ZRP simulation on JiST/SWANS

Performance Observations & Summary

Aaron Beach

Northwestern University

ABSTRACT

Routing in ad-hoc networks is a complex task for which multiple approaches have been proposed. ZRP is a hybrid of "pro-active" and "reactive" routing methods. ZRP attempts to take advantage of pro-active routing reliability within a scalable size zone and to allow scalability beyond the zone through reactive routing. ZRP presents a framework that can be tuned and fitted with reactive and pro-active elements to match application. This paper presents the simulation of ZRP on JiST/SWANS. Results are presented and analyzed in order to understand the practicality of ZRP as a routing protocol for car-to-car mobile ad hoc networks on real world streets.

Keywords

Mobil Ad hoc Network (manet), car-to-car network, Zone Routing Protocol (ZRP), Simulation

1 INTRODUCTION

This paper will analyze the merit of current ZRP implementations concerning car-to-car Mobile Ad-hoc Networks (manet) which exhibit high link state variability. Section 2 gives a short introduction to the ZRP framework, more in depth discussions can be found in [1,2]. Section 3 introduces the JiST/SWANS (Java in Simulation Time / Scalable Wireless Ad hoc Network Simulator)[3] platform that is used to simulate ZRP on the network. Section 4 discusses how data/message traffic and car mobility are modeled. Section 5 presents results of the simulation and analyzes them to draw conclusions about the merit of ZRP concerning the application (other ZRP analysis can be found in [4]). Section 6 then offers solutions and future work suggestions that would proceed from our analysis.

2 ZRP OVERVIEW

Mobile ad-hoc networks are networks that involve mobile nodes communicating over wireless connections. The networks are characterized by variable topologies (connectivity). Therefore static routing paths, like those used in conventional routing, cannot be used.

The dynamic properties of links within a manet require active topology discovery and maintenance to do any kind of deterministic routing. Continual variability, along with the existence of sudden and possibly unidirectional connections, makes routing on these networks very complex.

Approaches, termed *pro-active*, work much like traditional link-state routing algorithms, these approaches attempt to maintain an up to date view of the network locally and use this information to route packets towards their destination. These approaches have been found not to scale well, requiring exponentially increasing control traffic in order to maintain their connection state information. Therefore, another approach was developed that could scale without the need of consistent control traffic: *reactive* routing. The idea behind reactive routing is to only find connection state when it is needed. When a message is to be sent, a node floods the network with a route request. Everyone that hears the request appends their address and forewords the request. Once the destination is found, the request is routed back to the source using the appended addresses. The source then uses the list of addresses to route the message to its final destination. This incurs an extra round trip latency and higher chance of route failure in exchange for control traffic that scales linearly (with the amount of messages to be sent).

ZRP attempts to fuse these two approaches and benefit from each of their strengths, while minimizing their weaknesses. The approach is based on breaking the topology down into zones. Within these zones routing can be scalable and optimized using pro-active routing. Beyond the zone, reactive routing can be made more efficient as a zone to zone request protocol, more efficiently pruning requests and extending each hop to cover multiple nodes. ZRP refers to it's proactive routing as the Intra-zone Routing Protocol (IARP), and the reactive routing as the Inter-zone Routing Protocol (IERP). A zone is defined for each node as the hop “radius” to any other node. Any node within ***ρ*** (radius) hops from a particular node is in its zone.

The Node Discovery Protocol (NDP) is used to manage connection state within the zone, usually involving periodic “heartbeat” messages and exchange of routing table information with neighbors.

The Bordercast Routing Protocol is used to “glue” the reactive and proactive routing elements together. A node is considered another node’s “border node” if it is exactly ***ρ*** (zone radius) hops away. Nodes use their “Intra-zone” (proactive) routing protocol to foreward “Inter-zone” (reactive) protocol messages to border nodes. This process is repeated between zones avoiding flooding redundancy within the zone.

3 JiST/SWANS

Java in Simulation Time/
Scalable Wireless Ad hoc Network.

JiST is an general purpose event driver simulator based on the idea of *virtual machine based simulation*. This very efficient simulation platform supports object oriented Java programming making it very easy to build upon.

SWANS is a layered network simulator built on top of JiST. Its goal is to simulate ad hoc networks in a way that is scalable and easily adaptable to your specific simulation needs. JiST achieves its scalability through a *Hierarchical Binning* spatial data structure[2]. The layered approach along with object oriented java allows easy and quick rewriting of any aspect of the network simulation.

4 SIMULATION MODELS

The simulation attempts to model cars moving on streets throughout a city in a realistic manner. For this a mobility model was developed to allow realistic car movement that includes following speed limits, slowing down before turns or when other cars are in front of you, and occasionally having to stop at intersections for red lights. The mobility model is discussed in more depth in the following. Street maps from the US Census Bureau’s Tiger/LINE data[5]. Streets are broken down into small sections and classified by street-type, this data is then parsed and used as routing paths for our car mobility. The mobility model is discussed in depth here[?].

5 RESULTS & ANALYSIS

The simulation is analyzed by varying three main environment factors and looking at particular performance metrics. (The network can be viewed as a graph the words “car” and “node” may be used interchangeably.)

Factors:
- Car Density: cars within a fixed size map (.8 sq miles)
- Send Rate: rate of message injection into the system
- Zone Radius: Max node-hop distance within zone

Metrics:
- Message Delivery Ratio: % of messages that reach destination
- Mac Errors: Mac level errors such as collision
- IERP success ratios (Request, Response, Send)

Figure 1 shows the relationship between node density and message delivery %, showing that there indeed exists an optimal node density. Why would this exist? Because delivery ratio requires that the graph of nodes be connected (node degree) and avoid collisions. As the map begins to grow in density from zero nodes collisions are rare and connectivity is low, most messages are not delivered due to the fact that no one is within range. When the graphs becomes more and more connected messages begin to be delivered, however as density continues to increase suddenly we see a drop in message delivery ratio. By lowering the send rate, and lowering the chance that two nodes’ transmissions collide, we can achieve a higher number of nodes delivering optimal performance.

Figure–1

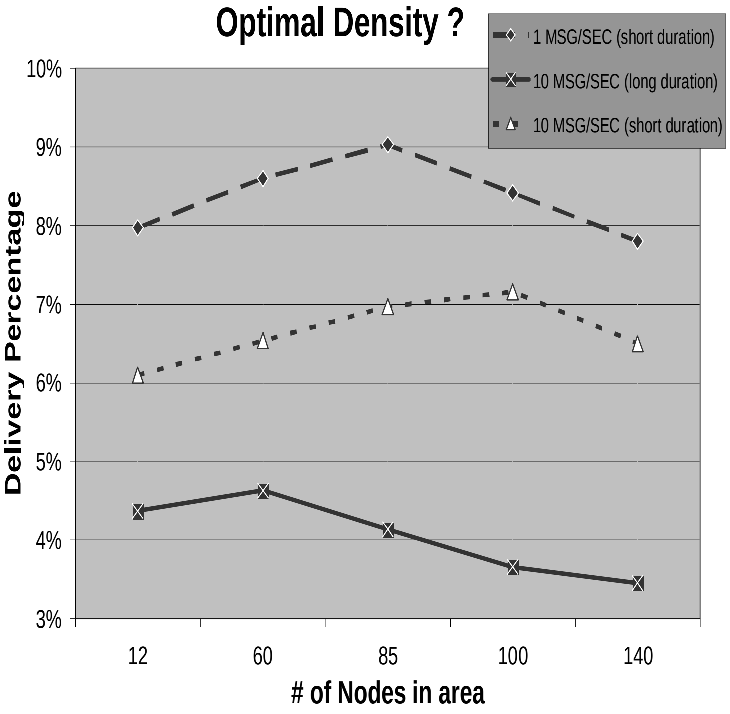

However, due to the simulation starting with a clean slate, there is initially little message forwarding traffic in the system, once messages begin to be sent, message density increases and therefore increases collision ratios. It was actually found that when allowed to live for a long time and to retransmit route requests, some reactive routing attempts were requiring overall (request ,response ,send) durations in the minutes (even with message pruning). This allowed message density to build and continue to degrade message delivery percentages to equilibrium of about 4%, depending on node density.

To understand the message routing failure, one must look more carefully at where ZRP mixes the two types of routing, then maybe one can understand what in the protocols is failing. For this the radius was configured to different sizes.

In figure-2 we can see that both proactive and reactive protocols have their merits. The interaction between zone radius and the protocols can be understood to mean the following: When the radius is zero you have complete reactive routing; when the radius is equal to the greatest number of hops on the shortest path between any two nodes in the system (i.e. all nodes are in all other node's zone) then routing is completely pro-active. Figure-2 shows that a mix of small pro-active zones is optimal, yet still rather ineffective.

Figure-3 and Figure-4 show trends in mac collisions and the number of messages that are reactively routed (IERP). As the system changes from predominantly reactive to proactive you can see that along with a decline in reactive routing there is sharp increase in mac collisions representative of the un-scalable behavior of pro-active routing. Collisions become much more common as Intra-zone control traffic increases.

We can see that the system cannot work with proactive routing beyond a certain zone size, however why doesn't a greater amount of reactive routing still not achieve a decent message delivery ratio?

First, an understanding of the IERP process. There are three main parts you can think of the IERP route request as:

1. →Flood Request (source)
2. ←Receive request and Respond (destination)
3. →Receive response and route message (source)

Figure-2

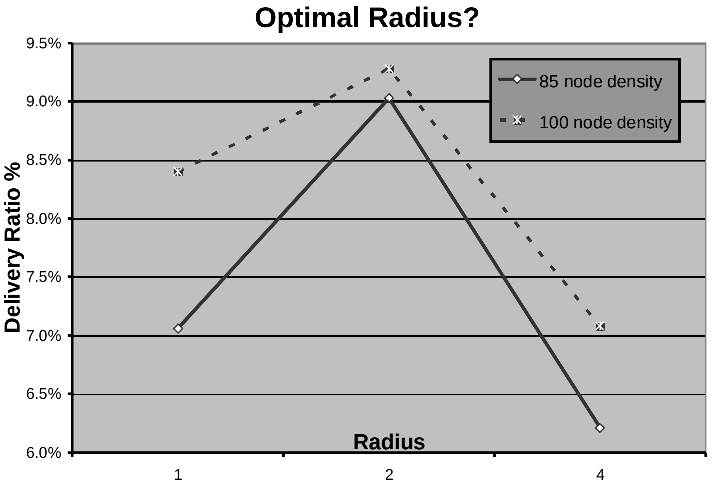

Figure-3

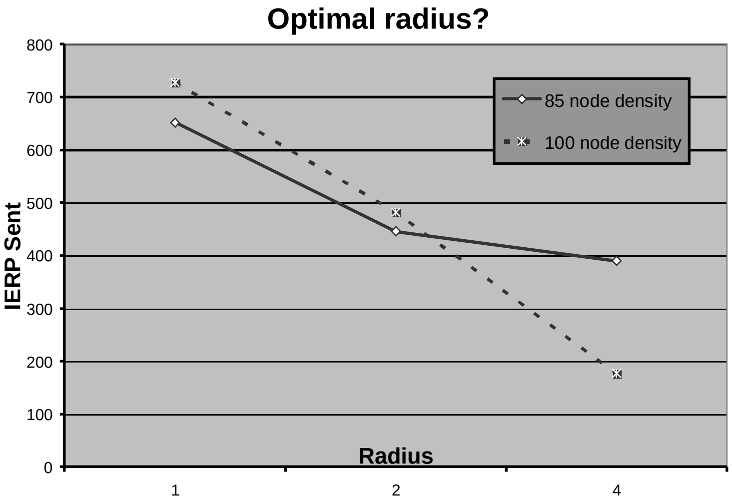

Figure-4

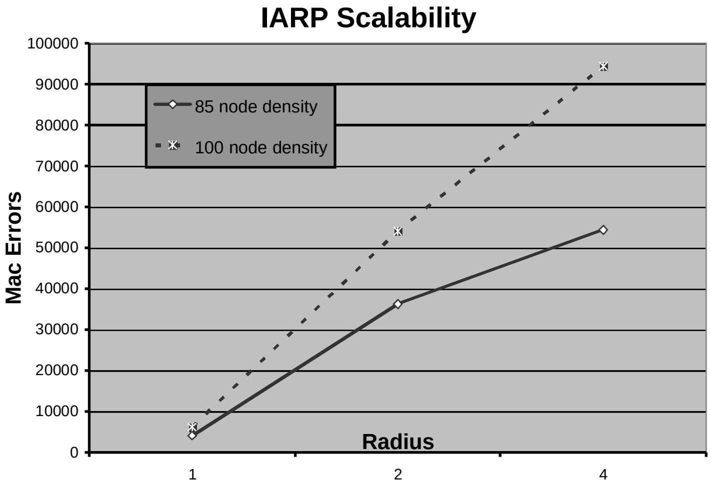

The simulation however showed that reactive routing requires links to be static for a certain amount of time (at least a round trip for route discovery). However, this was not the case in the mobile network. Cars were often moving in opposite or perpendicular directions and only within contact for very short periods. While small clusters (small zones) that travel in the same direction can maintain a stable distance, the majority of car interaction cannot be counted on for stable connections for a very long duration at all. It is actually quite intuitive that car interaction at low densities would exhibit this behavior.

This network property made reactive routing nearly impossible; to understand this, consider each part of IERP.

Flood Request: actually a much larger percentage of these messages successfully reached near the destination successfully, which is attributed to the fact that flooding does not depend on static routes. We typically observed in the range of 60-95% success rate depending on zone radius and car density, which when tuned properly can produce high percentages of flood requests reaching a destination (at least once). This is due to the fact that as density increases connectivity increases (the network graph average degree increases).

Route Response: This is where reactive routes were first used, and when routes are discovered to be rarely static. Depending on network configuration 10-25% success rates were observed, most of these attributed to short routes 1 zone away, which due to BRP did not actually rely on the source route to return to the sender.

Route Message Successfully: We see a somewhat higher delivery ratio in this stage, which implies that if a route is stable throughout one full traversal it has a better chance of being stable throughout subsequent traversals. We typically observed (depending on network configuration) in the range of 25-75% success routing on routes that had already been returned successfully. Remember, many of these cases are accounted for by single zone routing.

Can it then be said optimistically (based on our results) that given good network conditions and configuration, reactive routing may achieve approximately (95% * 25% * 75%) 18% message delivery?
Not in any of our tests. Through find-tuning of the network and traffic shaping reactive routing may be able to achieve higher levels of message delivery. However, our simulation was not even able to reach to the 300-400 car density that we had planned for the .8 sq mile map area (that can hold a max of approximately 3000 cars in bumper to bumper traffic). The simulation would become overrun with collisions and zone control traffic such that it became useless to try to route any messages and our event simulator was overrun with rapid collisions and re-transmits of heartbeat messages. Some of these collision problems may have been controlled through limiting zone size however, overlapping zones would still collide with each other, and we have seen that reducing the zone radius will not solve the problem (leaving us with purely reactive routing).

Reactive routing must be able to achieve long distance routing to really be useful, however given a link failure rate of x over time t (t being the time to traverse one link), on a route of n links, the success ratio S can be modeled by the following equation:

$$S = (1 - t*x) * (1 - 2t*x) * \ldots * (1 - n*t*x)$$

For modest numbers
$t*x = .1$ $\quad$ $n = 4$
$S = .3 = (1 - .1)(1 - .2)(1 - .3)(1 - .4)$

Given this relationship between route length and route failure, it is obvious that a highly variable topology would result in high route failure for reactive routing. This shows that while not producing exponential amounts of traffic, reactive routing does not scale due to route failure increasing exponentially.

6 CONCLUSIONS

The simulation has exposed the following facts about routing protocols for use in car-to-car mobile ad hoc networks:

- Pro-active routing can be useful, but only in a small enough zone-radius to avoid control traffic overloading the network. Any use of pro-active routing should be accompanied by strict bounds on the levels of control traffic and size of the proactive zone.

- Protocols that rely on traditional reactive route discovery require a certain level of link stability that can not be guaranteed by a realistic car-to-car manet, and therefore do not scale in these networks.

- As car-to-car network grow denser overall connectivity increases, however average link stability does not, this accounts for high levels of flood requests reaching their destination while static routes failed.

- A routing protocol for a car-to-car manet should not rely on link stability, however it may take advantage of high levels of connectivity as the network becomes more dense. The routing protocol must not require large amounts of control traffic to find routes to destinations. Current pro-active, reactive, and hybrid protocols fail to achieve these requirements.

7 Future Work

If a real car to car manet is to be realized then these problems must be solved.

Future work may be able to solve these problems by taking a location aware approach. The existence of location aware nodes would introduce a method for routing that does not require large amounts of control traffic. Location aware routing may be able to take advantage of high connectivity in spite of low link stability and could still possibly benefit from pro-active routing within limited zones. Reactive routing may also benefit by replacing the route instead with a location that is not dependent on link stability. However, given the new approach, modifying existing ad hoc protocols may be less advantageous than applying pure geographic routing schemes.

8 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] http://jist.ece.cornell.edu/

[4] *Determining the Optimal Configuration for the Zone Routing Protocol* Haas, Z.J., Pearlman, M.R., IEEE JSAC, August, 1999.

[5] http://www.census.gov/geo/www/tiger/