ensure that the zones overlap enough to allow for route-redundancy. Note, however, that reducing $\rho \rightarrow 0$ effectively turns the ZRP into a completely reactive protocol, obsoleting the advantages gained from its hybrid nature: all routing is done on-demand (using IERP), as no node is able to contact another node using IARP.

The results of [16] showed that the IARP traffic grows with the number of nodes in a given zone, while increased mobility of a the nodes increases IERP traffic: as nodes move, the routes between zones break and need to be “re-discovered”. Increasing the number N of nodes in the global network has only limited effect on the amount of pro-active traffic, since pro-active IARP updates are local to a zone. In general, it can be stated that “larger zones provide more efficient queries, which compensates for the IARP maintenance cost”([2]).

A more detailed analysis of ZRPs

general performance can be found on [16] and [17].

5.1 Performance gain on the example of Bordercasting

In this section, we will show how BRP can minimize the number of broadcasts significantly. First let us consider a network as shown in Figure 11.

The first 8 nodes connected to the node in the middle should be seen as that node’s peripheral nodes (that is, interior nodes are not shown in this example), the 8 nodes on the outside as the peripheral nodes of the extended zone. As we can see, the total number of queries if no BRP would be used would sum up to be 40!




Figure 11: Query Flooding

Now we introduce selective bor-

dercasting. A large number of broadcasts can be avoided, since many of the zones overlap and peripheral nodes connect to the same nodes in the extended zone. As Figure 12 shows, the number has been brought down to 16, a decrease of over 50%!




Figure 12: Selective Bordercasting

The tremendous advantage of using Selective Broadcasting in the IERP is visualized in Figure 13, as adopted from [16], [17] and others.

6 Summary

The Zone Routing Protocol (ZRP) provides a hybrid routing framework, in which each node maintains local routes within its zone in a pro-active manner, while interzone communication is performed in a reactive manner. In order to improve performance of the reactive IERP and avoid having to rely

on flooding all neighboring nodes, thus risking exhausting the available bandwidth, it makes use of a bordercasting protocol.

It is important to note that all of the components of the ZRP provide the necessary flexibility optimal for nodes in a MANET:

- each node may choose a different zone radius ρ , according to its signal strength, transmission power, velocity etc.
- ZRP provides support for unidirectional links, which frequently occur in MANETs
- local neighbors are detected using NDP, which can be implemented on the link-layer. Even though this requires periodic messages to maintain an accurate overview of the interior nodes, there is not much of an overhead since this is performed within the (comparably) small local zone
- since only link-layer unicasts are presumed to be reliable and in sequence, these properties are no requirements for other traffic
- the protocols that make up the ZRP are implementation specific. That is, one MANET with certain requirements may adjust AODV for use as the IERP and OLSR

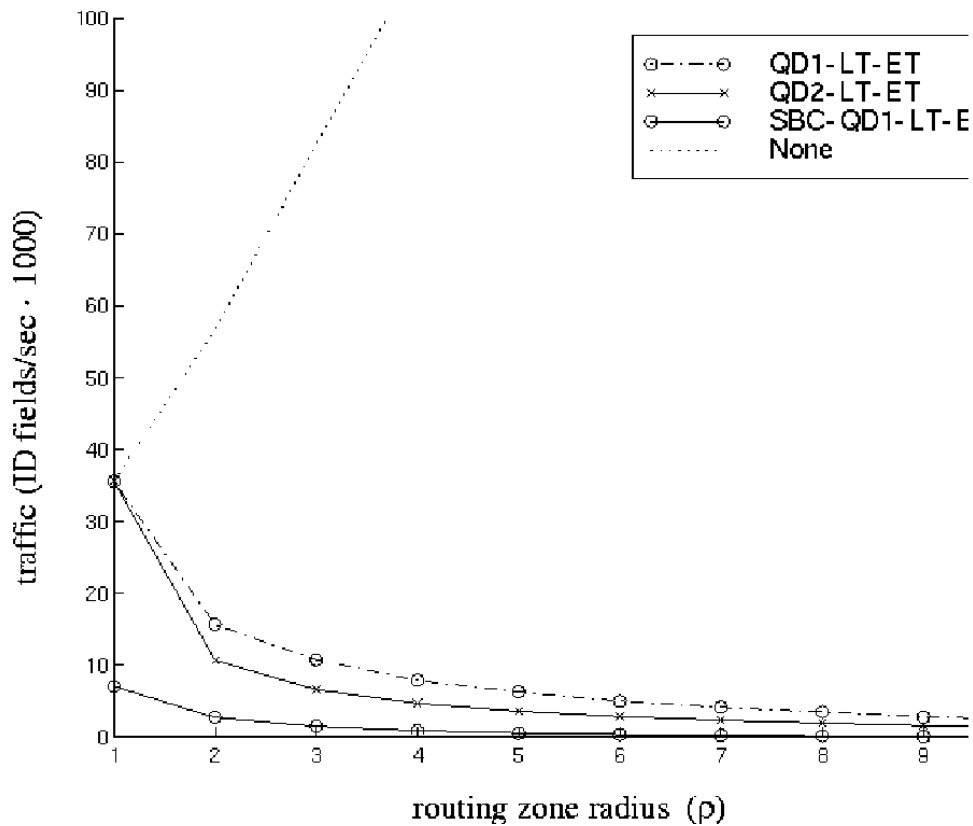


Figure 13: IERP Traffic per Route Discovery

as the IARP, while in another situation it might be favorable to use TBRPF as the IARP and DSR as the IERP.

- by basing all components on the IP, ZRP can provide all the services customers are used to from wired networks. Another advantage is the possibility of implementing security aspects on the IP layer (for example through

the use of IPsec), thus removing costly overhead from the wireless protocol

- vast performance gain through the use of link-optimization, selective bordercasting, Query Detection (1 and 2), Early Termination and Loopback Termination (see Figure 13)
- the ZRP can determine the appropriate zone radius dynami-

cally on-demand (as shown in [17]; propagation of the new routing zone is conceivable but would need to be implemented)

In contrast to other MANET routing protocols, ZRP, as mentioned above, provides a hybrid approach, a framework of protocols. Thus, it does not directly conflict or compete with any of the protocols described in [8], [9], [10], [11] and [12], but is able to take advantage of each of those protocols strengths, depending on the situation, requirements and implementation.

Also, ZRP is more suitable than other protocols for large networks spanning diverse mobility patterns by providing the benefits of both reactive and pro-active routing in a flat network that takes advantage of a near-hierarchical approach.

The Zone Routing Protocol is formally described by the Mobile Ad-Hoc Networks (MANET) Working Group in an IETF Internet Draft to expire in January 2003 ([1]). Its components IARP, IERP and BRP are described in the IETF Internet Drafts [3], [4] and [5] respectively.

References

- [1] *The Zone Routing Protocol (ZRP) for Ad Hoc Networks* Haas, Z.J., Pearlman, M.R., Samar, P., IETF Internet Draft, draft-ietf-manet-zone-zrp-04.txt, July 2002.
- [2] *Zone Routing Protocol* Nicklas Beijar, <http://www.tct.hut.fi/opetus/s38030/k02/Papers/08-Nicklas.pdf>
- [3] *Intrazone Routing Protocol (IARP)* Haas, Z.J., Pearlman, M.R. and Samar, P., IETF Internet Draft, draft-ietf-manet-iarp-02.txt, July 2002.
- [4] *Interzone Routing Protocol (IERP)* Haas, Z.J., Pearlman, M.R. and Samar, P., IETF Internet Draft, draft-ietf-manet-ierp-02.txt, July 2002.
- [5] *Bordercasting Resolution Protocol (BRP)* Haas, Z.J., Pearlman, M.R. and Samar, P., IETF Internet Draft, draft-ietf-manet-brp-02.txt, July 2002.
- [6] *Mobile Ad-Hoc Networks (MANET) Charter* IETF, <http://www.ietf.org/html.charters/manet-charter.html>
- [7] *Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations* Corson, S., Macker, J., Network Working Group, January 1999, RFC2501
- [8] *Ad Hoc On Demand Distance Vector (AODV) Routing* Perkins, C.E., Belding-Royer, E.M., Das, S.R., IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-12.txt>
- [9] *The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks* Johnson, D.B., Maltz, D.A., Hu, Y.-C., Jetcheva, J.G., IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-07.txt>
- [10] *On-Demand Multicast Routing Protocol (ODMRP) for Ad-Hoc Networks* Yi, Y., Lee, S.-J., Su, W., Gerla, M., IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-odmrp-04.txt>
- [11] *Topology Broadcast based on Reverse-Path Forwarding (TBRPF)* Ogier, R.G., Lewis, M.G., Templin, F.L., Bellur, B., IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-tbrpf-06.txt>

-
- [12] *Landmark Routing Protocol (LANMAR) for Large Scale Ad Hoc Networks* Gerla, M., Hong, X., Pei, G., IETF Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-lanmar-05.txt>
- [13] *Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers* ACM SIGCOMM, vol.24, no.4, October 1994 <http://www.iprg.nokia.com/charliep/txt/sigcomm94/paper.ps>
- [14] *Optimized Link State Routing Protocol* Clausen, T., Jacquet, P., Laouiti, A., Minet, P., Muhlethaler, P., Qayyum, A., Viennot, L., IETF Internet Draft, <http://menetou.inria.fr/olsr/draft-ietf-manet-olsr-07.txt>
- [15] *Providing Ad-Hoc Connectivity with Reconfigurable Wireless Networks* Haas, Z.J., Pearlman, M.R., Proceedings of the ACM SIGCOMM '98, September, 1998.
- [16] *The Performance of Query Control Schemes for the Zone Routing Protocol* Haas, Z.J., Pearlman, M.R., IEEE/ACM Transactions of Networking, Vol. 9, No. 4
- [17] *Determining the Optimal Configuration for the Zone Routing Protocol* Haas, Z.J., Pearlman, M.R., IEEE JSAC, August, 1999. http://www.ee.surrey.ac.uk/Personal/G.Aggelou/PAPERS/jsac99_paper.pdf