# FLOODING

Subject: flooding
Date: Thu, 1 Jun 2000 18:19:00 -0400 (EDT)
From: Dmitri Deshun Perkins <perkin27@cse.msu.edu>
To: manet@itd.nrl.navy.mil

Can anyone comment on the relative performance of "flooding" as a routing algorithm in ad hoc networks? Have any studies been done that discuss when(under what conditions) it might be advantageous to simply use flooding(at least for control pkts)?

Please point me to any papers that discuss this topic. I briefly searched the archives and did not locate a discussion directly related to the above questions. Please excuse me if this topic has already been discussed. Feel free to point me to the correct archives.

Subject: Re: flooding
Date: Thu, 01 Jun 2000 19:10:22 -0400
From: Joe Macker <macker@itd.nrl.navy.mil>
To: Dmitri Deshun Perkins <perkin27@cse.msu.edu>, manet@itd.nrl.navy.mil

For what its worth to you, a flooding performance comparison was included in early simulation of ILS (ideal link state) and TORA. The intent was to qualitatively see how well flooding was doing in comparison to the other algorithms.

http://tonnant.itd.nrl.navy.mil/tora/tora_sim.html

As far as control packets, many algorithms do use forms of broadcast forwarding for some control packets, queries ,etc (also controlled or scoped flooding algorithms have been proposed..even for data ...(see the DSR ID) ..others?)

Subject: Re: flooding
Date: Fri, 02 Jun 2000 12:05:00 +0200
From: Anis LAOUITI <anis.laouiti@inria.fr>
Organization: INRIA
To: Dmitri Deshun Perkins <perkin27@cse.msu.edu>, manet@itd.nrl.navy.mil

You can take a look to this Research Report
"Multipoint relaying : An efficient technique for flooding in mobile wireless networks"
ftp://ftp.inria.fr/INRIA/publication/publi-ps-gz/RR/RR-3898.ps.gz

This technique is used in OLSR protocol (Optimzed Link State Routing) to flood control information in the network by reducing the number of retransmissions while forwarding a broadcast packet.

http://www.ietf.org/internet-drafts/draft-ietf-manet-olsr-01.txt

Subject: RE: flooding
Date: Fri, 2 Jun 2000 10:06:54 -0400
From: Alan Amis <AAmis@telogy.com>
To: "'Dmitri Deshun Perkins'" <perkin27@cse.msu.edu>, manet@itd.nrl.navy.mil

Hi Dmitri,

Based on previous responses I would assume that you are considering flooding as a means to actually route the control packets. However, if you are looking at flooding to provide the underlying infrastructure for routers in ad hoc networks I can suggest you take a look at a heuristic described in "Max-Min D-Closure Formations in Wireless Ad-Hoc Networks". This paper can be found in the proceeding of INFOCOM 2000. The main juice of this paper is that leaders (or effectively routers) are determined with no node in the network more than D hops away from its leader. The very attractive part of this heuristic is that it runs in D message rounds, not a function of how many nodes are in the network.

Alan Amis

Subject: flooding(point of clarification)
Date: Fri, 2 Jun 2000 11:13:51 -0400 (EDT)
From: Dmitri Deshun Perkins <perkin27@cse.msu.edu>
To: Alan Amis <AAmis@telogy.com>
CC: manet@itd.nrl.navy.mil

Hi Alan,

Thanks for the reply. I just wanted to give a point of clarification to my previous e-mail(flooding). At this point, I am just exploring the issues(routing, security, etc) of ad hoc networks as a possible dissertation topic. As I discuss the routing issues with others(namely, members of my advisory committee and students that attend weekly research meetings), the first question(based only on intuition) is always, "What's wrong with simply using flooding?" I suspect that many of you must get this question as well.

Thus, I have been attempting to (quantitatively) answer this question via simulation and wanted to reference other work related to flooding in an ad hoc environment. I have read much of the info concerning routing in an ad hoc environment and was just curious to know if flooding had been considered/rejected as a possible routing solution and if so, why? Again, thanks for your comments.

Thanks, Perkins

Subject: Re: flooding(point of clarification)
Date: Fri, 02 Jun 2000 12:21:46 -0400
From: Kenneth Brayer <kb@mitre.org>
Organization: The MITRE Corporation
To: Dmitri Deshun Perkins <perkin27@cse.msu.edu>
CC: Alan Amis <AAmis@telogy.com>, manet@itd.nrl.navy.mil

Dimitri

In a fixed network with fixed topology and using shortest path routing you can calculate the traffic level that a network can carry, depending on data rates, error rates, capacities of equipment, etc. If you use flooding you send everything everywhere. That is you send many more packets than are necessary. That is you deny your customer the use of some of the networks capacity. In a MANET you are less efficient at using the network than in a fixed network because the topology is changing and you need to find the path to the packet's destination. Again if you use ONLY flooding you will minimize the total traffic handling capacity of the network. Some flooding may be necessary but it is advantageous to use as little as possible.

I was among the first researchers in this area in 1976-1981 before the words MANET and Ad-Hoc were coined. If you look at my routing protocol you will see I used a little flooding at startup but then tried to make the network look as much like fixed path routing as possible. Parts of my protocol have appeared in the internet protocols and in many of the Ad-Hoc protocols under study today.

A thesis using only flooding is not useful, in my opinion, as it will give a minimum capability network and is not new. You have to go back to Paul Baran's papers on Hot Potato Routing and work that followed it. I suggest you run the Science Citation Index on Paul Baran and find the papers that reference him. This should uncover papers on flooding. You might also contact Norm Abramson at U. of Hawaii. He invented the Aloha Net back in the 1970s and is probably knowledgeable on flooding papers.

A list of my Ad-Hoc papers follows (sorry there are no computer copies, my work preceded on line databases): The Proc. IEEE paper is a good place to get a view of my work.

Networking of HF Radio Transmission
K. Brayer, IEEE MILCOM'86, Monterey, CA, October 1986

Packet Switching for Mobile Earth Stations via Low-Orbit Satellite Network
K. Brayer, Proceedings of the IEEE, November 1984

Autonomous Adaptive Local Area Networking: Ring Communications via
Point-to-Point Implementation
K. Brayer, IEEE INFOCOM'84, San Francisco, CA, April 1984

An Adaptive Computer Communication Network Designed with Decentralized Control
K. Brayer, IEEE Communications Magazine, November 1983

Adaptive Networking of Variable Topology Satellite Networks
K. Brayer, Sixth International Conference on Digital Satellite Communications,
Phoenix, AZ, September 1983

Routing in a "mobile" network - fact or fantasy?
K. Brayer, Data Communications, August 1983

Implementation and Performance of Survivable Computer Communication with
Autonomous Decentralized Control
K. Brayer, IEEE Communications Magazine, July 1983

Geographically Distributed Control of a Microprocessor Communications Network
K. Brayer, IEEE Communications Conference, Boston, MA, June 1983

Implementing Computer Communications with OEM Microprocessors:
Survivable Network Routing System
K. Brayer, IEEE Globecom'82, Miami, FL, November 1982

Implementing Computer Communications with OEM Microprocessors:
Survivable Routing System Performance
K. Brayer, IEEE Globecom'82, Miami, FL, November 1982

Survivable Computer Communications Routing using Decentralized Control
K. Brayer, IEEE International Conference on Circuits and Computers,
New York, NY, September 1982

Simulation via Implementation with Applications in Computer Communication
K. Brayer, V.S. Lafleur, G.H. Simpson, 15th Simulation Symposium, Tampa,
FL, March 1982

Subject: Re: flooding(point of clarification)
Date: Fri, 02 Jun 2000 18:24:11 +0200
From: jacquet@menetou.inria.fr
To: Dmitri Deshun Perkins <perkin27@cse.msu.edu>, Alan Amis <AAmis@telogy.com>
CC: manet@itd.nrl.navy.mil

Clearly flooding data in the network is resource over-consuming. If your network has 100 nodes, then each data packet needs to be repeated 100 times, thus reducing your efficient bandwidth by a factor 100.

Philippe
PS: I also have problem on how the IP stack will support this.

Subject: Re: flooding(point of clarification)
Date: Fri, 02 Jun 2000 14:02:19 -0400
From: Ram Ramanathan <ramanath@bbn.com>
To: manet@itd.nrl.navy.mil

I wonder if it is appropriate to dismiss flooding so. Surely, whether flooding is viable in a network or not depends on several parameters including size, radio data rate, traffic load and mobility.

For instance, consider a 100 node network of 11 Mbps radios (COTS today). Suppose the only traffic were situational awareness traffic at 3 kbps per node. Assume further that channel access and header overhead etc. reduce the raw data rate of 11 Mbps by a factor of 20 (pessimistically), to yield an effective data rate per node of 550 kbps. Each node then runs comfortably at about 50% utilization. Flooding is an adequate solution.

Moreover, as has been pointed out previously on this working group and several papers, if the mobility is extremely high flooding may be the *only* way to deliver data.

I am not advocating flooding, or implying that one should tailor network design to low data rate traffic. I am merely pointing out that before one goes off and designs complex schemes, one should examine the viability bounds for the simplest possible solution for the problem -- and Dimitri is asking the right set of questions in this context. Indeed, questions that should have been asked (as Dimitri points out) and answered long ago, but no answers seem to be around? Further, as technology improves (e.g., the data rate), it might be worthwhile examining such brute force methods for limited scenarios.

Dimitri: Some recent work that may be relevant on this topic (although perhaps not answering the question directly)

Ho et al, "Flooding for Reliable Multicast in Ad hoc networks", DIAL-M Proceedings, 99

Ni et al, "The Broadcast Storm Problem...", Proc. Mobicom 99.

and citations therein.

Also check out NOAHnet -- this uses data flooding for wired nets.
Go to your favorite search engine (I used Google.com) and search for NOAHnet.

I would be interested in seeing your analysis of flooding when it is done!

-Ram.

Subject: Re: flooding
Date: Fri, 2 Jun 2000 11:10:06 -0700 (PDT)
From: Katia Obraczka <katia@isi.edu>
To: perkin27@cse.msu.edu
CC: manet@itd.nrl.navy.mil

Dmitri,

We have done a comparative study of flooding in the context of multicast routing in MANETs. A paper with our preliminary results appeared in the DIALM99 workshop (soft copy available from http://www.isi.edu:80/people/katia/dialm.ps.gz). We have just written a report on more recent results we got when comparing flooding with other ad hoc multicast routing protocols. If you are interested, I can send you a pointer to that as well. There is also a study by SJ Lee et al. from UCLA comparing their multicast routing protocol ODMRP with other ad hoc multicast protocols, including flooding. If interested I can send you a pointer to that as well.

Katia

Subject: Re: flooding(point of clarification)
Date: Fri, 02 Jun 2000 14:36:17 -0400
From: Kenneth Brayer <kb@mitre.org>
Organization: The MITRE Corporation
To: Ram Ramanathan <ramanath@bbn.com>
CC: manet@itd.nrl.navy.mil

Ram
Generally some flooding is needed at some time and the more fragmented the net the more you might need. Your example is perfectly fine but is that the only traffic pattern. In general my customer wants a wide variety of traffic in a large complex net so if I can find an approach which preserves capacity for the user that is what I will do. Thus, I seek to minimize flooding, overhead messages, connectivity table distribution etc. There is a large body of literature on flooding but it is more than 30 years old and exists only in hard copy in libraries. I think the IEEE literature search engines available through most libraries will find those citations.

Ken Brayer
The MITRE Corporation

Subject: Re: flooding
Date: Sat, 3 Jun 2000 14:16:20 +0800
From: "w.peng" <wpeng@nudt.edu.cn>
To: "Dmitri Deshun Perkins" <perkin27@cse.msu.edu>
CC: <manet@itd.nrl.navy.mil>

Dmitri,

I've done some research on the broadcast problem in mobile ad hoc networks. I've developed two approaches to reduce the broadcast redundancy and compared their performance with flooding. If you are interested in my work, you can take a look to my three papers:

1) "Efficient Broadcast in Mobile Ad Hoc Networks Using Connected Dominating Sets". This paper has been accepted for publication in ICPADS'2000.

2) "AHBP: An Efficient Broadcast Protocol for Mobile Ad Hoc Networks". accepted by Journal of Computer Science And Technology (publicated in Beijing, China).

3) "On the Reduction of Broadcast Redundancy in Mobile Ad Hoc Networks". will appear as a poster in MobiHOC'2000.

Because my homepage is under construction, so write to me if you want a copy of my paper.

Wei Peng
Department of Computer Science,
Changsha Institute of Technology,
Changsha, Hunan, 410073, P.R.China
Email: wpeng@nudt.edu.cn

Subject: Re: flooding
Date: Sat, 03 Jun 2000 12:03:50 +0100
From: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
Organization: University of Surrey, Guildford, England
To: perkin27@cse.msu.edu, manet <manet@itd.nrl.navy.mil>

Dimitri, you may want to have a look at our RDMAR protocol where localisation of Route Discovery as well as of Route Repair is achieved, thus reducing significantly the overhead of flooding the entire network area.

See my homepage for some RDMAR papers
http://www.ee.surrey.ac.uk/Personal/G.Aggelou/publications.html or the IETF MANET WG homepage for the IETF draft (http://www.ietf.org/html.charters/manet-charter.html).

Regards, George N. A.

Subject: Re: flooding(point of clarification)
Date: Sat, 03 Jun 2000 15:08:33 +0100
From: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
Organization: University of Surrey, Guildford, England
To: Ram Ramanathan <ramanath@bbn.com>, manet <manet@itd.nrl.navy.mil>

Ram a few comments on your thoughts:

> For instance, consider a 100 node network of 11 Mbps radios (COTS today).
> Suppose the only traffic were situational awareness traffic at 3 kbps per
> node. Assume further that channel access and header overhead etc. reduce the

> raw data rate of 11 Mbps by a factor of 20 (pessimistically), to yield an
> effective data rate per node of 550 kbps. Each node then runs comfortably
> at about 50% utilization. Flooding is an adequate solution.

1) If, according to your assumptions, each node runs at about 50% of its nominated data rate, don't you think this is a major reason to avoid flooding then? Do you believe 50% effective utilisations is a good performanmce indication for protocol design?

2) Uncontrolled propagation of control messages implies that:

a. some areas of the network topology are unecessarily disturbed. Hence, even if we accept the 50% effective utilisation for the set of nodes that must propagate this control messaging, say during a Route Discovery process, for the successful discovery of a route, then, however, this 50% is not an acceptable figure for the rest of nodes which even if they don't propagate the route discovery signalling, the discovery will be successful. For example, by localising the route discovery signalling such that only the nodes between the source and destination of the discovery are disturbed, then certainly the rest of the network nodes will enjoy much better channel utilisation rates.

b. (adding to a.), in CDMA-based networks where interference levels increase as a function of the user activity (and ), the extra control messaging will increase user interference. Therefore, system capacity will be reduced.

c. the probability of collisions at the channel access layer also increases. This in turn implies more retransmissions and the such.... Which in turn implies more delay on per packet forwarding.

By any means, your scenario does not show any tradeoffs as what the benefits of using flooding would be _compared_ to when flooding is not used. Therefore, I would say that your conclusion "Flooding is an adequate solution" is not true.

Regards, George N. A.

Subject: Re: flooding(point of clarification)
Date: Sat, 03 Jun 2000 22:22:57 -0400
From: Ram Ramanathan <ramanath@bbn.com>
To: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
CC: Ram Ramanathan <ramanath@bbn.com>, manet <manet@itd.nrl.navy.mil>

George,

Your repeated reference to "control messages" tells me you have completely misunderstood me. I was talking about flooding every DATA packet to everywhere in the network. There ARE NO control messages in a "flooding protocol". As I mentioned in my previous note, flooding is clearly not what I would want to use as my routing mechanism all the time. However, understanding in what regions of the n-dimensional space of (datarate, load, mobility, size ..) flooding works, and in what regions it breaks would be valuable -- the example was one point in such a space, to make my point (no pun:)

Again, what I suggest is that data flooding might be competitive when 1) mobility is so high that no routing protocol can converge fast enough 2) when offered_load/data_rate is sufficiently low -- this could happen because load is low or as data_rate increases (as is the industry trend)

It seems like you've missed the example too, so here are some inline comments.

> > For instance, consider a 100 node network of 11 Mbps radios (COTS today).
> > Suppose the only traffic were situational awareness traffic at 3 kbps per
> > node. Assume further that channel access and header overhead etc. reduce the
> > raw data rate of 11 Mbps by a factor of 20 (pessimistically), to yield an
> > effective data rate per node of 550 kbps. Each node then runs comfortably
> > at about 50% utilization. Flooding is an adequate solution.
>
> 1) If, according to your assumptions, each node runs at about 50% of its
> nominated data rate, don't you think this is a major reason to avoid
> flooding then? Do you believe 50% effective utilisations is a good
> performanmce indication for protocol design?

Huh? Are you saying that 50% is too much or that it is too little? If you are saying 50% is too much, i.e., people should run a link at less than this, consider that most ISPs use 50% as a conservative, safe number.

If you are saying 50% is too little, you have missed this completely. I use utilization, informally, as load/capacity. In the example, the node effective capacity is 550 kbps. The load is 3kbps*100 (since every packet is sent exactly once (using sequence numbers to purge duplicates) by every node in the network) or 300 kbps. Thus, the utilization is 54.5%.

If you want to run it at a higher utilization, sure, that is fine, it supports even higher traffic. For instance, at 90% utilization, it supports 5 kbps at every node. In other words, if you think I shuld be saying 90% instead of 50%, just replace 3kbps by 5kbps in the exampel.
>
> 2) Uncontrolled propagation of control messages implies that:
>
What control messages (read above)?!

I'll ignore the points a,b,c since it seems like you're talking about the benefits of "constrained flooding" of routing control messages - a noble thought, but largely irrelevant to the discussion.

> a. some areas of the network topology are unecessarily disturbed.
> Hence, even if we accept the 50% effective utilisation for the set of
> nodes that must propagate this control messaging, say during a Route
> Discovery process, for the successful discovery of a route, then,
> however, this 50% is not an acceptable figure for the rest of nodes
> which even if they don't propagate the route discovery signalling, the
> discovery will be successful. For example, by localising the route
> discovery signalling such that only the nodes between the source and
> destination of the discovery are disturbed, then certainly the rest of
> the network nodes will enjoy much better channel utilisation rates.
>
> b. (adding to a.), in CDMA-based networks where interference levels
> increase as a function of the user activity (and ), the extra control
> messaging will increase user interference. Therefore, system capacity
> will be reduced.
>
> c. the probability of collisions at the channel access layer also
> increases. This in turn implies more retransmissions and the such....

> Which in turn implies more delay on per packet forwarding.
>
>
> By any means, your scenario does not show any tradeoffs as what the

Didn't say I was showing tradeoffs....

> benefits of using flooding would be _compared_ to when floding is not
> used.
> Therefore, I would say that your conclusion "Flooding is an adequate
> solution" is not true.

Of course it is an adequate solution for the particular scenario. Of course it is not in general.

Cheers,-Ram.

Subject: Re: flooding(point of clarification)
Date: Sun, 04 Jun 2000 11:53:07 +0100
From: "George N. Aggelou" <g.aggelou@eim.surrey.ac.uk>
Organization: University of Surrey, Guildford, England
To: Ram Ramanathan <ramanath@bbn.com>, manet <manet@itd.nrl.navy.mil>

Ram,

My previous comments address some problems of flooding "control signalling" in a MANET as your wording "traffic" in your first message was interpreted as "routing-related control + data traffic".

Apologies for the confusion then.

Best, George N.A.

Subject: Re: flooding(point of clarification)
Date: Sun, 4 Jun 2000 08:25:09 -0500 (CDT)
From: Nitin H Vaidya <vaidya@cs.tamu.edu>
To: ramanath@bbn.com
CC: manet@itd.nrl.navy.mil

Ram:

>> From: Ram Ramanathan <ramanath@bbn.com>
...
>> 2) when offered_load/data_rate is sufficiently low -- this could happen
>> because load is low or as data_rate increases (as is the industry trend)

I agree with the first part of your statement. With sufficiently low load, the cost of finding/maintaining routes could be higher than simply flooding a packet, so flooding could be more effcient.

I am not so sure about the second part of your argument ... Correct me if I am mistaken, but I guess what you are saying is that if there is plenty of bandwidth available, the no harm in using it inefficiently. This may not always make sense. For instance, what is bandwidth is abundant, but energy supply isn't ?

Just nitpicking -:)

- nitin

Subject: Re: flooding
Date: Sun, 04 Jun 2000 17:30:25 -0700
From: Richard <ogier@pit.erg.sri.com>
To: Dmitri Deshun Perkins <perkin27@cse.msu.edu>
CC: manet@itd.nrl.navy.mil

> Can anyone comment on the relative performance of "flooding"
> as a routing algorithm in ad hoc networks? Have any studies
> been done that discuss when(under what conditions) it might
> be advantageous to simply use flooding(at least for control
> pkts)?
>
> Please point me to any papers that discuss this topic.
>

A simulation study comparing flooding (for disseminating link-state updates) to TBRPF (Topology Broadcast based on Reverse Path Forwarding) and to some other routing protocols was presented in the following paper:

R.G. Ogier, ``Efficient Routing Protocols for Packet-Radio Networks Based on Tree Sharing,'' Proc. Sixth IEEE Intl. Workshop on Mobile Multimedia Communications (MOMUC'99), November 1999.

TBRPF was originally presented in the following paper, but the above paper used more accurate simulation models.

B. Bellur and R. G. Ogier, ``A Reliable, Efficient Topology Broadcast Protocol for Dynamic Networks,'' Proc. IEEE INFOCOM '99, New York, March 1 999.

These papers can be found at http://www.erg.sri.com/publications.html

TBRPF and flooding both provide each node with the state of every link in the network (or within a cluster if hierarchical routing is used), but TBRPF does it more efficiently by sending each update along a min-hop tree rooted at the source of the update. (Thus, for example, leaves of the tree need not forward updates.) TBRPF dynamically updates these broadcast trees using the link-state information that is received along the trees.

The simulation results in the MOMUC'99 paper show that TBRPF generates between 36% and 85% less update/control traffic than flooding, depending on the density of the network topology and on whether unicast or broadcast link-level transmissions are used.

The MOMUC paper also presents a protocol called FTSP (Full Tree-Sharing Protocol), which is very similar to the ORA (optimal routing) version of STAR. The simulation results also show that TBRPF generates between 14% and 58% less update/control traffic than FTSP, despite the fact that FTSP provides each node with only partial topology information. (TBRPF has not yet been compared to other MANET protocols such as TORA, DSR, and AODV.)

Richard Ogier

Subject: Re: flooding(point of clarification)
Date: Mon, 05 Jun 2000 11:55:05 -0400
From: Ram Ramanathan <ramanath@bbn.com>
To: Nitin H Vaidya <vaidya@cs.tamu.edu>
CC: ramanath@bbn.com, manet@itd.nrl.navy.mil

> ...
> >> 2) when offered_load/data_rate is sufficiently low -- this could happen
> >> because load is low or as data_rate increases (as is the industry trend)
>
> I agree with the first part of your statement.
> With sufficiently low load, the cost of finding/maintaining
> routes could be higher than simply flooding a packet,
> so flooding could be more effcient.
>
> I am not so sure about the second part of your
> argument ... Correct me if I am mistaken, but I guess what you
> are saying is that if there is plenty of bandwidth available, the
> no harm in using it inefficiently. This may not always make sense.

More precisely, I was saying that for a given set of requirements/constraints I'd pick the simplest solution, and that flooding may be the simplest solution in some cases, e.g, when you need only to support low load and you have high bandwidth (capacity).

> For instance, what is bandwidth is abundant, but energy
> supply isn't ?

Yes, of course, energy may or may not be an issue. This is but one of several resource constraints (e.g. buffers may be another) that have to be met.

-Ram.

Subject: flooding (Epidemic Routing)
Date: Mon, 05 Jun 2000 12:07:21 -0400
From: Amin Vahdat <vahdat@cs.duke.edu>
Organization: Duke University
To: manet@itd.nrl.navy.mil

In the context of flooding protocols in ad hoc networks, we recently conducted a study on the effectiveness of "controlled" flooding for partially connected ad hoc networks. In the paper, we outline a number of scenarios where a connected path from source to destination is potentially never available, making traditional ad hoc routing protocols inappropriate. The idea is to use epidemic algorithms and logical time vectors to exchange unseen messages among connected neighbors, and using eventual connectivity among "connected" islands to deliver messages to their destination. For the set of experiments we conducted, our algorithms show excellent delivery rates, average delivery times that scale with how "unconnected" the network is, and moderate per-node resource consumption.

One direction we are pursuing is the applicability of the algorithms to multicast in ad hoc networks.

The paper is available at:

http://www.cs.duke.edu/~vahdat/ps/epidemic.pdf

The title and abstract are included below. I would appreciate any comments on the paper.

--
Amin Vahdat
Assistant Professor
Duke University
http://www.cs.duke.edu/~vahdat

Title: Epidemic Routing for Partially-Connected Ad Hoc Networks

Abstract:

Mobile ad hoc routing protocols allow nodes with wireless adaptors to communicate with one another without any pre-existing network infrastructure. Existing ad hoc routing protocols, while robust to rapidly changing network topology, assume the presence of a connected path from source to destination. Given power limitations, the advent of short-range wireless networks, and the wide physical conditions over which ad hoc networks must be deployed, in some scenarios it is likely that this assumption is invalid. In this work, we develop techniques to deliver messages in the case where there is never a connected path from source to destination or when a network partition exists at the time a message is originated. To this end, we introduce Epidemic Routing, where random pair-wise exchanges of messages among mobile hosts ensure eventual message delivery. The goals of Epidemic Routing are to: i) maximize message delivery rate, ii) minimize message latency, and iii) minimize the total resources consumed in message delivery. Through an implementation in the Monarch simulator, we show that Epidemic Routing achieves eventual delivery of 100% of messages with reasonable aggregate resource consumption in a number of interesting scenarios.

Subject: Re: flooding
Date: Mon, 5 Jun 2000 15:03:30 -0400 (EDT)
From: "A. Bruce McDonald" <tudball@lis.pitt.edu>
To: manet@itd.nrl.navy.mil

I think that most people would agree that even the *best* routing algorithm will eventually break if the topology changes rapidly enough (it would be interesting to learn what the *fundamental* limitations are---if we could characterize such a thing). I would hasten to add that even the simplest of all routing algorithms, flooding, will itself break down without considerable control. This seems to be an often overlooked reality of routing under extreme conditions. Flooding reliably in the best of circumstances can be difficult. Hence, even flooding may not be as simple as we'd like it to be under these circumstances.

--Bruce McDonald, University of Pittsburgh