OVERHEAD ANALYSIS

Subject: Overhead analysis
Date: Tue, 27 Jun 2000 14:48:51 +0200 (CEST)
From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: manet@itd.nrl.navy.mil

Dear MANet members,

After the discussions on the mailing, we have been thinking quite a lot with Philippe on overhead. We have developed some analysis about the relation between routing performance (overhead) and networks parameters (size, mobility, traffic). Our aim is to identify the scenarios which favor such and such MANet protocols. We consider a mobile ad-hoc network with N nodes (N assumed large, say N>50 nodes).

We have identified two classes of protocols which more or less fit the reactive and proactive boundary:

1. the flooding protocols (AODV/DSR, TORA)
2. the hello protocols (DSDV,OLSR)

The performance are quantified by the quantity of overhead generated by each protocols. We have identified two source of overheads:

A. the control overheads
B. the route overheads

Our main findings are the following:

-Both classes have O(N^2) control overheads which is no surprising from graph theory point of view. Flooding algorithm and hello algorithm have different factors in front of N^2 which actually depends on traffic, geometry, mobility. One can find scenarios in favor of flooding and other scenarios in favor of hello.

-Flooding protocols may not provide optimal routes (in term of hop number) the average deviation between optimal routes can be significant (a factor of 33% in dense 1D networks, more in 2D and 3D). The overhead generated by the extra transmission of data may be important since it is proportional to data traffic.

A research report detailing these results is available:
http://menetou.inria.fr/~viennot/postscripts/overhead.ps.gz (166 Kb)

Comments are welcome,

Laurent

Subject: Re: Overhead analysis
Date: Fri, 30 Jun 2000 07:47:45 -0400 (EDT)
From: Zygmunt Haas <haas@ee.cornell.edu>
To: Laurent Viennot <Laurent.Viennot@inria.fr>
CC: manet@itd.nrl.navy.mil

Hi Laurent,

There is another "class" of protocols, the hybrid class, which is a mix of proactive-reactive behavior. Examples of such protocols include the Zone Routing Protocol (ZRP) and the Fisheye protocol.

As you have stated, pure proactive/reactive protocols may not present "optimal" performance for a specific network operational conditions (mobility and traffic). The bigest advantage of the hybrid class is that it can adapt itself to the network operational conditions. More specifically, by dynamically adjusting the ratio of the proactive vs. reactive behavior, the amount of overhead control traffic can be *drastically* reduced. In the Zone Routing Protocol, this is achieved through the use of a single parameter - the Zone Radius, which is updated based on continual measurement of the mobility (the rate of creation/destruction of links with a node's neighbors) and the amout of traffic a node sees. (More on updating of the zone radius could be found in our paper in JSAC / August 1999.)

I guess, what I am trying to say is that the hybrid class of protocols is the answer to the conclusions of your study - that in order to "cover" a broad range of MANET topologies (size of networks) and operational conditions, a *hybrid* strategy should be used.

Zygmunt.

Subject: Re: Overhead analysis
Date: Mon, 03 Jul 2000 14:27:18 +0200
From: jacquet@menetou.inria.fr
To: Zygmunt Haas <haas@ee.cornell.edu>, Laurent Viennot <Laurent.Viennot@inria.fr>
CC: manet@itd.nrl.navy.mil

Hi, Zygmunt,

It is not exactly the conclusion of our study. It looks like that hybrid protocol overheads are roughly the sum of flooding protocol overheads and hello protocol overheads. They would be very likely always above. This is the reason why we a priori prefered not to include it in the first version of our study.

Philippe and Laurent

PS: Laurent just got a baby yesterday, so I replace him on the spot.

Subject: Re: Overhead analysis
Date: Mon, 03 Jul 2000 15:38:33 +0100
From: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
Organization: University of Surrey, Guildford, England
To: jacquet@menetou.inria.fr, manet <manet@itd.nrl.navy.mil>

Philippe how about including in your results a routing scheme that uses localisation for route discovery as well as for the route repair of broken data paths.. Such as scheme could be the RDMAR protocol..-:)

Although RDMAR is referred in your report, I would be very much interested to see it in your analysis too.

Regards,
George A.

Subject: Re: Overhead analysis
Date: Mon, 3 Jul 2000 10:48:35 -0400 (EDT)
From: Zygmunt Haas <haas@ee.cornell.edu>
To: jacquet@menetou.inria.fr
CC: Laurent Viennot <Laurent.Viennot@inria.fr>, manet@itd.nrl.navy.mil

Hi Philippe,

Hmmm, this is not so. Hybrid protocols, through their ability to adjust to the network operational conditions will *always* result in less control traffic than pure proactive or pure reactive protocols. This statement is nearly self-explanatory, as *hyrbrid protocols degrade to either pure reactive or pure proactive behavior, based on their parameter settings.*

Let me comment on the last statement: in the Zone Routing Protocol (ZRP), for instance, there is a single parameter to set - the Zone Radius (ZR). If ZR=1, then ZRP degrades to pure reactice protocol, while when ZR=\infty, then ZRP degrades to pure proactive protocol. The characteristic curves of ZRP look as follows: (see for instance the paper: [M.R. Pearlman and Z.J. Haas, "Determining the Optimal Configuration of for the Zone Routing Protocol," IEEE JSAC, special issue on Ad-Hoc Networks, vol. 17, no.8, August 1999])

total traffic

```
^
|*
| *                     *
|  *                 *
|   *             *
|   *           *
|    *        *
|     *     *
|      *   *
|       ***
|
---------|--------------> Zone Radius
    the optimal ZR
```

The protocol should dynamically adjust the ZR, so as to minimize the total control traffic. Depending on the operational conditions of the network, the reduction of the hybrid scheme can be "quite dramatic" - please see the above paper for more quantitative results.

Thus, repeating myself, the hybrid protocols yield *much* reduction in the control traffic and are a natural choice in an environment with changing operational conditions, as well as for a broad range of MANETs.

Zygmunt.

Subject: Re: Overhead analysis
Date: Mon, 03 Jul 2000 19:29:19 +0200
From: jacquet@menetou.inria.fr
To: Zygmunt Haas <haas@ee.cornell.edu>
CC: Laurent Viennot <Laurent.Viennot@inria.fr>, manet@itd.nrl.navy.mil

Dear Zygmunt,

I feel uncomfortable to talk about ZRP because it does not specify the link state algorithm it uses for intra-zone routing. If you use non-optimized OSPF link state, then I will agree with your plot, but if you use another optimized link state, like OLSR for example, you will obtain a completely different plot.

As we see in our study, the leading overhead in OLSR is in the hello transmissions. Therefore ZRP implemented on OLSR will likely do worse than OLSR or flooding protocols as soon ZR is greater than 1 and smaller than infinity, because ZRP would include both hellos cost for intra-zone and flooding cost for inter-zone. We also expect that ZRP on OSPF will do worse than ZRP on OLSR.

For example, I would expect a plot like this one, which depends on traffic scenarios, mobility, etc... But I am not sure at 100 %, it needs more work and we prefer to skip this in our first version.

```
total traffic ZRP on OLSR
        ^
        |
        |
        |      * * * * *
        |   *           *
        |  *               *
        |
        | *                  *
        |
        |
        |
        ---|------------------|-----> Zone Radius
            the optimal ZR
```

Best regards, Philippe

PS: Laurent will come back soon, I will not able to continue this interesting discussion because I leave for a one week travel.

Subject: Re: Overhead analysis
Date: Mon, 3 Jul 2000 15:50:13 -0400 (EDT)
From: Zygmunt Haas <haas@ee.cornell.edu>
To: jacquet@menetou.inria.fr
CC: Laurent Viennot <Laurent.Viennot@inria.fr>, manet@itd.nrl.navy.mil

Hi Philipe,

> As we see in our study, the leading overhead in OLSR is in the hello transmissions.

I would agree with this.

> Therefore ZRP implemented on OLSR will likely do worse than
> OLSR or flooding protocols as soon ZR is greater than 1 and smaller than
> infinity, because ZRP would include both hellos cost for intra-zone and
> flooding cost for inter-zone.

No, this is not correct. For any specific intra-zone protocol, one can obtain a *minimum* by choosing the right ZR. The point that you are missing is the way the inter-zone is implemented in ZRP - it is NOT a regular flooding. Rather, the inter-zone is implemented through bordercasting (which, in our newer version, is done through multicasting). Thus the ZRP cost is NOT a simple addition of the intra-zone and flooding costs. In other words, indeed, as the ZR increases, the cost of the intra-zone increases, BUT the cost of the inter-zone decreases. ***This is the whole idea of the hybrid protocol like ZRP. ***

Again, the total control traffic in ZRP as a function of the Zone Radius behaves as follows:

```
total traffic
        ^
        |*
        | *                     *
        |  *                 *
        |   *              *
        |   *            *
        |    *         *
        |     *      *
        |      *   *
        |       ***
        |
        -|-------|----------------|--> Zone Radius
        pure  optimal ZR         pure
      reactive                 proactive
      behavior                 behavior
```

Again, please refer to our paper: [M.R. Pearlman and Z.J. Haas, "Determining the Optimal Configuration of for the Zone Routing Protocol," IEEE JSAC, special issue on Ad-Hoc Networks, vol. 17, no.8, August 1999] - the interplay between the cost of the intra-zone and the inter-zone routing is clearly explained there.

> We also expect that ZRP on OSPF will do worse than ZRP on OLSR.

This is possible - optimizing the intra-zone protocol will, indeed, result in lower total ZRP cost. But this is not the issue. The issue is that once you have chosen the intra-zone (proactive protocol, whether OSPF, OLSR, or any other you-favorite-procative protocol), the hybrid mechanism allows you to adjust to the operational conditions of the network.

Let me try to clarify the last statement again. For a highly mobile network, a hybrid protocol with a smaller ZR will do better (as the cost of maintaining the states in a zone (through proactive protocol) will be too costly). On the other hand, if a network is nearly stationary and there are no changes, large ZR (in the limit the whole network) will lead to smaller total overhead. Again, note that the total cost is the sum of the intra-zone and inter-zone routing costs. The first increases, while the second decreases as the ZR increases.

I think that your confusion stems from the fact that you overlooked the complementary behavior of the two costs.

Regards, Zygmunt.

Subject: Re: Overhead analysis
Date: Mon, 03 Jul 2000 16:22:11 -0700

From: Richard <ogier@pit.erg.sri.com>
To: Zygmunt Haas <haas@ee.cornell.edu>
CC: jacquet@menetou.inria.fr, Laurent Viennot <Laurent.Viennot@inria.fr>, manet@itd.nrl.navy.mil

Zygmunt,

> Let me try to clarify the last statement again. For a highly mobile
> network, a hybrid protocol with a smaller ZR will do better (as the cost
> of maintaining the states in a zone (through proactive protocol) will be
> too costly).

Is the above statement true even if the traffic demand is heavy and uniformly distributed? I think this was discussed before. In this case, in a highly mobile network with heavy uniform traffic, each node must be provided with frequently updated paths to all destinations, which is what a proactive protocol provides. And an efficient proactive protocol (e.g., OLSR or STAR or TBRPF) may be able to do this more efficiently than a reactive protocol.

Richard

Subject: Re: Overhead analysis
Date: Tue, 4 Jul 2000 18:59:31 -0400 (EDT)
From: Zygmunt Haas <haas@ee.cornell.edu>
To: Richard <ogier@pit.erg.sri.com>
CC: jacquet@menetou.inria.fr, Laurent Viennot <Laurent.Viennot@inria.fr>, manet@itd.nrl.navy.mil

Dear Richard,

You have asked a very interesting and important question. I have been a bit "simplistic" in my previous postings - my goal was to justify the claims that the hybrid protocols can do better, sometimes drastically better, than pure reactive or proactive protocols.

However, in fact, there are *two* parameters that control what the "optimal" mixture of practive vs. reactive behavior of a protocol should be: the mobility of the nodes (how fast nodes move relative to their neighbors, and, thus, how often they break the links with their neighbors) and the activity of the nodes (how often they initiate the route discovery process). For the sake of simplicity, let's assume that the values of these two parameters are constant throughout the network. (Of course, in practical cases this is far from being true, especially when the network is large.) Let me first state the claim and then comment.

The Claim: The larger the mobility of the network is and the smaller the activity, the more reactive the routing protocol should be. And, vice versa, the smaller the network mobility and the larger the activity, a more proactive protocol would lead to smaller total amount of control traffic.

The justification is as follows, if the mobility is large, using a proactive scheme would lead to a situation in which the resources (wireless, processing, etc) used to "learn" the topology of the network would be wasted, as by the time that the information about the topology is used, the topology has already changed due to the large mobility of the nodes. Thus, more reactive appraoch would be better here. And vice versa, if the mobility is small ...

Also, the larger the network activity is, a proactive approach would engage the route discovery process quite often. Thus learning the topology information would lead to less control traffic. Such learning of topology is, of course, done by a proactive protocol. And vice versa, if the network activity is small ...

Of course, practical cases are in between ... this is why some mixture of proactive/reactive behavior is so beneficial.

Relating this to the Zone Routing Protocol (ZRP), the larger the network mobility is and the smaller the network activity is, the more reactive the protocol should be and, thus, the smaller the Zone Radius (ZR) should be. In the limit, ZR=1, which is the pure flooding case - in which the changes of the network are so frequent that learning the topology is a total waste, since by the time it is used, the topology has already changed. Similarly, if the network activity is low, learning the network activity would not pay off, as the topology information is very infrequently needed.

On the other side of the spectrum, if the mobility in the network is low (a stationary network, in the limit) and the network activity is high, the protocol should be more proactive, and the ZR should be large - in the limit the ZR=the network diameter (and conceptually, ZR=\infty). The reasons are as stated above.

Happy 4th of July!

Zygmunt.

Subject: RE: Overhead analysis
Date: Thu, 06 Jul 2000 15:40:59 +0100
From: "Groten, D." <D.Groten@kpn.com>
To: Laurent Viennot <Laurent.Viennot@inria.fr>, "'jacquet@menetou.inria.fr'" <jacquet@menetou.inria.fr>
CC: manet@itd.nrl.navy.mil, Zygmunt Haas <haas@ee.cornell.edu>

Hi Philippe and Laurent,

I have a reaction on your paper concerning the terminology you use: on the one hand, we can distinguish between on-demand (reactive) and table-driven (proactive) protocols. On the other hand, you have identified the first category as "flooding", and the second category as "hello". I find this quite confusing: both types of protocols flood routing information over the network. And while "hello"-type neighbour discovery is mandatory in a table-driven protocol, it may also be used in an on-demand protocol in order to establish and repair routes more quickly.

Of course, strictly speaking, a purely on-demand protocol cannot use "hellos" because such neighbour discovery is performed pro-actively. But from the point-of-view of the general concept, the distinction between on-demand (end-to-end routes are only found when necessary) and pro-active (end-to-end routes always available) is more easy to grasp than that between hello and flooding.

I have two questions:
1. Am I correct when concluding that your analysis is based purely on the extreme case where on-demand actually is equivalent to flooding of the entire network (no hellos to help)?
2. Is the combination of a simple neighbour discovery (via pro-active hellos) and on-demand route discovery equivalent to ZRP with zone radius 1?

Regards, Dirk.

Subject: RE: Overhead analysis
Date: Fri, 7 Jul 2000 16:02:35 -0400
From: Marc Pearlman <pearlman@ee.cornell.edu>
Organization: Cornell University EE

To: "'Groten, D.'" <D.Groten@kpn.com>
CC: "manet@itd.nrl.navy.mil" <manet@itd.nrl.navy.mil>

Hi Dirk,

> 2. Is the combination of a simple neighbour discovery (via pro-active
> hellos) and on-demand route discovery equivalent to ZRP with zone radius 1?

Yes.

Technically speaking, a purely reactive ZRP configuration would correspond to a zone radius == 0 hops. In such a case, there is no neighbor discovery and route requests (RREQs) are distributed via unreliable shared-channel neighbor broadcast.

Sometimes, unreliable shared channel neighbor broadcast is not advisable (due to low success rate) or even available (in a multiple channel system). In these situations, neighbor discovery is needed, either to track neighbor ACKs from a *reliable* neighbor broadcast or to resolve the neighbor broadcast into individual unicasts to each neighbor. If this neighbor information is provided to the ZRP, then the zone radius will be == 1 hop

Within this context, we sometimes refer to a zone radius == 1 as being purely reactive. What we really mean is that this is the *most reactive* configuration possible, given that neighbor discovery is already present to support other networking services.

regards, Marc

Subject: RE: Overhead analysis
Date: Mon, 10 Jul 2000 14:38:58 +0200 (CEST)
From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: manet@itd.nrl.navy.mil

Hi Dirk,

A word about our terminology: We call flooding protocols, protocols that construct routes by flooding like AODV.

We call hello protocols those that discover the topology by neighborhood hellos (hellos containing lists of neighbors) and that use the knowledge of the topology to optimize broadcasting like OLSR (see the nice properties of multi-point relay sets).

This can be seen as a refinement of the reactive/pro-active classification. We have chosen a different terminology to insist on a classification according to the way routes are constructed.

> I have two questions:
> 1. Am I correct when concluding that your analysis is based purely on the
> extreme case where on-demand actually is equivalent to flooding of the
> entire network (no hellos to help)?

No.

> 2. Is the combination of a simple neighbour discovery (via pro-active

> hellos) and on-demand route discovery equivalent to ZRP with zone radius 1?

I guess yes, Zygmunt is the right person for answering that.

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 14:48:27 +0200 (CEST)
From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: Zygmunt Haas <haas@ee.cornell.edu>, Richard <ogier@pit.erg.sri.com>
CC: manet@itd.nrl.navy.mil

Dear Zygmunt and Richard,

Sorry for answering so late but I was taking care of my new born child.

Both protocol approaches behave differently according to mobility and activity mainly. To summarize: flooding (reactive) protocols react better to high mobility and hello (pro-active) protocols react better to high activity. We agree on this.

The topology of the network is also important because it determines how much hello protocols can optimize broadcasting. For example, the random graph model is a difficult case for hello protocols.

We did not analyze ZRP because it depends greatly on the pro-active protocol used. My understanding is that the curve of your JSAC paper supposes a DSDV like protocol and that Philippe obtains a different curve when OLSR is used.

Something I do not understand is why your curve does not have a discontinuity at ZR=0 (or ZR=1) since pure flooding do not include hello cost and ZRP with ZR=2 includes the hello cost of the pro-active protocol. I would have expected a curve like this :

```
total traffic
        ^
        |
        | *                     *
        |  *                 *
        |   *              *
        |   *            *
        |*   *         *
        |.    *      *
        |.     *   *
        |.      ***
        |.
        -.-------|--------------> Zone Radius
         .  the optimal ZR
         .
         .<--- hello packets ---->
         .
   no hello (pure reactive)
```

laurent

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 14:52:31 +0200 (CEST)

From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
CC: manet <manet@itd.nrl.navy.mil>

Dear George,

With use of localisation, RDMAR can optimize the flooding cost by restricting the set of nodes that re-emit a flooding packet. To extend our analysis to RDMAR, you just need to evaluate the average proportion of nodes that participate to a flooding.

Laurent

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 09:10:27 -0500 (CDT)
From: Nitin H Vaidya <vaidya@cs.tamu.edu>
To: Laurent.Viennot@inria.fr
CC: manet@itd.nrl.navy.mil

hi Laurent

>> From: Laurent Viennot <Laurent.Viennot@inria.fr>
...
>> If planar coordinates are available to each node, what about limiting
>> the flooding to an even narrower area, for example a rectangle
>> containing the source and the estimated position of the destination ?
>> Do you have some hints about such ideas ?

Yes ... you might want to take a look at the paper on Location-Aided Routing from the 1998 MobiCom, Dallas.

- nitin

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 15:30:31 +0100
From: Nikos Triantafillis <N.Triantafillis@cs.ucl.ac.uk>
To: Laurent.Viennot@inria.fr
CC: manet@itd.nrl.navy.mil

>If planar coordinates are available to each node, what about limiting
>the flooding to an even narrower area, for example a rectangle
>containing the source and the estimated position of the destination ?
>Do you have some hints about such ideas ?

There is some work done I believe in this area where flooding is limited to an even narrower space eg: Young-Bae Ko and Nitin H. Vaidya. "Location-Aided Routing (LAR) in Mobile Ad Hoc Networks". Proceedings of Mobicom '98 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking, pp. 66-75, Dallas, TX, October 1998, including some other papers on geographic routing and maintenance.

Nick

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 16:24:45 +0100
From: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
Organization: University of Surrey, Guildford, England
To: Laurent.Viennot@inria.fr, manet <manet@itd.nrl.navy.mil>

> If planar coordinates are available to each node, what about limiting
> the flooding to an even narrower area, for example a rectangle
> containing the source and the estimated position of the destination ?
> Do you have some hints about such ideas ?

Yes! the relative position of two mobiles (and hence their relative distance) is determined throught the Relative Distance Estimatin (RDE) algorithm in RDMAR. RDE relies its operation on a stochastic model that bases its estimations only on the previous relative distance of the two mobiles. Our results (reported in the paper I mentioned earlier) are very promising as 95-98% hit ratios (i.e., correct estimatinos) are achieved. Note finally that RDMAR protocol _does not_ assume any system parameters such as planar coordinates or the actual velicity of nodes are available to the nodes.

Regards,
George A.

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 10:57:03 -0500 (CDT)
From: Nitin H Vaidya <vaidya@cs.tamu.edu>
To: Laurent.Viennot@inria.fr, manet@itd.nrl.navy.mil

Laurent:

>> > If planar coordinates are available to each node, what about limiting
>> > the flooding to an even narrower area, for example a rectangle
>> > containing the source and the estimated position of the destination ?
>> > Do you have some hints about such ideas ?

By the way, also take a look at the Query Localization scheme from 1999 MobiCom by Castaneda & Das. That scheme also limits route discovery to a small region (as does Location-Aided Routing LAR), but without using physical location information (unlike LAR).

Cheerio. – nitin

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 10:31:44 -0700
From: Richard <ogier@pit.erg.sri.com>
To: Laurent.Viennot@inria.fr
CC: Zygmunt Haas <haas@ee.cornell.edu>, manet@itd.nrl.navy.mil, ogier@erg.sri.com

Hello Laurent,

> Both protocol approaches behave differently according to mobility and

> activity mainly. To summarize: flooding (reactive) protocols react
> better to high mobility and hello (pro-active) protocols react better
> to high activity. We agree on this.

Yes, we agree on this. But what if there is both high mobility and high activity? My *guess* is that in this case, an efficient proactive protocol generates less control traffic. My reasoning is that every node needs to maintain frequently updated paths to every other node, and efficient proactive protocols do this more efficiently than flooding.

Richard

Subject: Re: Overhead analysis
Date: Mon, 10 Jul 2000 13:44:53 -0700
From: "Charles E. Perkins" <charliep@iprg.nokia.com>
Organization: Nokia Research Center
To: Richard <ogier@pit.erg.sri.com>
CC: manet@itd.nrl.navy.mil, ogier@erg.sri.com

Hello Richard,

> .... But what if there is both high mobility
> and high activity? My *guess* is that in this case, an efficient
> proactive protocol generates less control traffic.
> My reasoning is that every node needs to maintain frequently
> updated paths to every other node, and efficient proactive protocols
> do this more efficiently than flooding.

I believe this effect has been measured, although I do not have the reference, and indeed it is possible for proactive protocols to outperform on-demand protocols if all the nodes talk to all of the other nodes. It also helps if the nodes stand still, or if triggered updates are batched (i.e., not all sent separately).

Regards, Charlie P.

Subject: Re: Overhead analysis
Date: Tue, 11 Jul 2000 11:05:41 +0200 (CEST)
From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: Richard <ogier@pit.erg.sri.com>
CC: manet@itd.nrl.navy.mil

Richard writes:
>
> > Both protocol approaches behave differently according to mobility and
> > activity mainly. To summarize: flooding (reactive) protocols react
> > better to high mobility and hello (pro-active) protocols react better
> > to high activity. We agree on this.
>
> Yes, we agree on this. But what if there is both high mobility
> and high activity? My *guess* is that in this case, an efficient
> proactive protocol generates less control traffic.
> My reasoning is that every node needs to maintain frequently

> updated paths to every other node, and efficient proactive protocols
> do this more efficiently than flooding.
>

Yes, this is confirmed by our analysis. We can get regions of the plane mobility x activity favorable to each protocol flavor. This gives something like this :

```
activity
   ^
   |---------------|
   |               |
   |**             |
   |  * proactive  |
   |   *           |
   |    *          |
   |      ****     |
   |          *****|
   | reactive      |
   |               |
   ----------------|-------> mobility
```

An activity of 2 or 3 active routes per node is sufficient for proactive protocols to outperform on-demand protocols. All the nodes talking to all of the other nodes corresponds to an activity of N active routes per node. You don't need so high activity to see proactive protocols behaving better than reactive protocols.

You can find similar curves for the random graph model, the geometric model (strip and square) and the grid in our research report RR-3965 http://menetou.inria.fr/~viennot/postscripts/overhead.ps

Laurent

Subject: Re: Overhead analysis
Date: Tue, 11 Jul 2000 10:21:04 -0700
From: Richard <ogier@pit.erg.sri.com>
To: Laurent.Viennot@inria.fr
CC: manet@itd.nrl.navy.mil, ogier@erg.sri.com

Laurent,

Your analysis is interesting, and I am glad to see it confirms my intuition. But your plane diagram seems contrary to your earlier statement (which I agree with) that reactive protocols favor high mobility (or proactive protocols favor low mobility). Can you explain this discrepancy? I would expect that the diagram would look more like the following:

```
activity
   ^
  |---------------|
  |               |
  |               |
  |  proactive    |
  |            ** |
  |         **    |
  |    *****      |
  |****           |
  |    reactive   |
  |               |
  ----------------|-------> mobility
```

Regards, Richard

Subject: Re: Overhead analysis
Date: Wed, 12 Jul 2000 13:03:45 +0200 (CEST)
From: Laurent Viennot <Laurent.Viennot@inria.fr>
To: Richard <ogier@pit.erg.sri.com>
CC: manet@itd.nrl.navy.mil

Hi Richard,

Both diagrams are possible depending on the topology profile of the network. The size of the packets flooded for the reactive protocols versus the size the topology update packets for proactive protocols is also very important.

Laurent

Subject:
Re: Overhead analysis
Date: Wed, 12 Jul 2000 15:24:48 +0200
From: jacquet@menetou.inria.fr
To: manet@itd.nrl.navy.mil
CC: laurent.viennot@inria.fr

Dear Zygmunt

As promized, we tried to get into an analysis of ZRP in the framework of our report. We first try to analyse ZRP over DSDV in order to cross check with your paper. And second, we try to do the analysis of ZRP over OLSR. It comes that the two plots are very different.

To simplify we only consider mobility parameter. overheads are expressed in packet number per second (as in your paper). A more involved analysis will need to consider the bandwidth in packets. We also consider the case where the network is a narrow strip of length L, because it is easier to handle. We could also add the activity parameter but the formulas would be longer.

Unfortunately we don't have a complete analysis of DSDV, so that we borrow some results from your paper. What is important is the order of magnitude of the parameters before Z_R and 1/Z_R.

In summary we obtain for Z_R>0
 ZRP over DSDV: h*N+tau*N+mu*N^2*a*L*(1-Z_R/L)/Z_R+beta*mu*N^2*Z_R/L
 ZRP over OLSR: h*N+4*tau*N*Z_R+mu*N^2*a*L*(1-Z_R/L)/Z_R+4*mu*N*Z_R

Notice that OLSR transforms some pro-active O(N^2) overheads in O(N) overheads because of multipoint relay optimization.

As you can see the order change in overheads complexity significantly modifies the plot shapes. Indeed the optimal Z_R with DSDV is the closest value between 1 and L to sqrt(a*beta)*L. For OLSR the optimal Z_R is the closest value between 1 and L to sqrt(a*N*mu*L/(tau+mu)). When N increases, this value will be likely be L. Therefore the optimal ZRP will act as a
pure proactive scheme.

Warning: It does not mean that pure pro-active protocols systematically outperform pure reactive protocols, since our model only holds for Z_R > 0.

```
 >ZRP/DSDV
 >      ^
 >      |
 >      | *                     *
 >      |  *                 *
 >      |   *              *
 >      |   *            *
 >      |    *         *
 >      |     *      *
 >      |      *   *
 >      |       ***
 >      |
 >      ---------|--------------> Zone Radius
 >             the optimal ZR
>

> ^ZRP/OLSR
 >      |
 >      | *
 >      |   *
 >      |     *
 >      |       *
 >      |         *
 >      |           *
 >      |             *
 >      |               *
 >      |                 *
 >      -------------------|----> Zone Radius
 >             the optimal ZR
```

Best regards, Philippe
Detailled analysis for ZRP over DSDV
 1. The fixed overhead of ZRP/DSDV seems to be h*N+tau*N packets per second
 where

h: hello period
tau: periodic route update
N: total number of nodes in the network

Let mu be the mobility parameter defined in our repport. The variable overhead is in two components.

2. the interzone bordercast flooding: mu*a*L*N^2*(1-Z_R/L)/Z_R
a: number of active routes per node
L: half length of the strip
Z_R: zone radius

Notice that (1-Z_R/L) is the probability that a route length exceeds Z_R and mu*L is the rate at which a route of length L fails.

We assume that the flooding via bordercast is ideal and costs only N/Z_R (omitting transmissions over multicast trees and possible loop back in flooding due to non-overlapping multicast tree).

3. the intrazone topology change update: mu*beta*N^2*Z_R/L
beta*N*Z(R)/L is the number of packets generated by a neighbor change. Your paper suggests beta approximatively equal to 2. Is the rate of neighbor change be mu*N, as suggested in your paper?

Consequently the total overhead should be

h*N+tau*N+mu*a*L*N^2*(1-Z_R/L)/Z_R+mu*beta*N^2*Z_R/L

This quantity attains its maximum at Z_R=sqrt(a/beta)L independently of mu, which is confirmed by your paper.

Detailled analysis for ZRP over OLSR we have
1. fixed overhead: h*N+2*tau*N*Z_R packet per second
h: hello frequency
tau: TC-PDU frequency

Notice that the number of multipoint relays is actually finite (2 per node in a strip, 4 or little more in 2D) and that each TC-PDU must be repeated 2*ZR times to cover the intra-zone area. This is the intend of OLSR protocol to have O(1) multipoint relay set size (O(log N) in random graphs).

The variable overheads are in two components:
2. the interzone bordercast flooding: mu*a*L*N^2*(1-Z_R/L)/Z_R
3. the intrazone component: mu*4*Z_R*N

A node will generate an extra TC-PDU only if one of its two multipoint relays is changed, thus a rate of 2*mu*N. The TC-PDU is retransmitted ZR times in the intra-zone area.

The total overhead is therefore
h*N+2*tau*N*Z_R+mu*a*L*N^2*(1-Z_R/L)/Z_R+mu*4*Z_R*N

Notice that this overhead is significantly lower than the overhead with DSDV, since many terms in N^2 have been changed in N. The maximum is attained at the closest value of Z_R to sqrt(N*a*L/(2*tau+4*mu)), i.e O(sqrt(N)). When the number of nodes N is important the optimal value is very likely L.

Subject: RE: Overhead analysis
Date: Fri, 14 Jul 2000 15:25:26 -0400
From: Marc Pearlman <mrp12@cornell.edu>

Organization: Cornell University
To: <jacquet@menetou.inria.fr>, "manet@itd.nrl.navy.mil" <manet@itd.nrl.navy.mil>
CC: <laurent.viennot@inria.fr>, "'haas@ee.cornell.edu'" <haas@ee.cornell.edu>

Hi Philippe,

> In summary we obtain for Z_R>0
> ZRP over DSDV: h*N+tau*N+mu*N^2*a*L*(1-Z_R/L)/Z_R+beta*mu*N^2*Z_R/L
> ZRP over OLSR: h*N+4*tau*N*Z_R+mu*N^2*a*L*(1-Z_R/L)/Z_R+4*mu*N*Z_R

First of all, thank you for taking the time to extend your protocol analyses to the ZRP framework. I've looked over your ZRP traffic equations for the 1-D narrow strip, and they look fine to me. Your results seem to support the assertions that we've made thus far about the ZRP.

Decomposing your final traffic equations into their respective proactive and reactive components, we have:

	reactive	proactive
DSDV	mu*N^2*a*L*(1-Z_R/L)/Z_R	h*N + tau*N + beta*mu*N^2*Z_R/L
OLSR	mu*N^2*a*L*(1-Z_R/L)/Z_R	h*N + 4*tau*N*Z_R + 4*mu*N*Z_R

For the DSDV and OLSR implementations of IARP, the reactive (IERP) traffic is identical, and of the form C*(1/Z_R) - B. The DSDV and OLSR implementations of IARP generate different amounts of proactive traffic, but both expressions take on the form A*(Z_R)+ B.

Thus, the total ZRP traffic takes on the form A*(Z_R) + B + C*(1/Z_R). A and C are positive, B may be positive or negative. (Please excuse my reuse of the coefficients A,B,C. They're just intended to serve as place holders)

If you plot the equation A*(Z_R) + B + C*(1/Z_R), you'll *always* end up with a convex ("U-shaped") curve. Depending on the actual values of A,B and C, you may only get the increasing or decreasing section of the "U". An example of this is your ZRP/OLSR curve.

```
> ^ZRP/OLSR
 >      |
 >      | *
 >      |    *
 >      |       *
 >      |          *
 >      |             *
 >      |                *
 >      |                   *
 >      |                      *
 >      |                         *
 >      -------------------|----> Zone Radius
 >               the optimal ZR
```

The fact that your curve is strictly decreasing is due to your assumption of large N.

> For OLSR the optimal Z_R is the closest value between 1 and L to sqrt(a*N*mu*L/(tau+mu)). When N increases, this value will be likely be L. Therefore the optimal ZRP will act as a pure proactive scheme.

As N decreases, the optimal zone radius may be smaller than L, indicating that a hybrid approach outperforms purely proactive routing.

The important point to walk away with here is that the ZRP/DSDV and ZRP/OLSR curves aren't really shaped so differently. They both belong to the family of convex (U-shaped) curves. Given that OLSR is a very efficient proactive protocol, it seems reasonable to expect that *any* IARP implementation will produce a convex (U-shaped) traffic curve.

regards, Marc