Network Working Group
Internet-Draft
Expires: September 15, 2002

F. Baker
Cisco Systems

March 17, 2002

An outsider's view of MANET
draft-baker-manet-review-01

Status of this Memo

Copyright Notice

Abstract

This note addresses routing in chaotic non-engineered radio networks.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [2].

1. Overview and disclaimer

This note addresses routing in chaotic non-engineered radio networks. The "chaos" in these networks derives from a combination of device motion and interactions with the environment.

Wireless links are quite susceptible to time varying statistical behavior caused by many factors, including the physics of the propagation medium, inner city fading characteristics, shadowing (e.g., a person walking by a device), potential power control, etc induce effects that need addressing even in pseudo-static scenarios.

Examples of such networks vary from webs of radio-linked sensors distributed like seed by air drop, to the behavior of satellites in random orbits, to automotive applications in which cars and traffic lights are communicating nodes, to military applications such as battlefield communications among soldiers, unmanned reconnaissance platforms, vehicle-mounted devices, fixed bases, and field encampments.

Such networks have been the subject of significant research over the past several years, with numerous routing proposals, and offers to re-engineer TCP to make applications operate well in the network. The fundamental bent of this note, however, differs from this research in intent. Mobile ad hoc networks, or manets, are not seen as networks in their own right any more than local area networks are networks in their own right. Instead, manets are seen as localities within networks, much as LANs operate as the local access to a wider area Internet. The operation of manets in isolation is a special case of their operation as part of a larger network.

Taken from this perspective, the important question is not so much "please design a routing protocol that will be useful in a manet", as it is "please design a routing protocol that will be useful in a network that contains one or more manets".

2. IETF History and Work: the MANET Working Group

The MANET working group was chartered in 1997, to discuss and develop solutions for what were described as Mobile Ad Hoc Networks. Although it was chartered as an engineering group, one could argue that it was then and is now a research organization. There has been little if any commercial activity related to this type of network; activity has been focused in the research divisions of various companies, notably Nokia, Inria, SRI, Intel, and others, and with academic institutions such as UCSB, Rice, and so on.

2.1 Problem Statement

The problem statement that the MANET working group was given, which may be found at http://www.ietf.org/html.charters/manet-charter.html, says:

> A "mobile ad hoc network" (MANET) is an autonomous system of mobile routers (and associated hosts) connected by wireless links--the union of which form an arbitrary graph. The routers are free to move randomly and organize themselves arbitrarily; thus, the network's wireless topology may change rapidly and unpredictably. Such a network may operate in a standalone fashion, or may be connected to the larger Internet.
>
> The primary focus of the working group is to develop and evolve MANET routing specification(s) and introduce them to the Internet Standards track. The goal is to support networks scaling up to hundreds of routers. If this proves successful, future work may include development of other protocols to support additional routing functionality. The working group will also serve as a meeting place and forum for those developing and experimenting with MANET approaches.
>
> The working group will examine related security issues around MANET. It will consider the intended usage environments, and the threats that are (or are not) meaningful within that environment.

In general, a MANET network is very similar to any other Internet technology; one researcher, in a discussion of how to manage low signal-to-noise ratio channels, ruefully remarked that the researchers in the area frequently find themselves re-inventing wheels. Where it differs from standard routing, however, is the structure and characteristics of a low-power radio network.

As an example of the kind of interesting radio environment that can exist, consider the sad case of Alice, Bob, and Carol in figure 1. An environmental obstacle separates Alice and Bob, so their radios cannot "hear" each other, but Carol can "hear" them both quite well, as long as they do not happen to "speak" at the same time. Carol can interconnect them by repeating their messages; they might also be able to correct the problem by taking a few steps or lifting their radios - anything that would obviate the obstacle. Clearly, these devices are in close physical proximity, but their views of the network are very different, and their ability to use it differ markedly as well. As they move, or as their environment changes around them, this view of the network will also change - often appearing to change randomly.

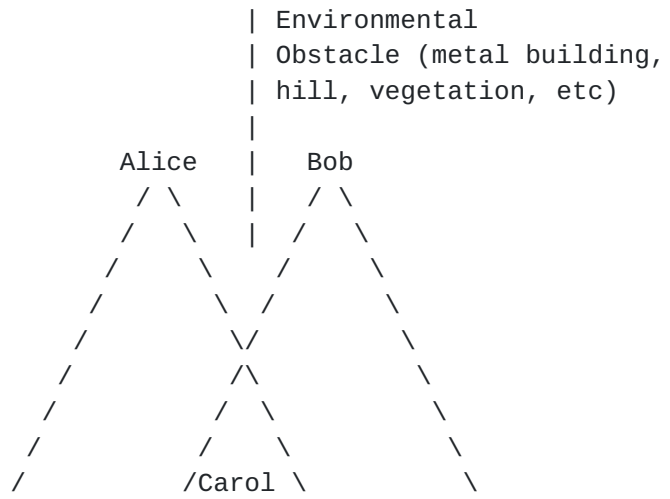

Figure 1: Varying views in a radio network

2.1.1 Neighbor sets

Unlike wired networks, each device in a radio network has a slightly different view of its world. From the router's perspective, a LAN is an essentially fixed set of routing neighbors, which changes only on administrative action, with additional end systems, which may come and go. It is therefore rational and desirable to have the routers elect one among them to perform coordination tasks - what is called a "designated router" in OSPF and IS-IS. In a MANET, however, any system might be called upon to relay traffic for others.

Signal quality between a wireless transmission source and a receiver, usually measured as a signal-to-noise ratio, can be a difficult to model and quantify. Although there are simple propagation models for ideal conditions that are represented by deterministic equations (e.g., free space ~1/r^2 and log distance path models ~1/r^n), the

complex mechanism of electromagnetic wave propagation can be attributed to reflection, diffraction, and scattering effects. Many mobile radio systems will operate in urban areas or within buildings where time varying signal fading and shadowing effects may naturally occur. Thus, real world propagation effects often result in a time varying function for received signal power from a source. These real world physics effects make signal strength prediction and short-term estimation of link quality more complex.

Since systems are in different locations, each system may have a different set of neighbors that it is able to communicate effectively with, which overlay each other haphazardly. For this reason, the rules that allow OSPF to reduce its flooding statistics from an exponential to linear behavior by electing a designated router to perform the job are unusable in a radio network.

2.1.2 Random Interconnection Topology

Another issue is the aspect of mobility, which is different from what has traditionally been termed "IP mobility". The concept in IP Mobility is that a device has a normal home in some topological part of a stable network, as indicated by its address, but may temporarily move somewhere else. That "address" then becomes something more like a name. A home agent translates it into a second address, which represents the device's current actual topological location, and the packet stream is forwarded there. The device may then advise its correspondent of its current topological "care-of" address, facilitating more direct routing. In a MANET, the address is tied to the device, not a topological location, as there is no fixed network infrastructure. When the addressed device moves, therefore, the motion changes the routing infrastructure. There is no question of the correspondent transmitting to the new care-of address, or of a home agent forwarding traffic from "the right place" to "somewhere else" along a dog-leg path; standard routing will get the packets there as a direct outcome as routing tracks the motion of the device.

Mobility is not an aspect of all MANETs; some varieties of sensor networks (such as forest fire sensors scattered by airdrop in the region of a fire) can be expected to be stationary once deployed. However, even in this case, topological relationships are arbitrary and unengineered. In applications where node mobility is in view, it can be haphazard, and in extreme cases can result in entire networks randomly partitioning and joining together.

The fundamental behavior of a manet is that a routing node carries with it an address or address prefix, and when it moves, it moves the actual address. When this happens, routing must be recalculated in accordance with the new topology.

This has ramifications for such normal behaviors as autoconfiguration of address prefixes and router IDs, which can be replicated in separate networks and will require resolution when they join. It also has ramifications for movement among what OSPF or IS-IS would call "areas"; if an address is "known" to be someplace and suddenly pops up somewhere else, it will need to change areas as well.

IP Mobility solves an issue in addressing caused by temporary mobility; MANET routing solves a routing problem in a network where mobility is normal. When mobility is solved using routing, addressing-based solutions are irrelevant.

2.1.3 Radio issues

The IEEE 802.11 radio networks that typically connect research manets have all of the radios on the same frequency, using a Carrier Sense Multiple Access (CSMA) discipline. In other words, if the receiver can tell that someone else is transmitting, it may attempt to not interrupt, but there is no guarantee that it will be able to sense collisions. In such cases, since all radios use the same frequency and spread spectrum patterns, the transmitters effectively jam each other.

One could imagine solving that using disciplines similar to that used in LDDI, wherein each system has a sequence number and transmissions are made in that order, to generate a form of Slotted ALOHA. What many of the MANET routing protocols are doing, though, is finding reasons to not have correlated reasons to transmit, such as acknowledging multicast messages, and relying then on randomization to evade the problem. Past research on such approaches suggests that it is helpful, but introduces complexities of its own as well.

There are, of course, other varieties of radios. Military radios may use Code Division Multiple Access (Spread Spectrum CDMA) or Time Division Multiplexing (TDMA) disciplines.

2.1.4 Convergence Requirements

Internet routing protocols, such as RIP, OSPF, IS-IS, and BGP, have always been developed on the assumption that networks proceed from a converged state to a converged state through epochal transitions such as changes to router configurations, loss or restoration of links, or loss or restoration of routers. For this reason, instability in networks is viewed with some alarm. OSPF and IS-IS were developed in large part because it was increasingly observed that existing distance-vector IGPs displayed unacceptably long convergence intervals or were not sufficiently resilient. The increased expressiveness of what were then called "variable length subnets",

and are now called Classless Inter-Domain Routing (CIDR) was also a significant factor. At this time, concern is raised in many quarters because the BGP4+ backbone displays significant instability and long convergence intervals.

MANET networks display exactly the opposite characteristic: due to node mobility and constantly changing neighbor interconnectivity, the network may display episodes in which it converges, but normally is in a state of flux. The question becomes what level of convergence is required: is it worthwhile to expend a great deal of effort to attempt to maintain a higher level of convergence, or is it better to accept partial convergence? The answer to that is not obvious, and most likely varies from network to network and application to application.

2.1.5 Unidirectional Routing

Manet research has blown hot and cold on the matter of directional routing.

Common routing protocols depend on bidirectional connectivity. Distance Vector protocols, for example, advertise what might be considered to be statements that "you can reach [this prefix] with [these attributes] via me, on the interface that you receive this message on". OSPF and IS-IS, while not making statements of that form, explicitly refuse to use links that lack bidirectional connectivity. They refuse to neighbor, and the SPF implementation checks that the far end of a link is reporting bidirectional linkage before accepting the extension of a route.

In a manet network, as previously discussed, a given relationship can be unidirectional. System A may be able to "hear" system B, but B not "hear" A, and it may make operational sense to allow A->B to be used as a message forwarding link. Few if any published protocols do this today, but it is raised as a desirable capability in some discussions.

The operational and protocol issues are immense. For a solution to support unidirectional links, either the sender on such a link must be sending messages for which it cannot determine whether any given target receives them, or it must have another path (perhaps also unidirectional) via which it receives routing information that tells it of its hearing neighbor. This fact limits the classes of protocols that can be used to deploy such a network, and applications that will find such a network useful. Operationally, the fact that a link is bidirectional is often the only way a system can know it is working at all.

2.1.6 Solution Approaches

There are a number of ways to solve these issues, as the number of proposals made to the MANET working group attests. They are commonly broken down into two broad classes: reactive protocols, which determine what route to use when the route is needed, and proactive protocols, that predetermine routes on the assumption that they may be needed.

Reactive protocols follow approaches such as source routing or some form of routing on demand. These are designed with the premises:

- Network locality is strong: most active routes are topologically local, within one or two hops.
- The application can work around occasional routing glitches if recovery is expedited
- While routing may change continuously in the global sense, individual routes generally survive long enough to perform common application tasks.
- The overhead of searching for a route when it is needed (which may take several round trip times) is acceptable.
- The ratio of multi-hop routes actively being discovered and maintained is small compared to the number of such possible routes within a manet area.
- Route exploration surges, which result from the movement of "keystone" nodes, are at an acceptable level.

If one accepts these premises, then it is reasonable to assume that one will search for paths when they are needed, and save them either in the source system or the intervening nodes.

Proactive protocols generally follow some form of link-state algorithm, such as SPF (Dijkstra) or of map-based explicit routing. These are designed with the premises:

- Network locality is indeterminate; routes of any length may be commonly used, or not at all.
- The application can work around occasional routing glitches, but recovery must be almost immediate.
- The constant route changes that happen globally may materially affect the correct operation of individual nodes.

- o The overhead of calculation and information flooding is acceptable, but the overhead of searching is not.

It is possible to mix the two models as well; a link state database could be maintained through the network, but inspected only when it changes the routing behavior of a network node known to be relevant to a route that is currently in use or a new route is needed.

2.2 Progress of the group

Since 1997, at least ten protocols have been proposed. These fall into several categories. Dynamic Source Routing (DSR) is similar in many respects to IEEE 802.5 Source Route Bridging. Ad hoc On-Demand Distance Vector (AODV) is a reactive protocol that introduces routing state in a network only when needed. Topology Multicast Reverse Path Forwarding (TBRPF) and Optimized Link State Routing (OLSR) are SPF-based protocols, which may be compared to OSPF or IS-IS. They differ from these in operational detail, and in the way they flood routing database information.

Security is an issue that none of these protocols has directly addressed, although some general analyses have been floated in the working group. Security flaws exist in many of them, which could be exploited; for example, DSR is subject to man-in-the-middle attacks, and according to one of the authors has experienced them (in the form of a lack of routing robustness when stations move) in field-testing.

Similarly, scaling is an issue that has been dealt with only on the surface. The stated goal of these protocols is "scaling up to hundreds of routers"; whether or not the features that allow this level of scaling will in turn enable scaling to thousands or tens of thousands of routers remains to be shown. The difference between proactive and reactive protocols is intended to address some issues in scaling, with different trade-offs. A reactive protocol might be appropriate in a network where most connectivity is local and non-local routes tend to remain fairly stable for the duration of a typical session; the router maintains no state that is not in current use, and is willing to perform an expensive set-up when it needs non-local routing state. A proactive protocol might be appropriate in network in which non-local communications are normal and route maintenance must be rapid. The trade-off is that in a proactive protocol, topological turbulence causes nodes to constantly store, propagate, and adjust routing information that has no current utility.

Quality of Service (important for voice applications) is also not addressed, except in AODV. There is a draft that describes QoS use of the routing protocol, which would have it seek a path in which

certain bandwidth and delay bounds are met, and in which the request for a route would fail if its conditions cannot be satisfied. QoS routing is, of course, seen as a research topic by much of the IETF community, due to a lack of commercial demand and the difficulty of the problem in a destination-routed connectionless network [6][10].

Distributed address autoconfiguration in manet is non-trivial due to the need for multi-hop DAD algorithms. There have been discussions with zeroconf participants to explore possibilities and issues.

While extending any of these protocols to IPv6 is straightforward, in publicly available documents, only AODV has materially addressed IPv6. There are drafts on stateless autoconfiguration of IPv6 networks in a MANET, but it is independent of the routing protocol, and apply to non-routing hosts which neighbor with routers, rather than to systems capable of forwarding packets. OLSR mentions how it could be extended to address IPv6. Likewise, TBRPF states (section 9.7.2) that

> Transition mechanisms described in the IETF NGTRANS working group (e.g. ISATAP) enable IPv6 operation over IPv4 routing infrastructures. ISATAP [19] can be used on TBRPF MANETs to enable automatic IPv6-in-IPv4 operation regardless of route changes due to mobility. Future versions of this draft will specify a native IPv6 capability for TBRPF using mechanisms similar to those specified in [21]. Packet formats which implement such mechanisms will use 4-octet router ID's instead of 16-octet IPv6 addresses for greater efficiency.

DHCP is not mentioned in any posted draft, although there are an argument that some form of address assignment protocol adapted to MANET networks is required. IP Mobility is not addressed either.

An initial requirements document has been published, as RFC 2501 [7]. DSR and AODV have been proposed to the IESG for publication as Experimental RFCs. No other drafts have been sent to the IESG.

2.3 Probable directions

The Working Group expects to publish several protocols as Experimental, including DSR and AODV, but expect to take one reactive and one proactive protocol onto the IETF standards track. These will likely be AODV and OLSR.

In 1997, the working group chairs asked the authors of OLSR to publish their work in the IETF context (although one of the working

group chairs is the author of a competing proposal), because they considered it a well architected solution to the problem. Although some details remain to be worked out, they still consider it among the better proposals on the table.

AODV is likewise a quite workable solution, with an interoperability test of as many as fifteen academic implementations scheduled in March 2002. It alone, of the candidates, addresses IPv6 or Quality of Service issues.

TBRPF is interesting, and should work correctly, although the operational utility of some of its optimizations may be open to question in a given network. SRI has aggressively marketed TBRPF and its IPRs to the working group. The fact that a patent has been applied for on certain aspects is, however, severely limiting politically. If there is any way in which the IETF is absolutely predictable, it is that when confronted with a choice between a proposal encumbered with IPR issues and an unencumbered proposal, it will choose the unencumbered one.

The other proposals are either not as far along, have encountered problems, or have less traction in the working group.

3. Market issues

From the perspective of the marketplace, at this time there is little commercial demand for MANET-style protocols. This is not an issue in the protocols themselves; it is an issue of the applications in which they might be used. While interactive automotive mapping services are common in Japan and some European countries, these use direct-connect short-reach radio technologies or third generation wireless, rather than packet networks. Sensor networks remain the realm of research, and military uses are in research. As a result, not only are we limited by the lack of standards, but by a distinct lack of market interest.

In fairness, when IPv4 was first deployed, there was little commercial demand for it, either. Until the mid-1990's, Novell Netware and Apple AppleTalk had more commercial penetration and offered superior application features, and IBM SNA absolutely controlled the financial services sector. At this point, while few dispute that IPv6 gives more addresses, and therefore is a good solution to certain market issues, there is significant dispute that markets demand IPv6 deployment. Dismissing MANET-style routing as meeting with little market demand is at best shortsighted. Rather, since it meets some market requirements, the most sensible approach is to develop the capability experimentally and see if markets grow to depend on it.

4. Protocol Proposals

For completeness, I will now discuss various possible approaches to MANET routing as described in some of the leading protocols. I will first discuss the use of OSPF Version 2 [4]as a MANET protocol, which has both issues and opportunities. I will also discuss the proposals that I perceive to be leading in the MANET working group, for comparison.

4.1 OSPF Version 2 and IS-IS

From the perspective of a commercial vendor, the most obvious routing protocol to use for any application is one that is already implemented. For this reason, companies like Cisco and many of their customers would likely look first to such standard protocols as OSPF, IS-IS, or possibly proprietary protocols such as EIGRP, before assuming a special purpose protocol is required. It also serves as a point of perspective, defining terms and surfacing issues, which the remaining discussion may refer to. If a workable scenario can be found for OSPF or IS-IS in a MANET, interoperation with wired internet components, including wired networks within vehicles, wired bases, and the internet proper, becomes within the grasp of a MANET network, which is one of MANET's clear expectations in its charter.

IS-IS and OSPF V2 mirror each other in many respects. Each is an SPF-based protocol, which means that link connectivity information is advertised by each router in the network and maintained in a database by every node. Routes are then calculated through the network by each router separately, but in a consistent manner and using consistent information, which results in rapid convergence on a usable set of routes.

The interfaces in an IS-IS or OSPF network fall into four categories:

- Local Area Network: A LAN is viewed as a stable and consistent set of neighbors with consistent addressing within an address prefix, which can use LAN multicast or multicast technology for communications. This permits several optimizations over describing them as pairwise point to point connections.

- Non-Broadcast Multi-Access: OSPF defines an NBMA interface as a special case of a LAN, which does not support a multicast or multicast capability. It is primarily used in Frame Relay and ATM networks.

- Point to Point: A point-to-point interface is an interface on which there are exactly two neighbors.

o Point to Multipoint: OSPF defines a Point to Multipoint is a grouping of point-to-point relationships over a common interface. It is primarily used in Frame Relay and ATM networks.

Of these types, it will be observed that only two support a multicast or multicast capability - the LAN and the Point to Multipoint Interface. The fundamental issue relates to the process of bringing up a new routing neighbor. In an SPF-based routing network, all routing databases must be consistent to guarantee consistent results. As a result, it is necessary only for a router joining an operational network to synchronize with one router in order to obtain that database. The routers on a LAN elect one among them (called the "designated router") to perform synchronization tasks; as a result, rather than experiencing database flooding traffic on the order of the square of the number of routers on the LAN, that traffic is linear with the number of routers on the LAN. Points to point links, of course, require no such optimization.

In the design of OSPF, Frame Relay was originally viewed as being much like a LAN, with internal connectivity that need not be visible to the routing protocol. For this reason, Frame Relay was modeled as an NBMA network, using a designated router like a LAN to make traffic distributions linear. A problem was discovered in Frame Relay networks, however, which used switching equipment that did not support dynamic rerouting around failed internal trunks. In this case, a failure of the trunk connecting the designated router and its backup (shown in figure 2) would cause a designated router election, in which various systems elected different designated routers. OSPF's solution to this is to wait for a uniform election result before continuing, resulting in a complete failure of routing.

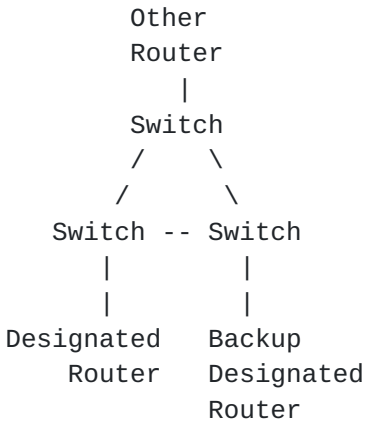

Figure 2: Routed IP network surrounding a Frame Relay network

The solution was to view a Frame Relay network as a bundle of point-

to-point connections, which was called a point to multipoint network interface. While this is subject to traffic volumes on the order of the square of the number of connected routers, the loss of an internal trunk does not result in the loss of external connectivity unless no connectivity exists.

In MANET environments, OSPF V2 or IS-IS are likely to encounter a number of challenges. The radio network is a multicast network, so it is tempting to think of it as a LAN, perhaps an 802.11 variant. However, in this environment several issues immediately result:

- When routers with different parameters on an interface, including area number or address prefix, find themselves in communication, they each assume that the other is misconfigured. As a result, they refuse to accept each other as routing neighbors.

- Because each router's view is slightly different, even among routers that choose to become neighbors, the designated router election has inconsistent and inconclusive results. While some sets of routers may converge on consistent designated router choices, the network does not, and routing is not even unstable - it is non-existent.

- If the router did calculate routes, other routers would understand from its advertisement that it was able to deliver traffic directly to any router using the same prefix, which would be untrue.

- Since flooding occurs away from the interface that information is received on, the only routers that will receive a given bit of routing information will be those within radio range of the originating router.

- If N routers advertise an LSA among themselves, in the average case each will send with a link state update or a link state acknowledge to the DR and from the DR to the others, for O(3N) messages.

- If a multicast link state update is sent, OSPF has each recipient respond with a unicast acknowledgement right away. In a CSMA network, this is a recipe for disaster; the various senders have a high probability of colliding. If acknowledgements or retransmissions are delayed for a random interval long enough to materially reduce the probability of collision, network convergence is delayed by the same amount.

- Since OSPF uses only provably bidirectional links, unidirectional links will be excluded from use.

The most straightforward repair within the existing specification is to consider the MANET to be a point-to-multipoint link, and allocate the interface addresses from a single large prefix per area. In this environment, routing through the MANET is straightforward, and the other issues are resolved. However, these issues remain:

- A router that is not configured for a certain OSPF area will not neighbor with routers in that area.
- In a multi-area, should a router change its area but retain the same prefix on the radio link, the prefix will appear to be in both areas, and devices in those or other areas will have incorrect routing to some subset of the addresses in that prefix.
- Link state updates can be multicast, but the acknowledgements are unicast. Thus, total transmissions are on the order of the square of the number of neighboring routers.
- Since OSPF uses only provably bidirectional links, unidirectional links will be excluded from use.

4.2 Ad hoc On-Demand Distance Vector (AODV) Routing

AODV is an example of a reactive protocol developed in the MANET context. The authors are Charles Perkins (Nokia Research Center), Elizabeth Belding-Royer (University of California, Santa Barbara) and Samir Das (University of Cincinnati).

It has, in draft 10, three messages: a Route Request, and Route Reply, and a Route Error. The Route Reply is essentially a route announcement or update, in the parlance of more traditional distance vector protocols; it says, "You can get to this IP Prefix via me". A Route Request, as its name suggests, searches for a route to an address. When a system needs a route from here to there, it emits a local multicast that floods to all systems within some number of hops away; those systems also learn from the Route Request a least hop count path to the originator of the request. If a copy of the Route Request gets to the target or to any system which has a route to the target, that system issues a Route Reply, which is forwarded along the best path to the source, and installs a route to the destination. This route has a timer on it; it will survive until a movement of one of the devices en route changes it, or until the timer expires.

A route reply is used in another way as well. It can optionally be periodically flooded to all neighbors within a certain distance; the specification refers to such behavior as a "hello", and suggests limiting the flood to directly accessible routers. In this way, all

neighboring systems normally have routes, and need not search among immediate neighbors.

The routers also keep state on any route that is in use; if a packet is sent from "here" to "there", each system en route tags the route with the fact that the previous hop router to "here" recently used the route. In the event that a route fails (the route is in use and times out, the next hop is lost, or the next hop issues a Route Error to it), it issues a Route Error to its neighbors that are using the route. This gets back to the source of the traffic. The source recalls how many hops away the destination was and issues a slightly wider diameter search, to set up a new route to the destination.

The protocol was originally specified to support IPv4 in a best effort model. It has been extended, in separate drafts, to carry IPv6 information, and to eschew routes that fail specified criteria. The latter is referred to as "QoS Routing", on the premise that a route which has no more than a certain percentage utilization of the link and no more than a specified worst case delay will deliver a specified quality of service.

One issue in robustness has been reported; it is possible to receive a Route Reply hello through a link that has a poor signal to noise ratio, and be unable to actually use the route for communication. Unfortunately, drivers may or may not report the signal to noise ratio, the signal to noise ratio does not necessarily translate into a statement that a certain percentage of traffic will survive a link, and mechanisms in 802.11b that should mitigate this are unimplemented and perhaps unimplementable in most drivers. Experimentation is ongoing with filters to detect and deal with this issue.

4.3 Optimized Link State Routing (OLSR) Protocol

OLSR is an example of a proactive protocol using Dijkstra's algorithm to calculate routes. Thomas Clausen, Philippe Jacquet, Anis Laouiti, Pascale Minet, Paul Muhlethaler, Amir Qayyum, and Laurent Viennot, all of INRIA Rocquencourt in France, originally developed it. Comparisons are made to OSPF, of the form "it is a simplified version of OSPF". It is fairer to say that it uses similar fundamental algorithms; it distributes connectivity information using a flooding algorithm, and maintains a route table calculated using the SPF algorithm. Unlike OSPF, the flooding algorithm is unreliable.

Fundamentally, the protocol consists of two message exchanges: hello messages and link state flooding (which includes both connectivity information and withdrawal of the same). Each system in the network emits a periodic hello, which lists the systems whose hellos it is hearing. If those systems can also hear it, the message identifies

bidirectional channels (channels which carry control traffic in both directions). As they listen to each other, they can determine that they may be in possession of information from some of their neighbors that other neighbors do not have; they are therefore also able to forward these link connectivity messages (or the withdrawal of those messages) to their peers. They can then run an SPF calculation to calculate the correct routes.

This breaks down in two places. One is that, since every system has a slightly different set of neighbors, every system can often justify repeating its message to someone. However, this logic results in far more relay transmission of the link state database than is actually necessary. A small subset of those relay systems is capable of delivering the same effectiveness in flooding. The difficult question is "what subset should be used?"

OLSR provides a way of resolving this, by asking each system to identify the neighboring system that seems most capable of giving it information about the part of the network it is not hearing from somewhere else, and designate that system as a MultiPoint Relay (MPR). The systems so designated form a lattice across the larger network, relaying routing information and multihop route messages among themselves, and relegating the other systems to a status more similar to that of hosts in the general Internet. This provides no area hierarchy, in the OSPF sense, but does provide a way to minimize the remulticast of routing information, and settles the network on a backbone of sorts. This backbone shifts, as the network itself shifts.

The other problem inherent in OLSR is the same robustness issue found in AODV. It is possible to receive a Hello through a link that has a poor signal to noise ratio, and be unable to actually use the route for communication. As with AODV, experimentation is ongoing with filters to detect and deal with this issue.

The robustness issue has another side effect, however, this more serious. Since flooding is unreliable and links are error-prone, there is a nontrivial chance that the information fails to be delivered everywhere. In such cases, routing may recover; the best route may not be calculated, but the network may succeed in calculating one that works. If routing does not, one can only hope that the route is not used until it is corrected.

Ongoing research is looking at the MPR determination heuristic and the use of filters to identify unacceptably lossy links.

4.4 Extensions to OSPF Version 3

OSPF Version 3 [12]is an extension of OSPF to IPv6, and uses IPv6 to accomplish its goals. It is quite similar to OSPF V2 in most respects, but an important consideration is that it uses the IPv6 link-local address for all inter-router links, and injects prefixes into an area in an LSA separate from the LSA used to construct the area's routing lattice. This reduces or eliminates complexities related to un-numbered links, choice of prefix, and so on, and adds some capabilities in prefix advertisement. Two properly configured routers can neighbor even if they have no prefixes in common, as a result. In a MANET, this is an important result.

MANET routing should be manageable in OSPF V3 if two extensions are adopted: area mobility and a "MANET" interface type with appropriate procedure and metric accommodation to the MANET network. If these two modifications are accepted, then the only remaining issue is that OSPF only uses bidirectional links, which is not necessarily bad.

4.4.1 Area Mobility

One issue in MANET routing using OSPF is what happens when a router finds itself faced with someone of a different area. For example, if a vehicle associated with one area goes around a hill to a region occupied by another, it still needs to communicate with its home base, but the only available connectivity may be through the new OSPF area. It is possible to configure the use of every possible area on the MANET interface, but this is problematic. It seems like a better approach would be to autodetect the area and join it. For scaling reasons, in some cases, a special "joining area" is also advisable.

Apart from administrative issues, autodetection is itself straightforward: as the device moves into the new area, it will start receiving hellos from new neighbors, which carry the configuration of the interface that they use. When configured to do so, and assuming that appropriate authentication has taken place, the router auto-creates an OSPF interface on the MANET interface that adopts those parameters. The hellos initiated on that OSPF interface will now neighbor with the new devices.

Several problems instantly materialize, however.

A router which is in two or more areas, in OSPF, is considered to be an Area Border Router. As an ABR, one of the areas it must support is area 0.0.0.0, the backbone area, and if there is only an indirect connection to other ABRs, a direct connection should be created using a virtual link. For several reasons, this is problematic in a manet. Unless the device is configured to be an ABR, it would be better if

it would advertise all of its prefixes in both areas, and depend on the multipath routing characteristics of OSPF to resolve the issue. This may be considered similar to running multiple instances of an OSPF process, and advertising all local prefixes in both.

A router which automatically discovers a new area needs an algorithm to determine when it should adopt or discard it, and therefore to create or collapse the auto-generated area configuration and database. The simplest approach I have come up with involves noticing whether a route exists to an ABR.

The fundamental principle is the principle of least change, coupled with the observation that OSPF often summarizes information into the backbone, creating a preferred, or "home", area for any given router. If a router has the option of communicating in its "home area" or another area, it should choose the home area, to maximize the scaling utility of summarization. If this does not result in connectivity to an ABR in its home area, however, it will only be able to communicate with routers in other areas by joining one or more of them. If it finds an ABR in one of those other areas, it can treat that area as its temporary "home" until it finds connectivity in its real "home area"; it should join that area and drop other non-home connectivity in which an ABR is found.

One issue with this is that it must join each area long enough to determine whether ABR connectivity exists, and stay in the area if ABR connectivity is absent. To minimize the pain of such exchanges, I propose an option or attribute on the OSPF Hello that indicates whether the device has connectivity to an ABR. A device trying to determine whether it need join an area can determine from this which area to join without the full exchange having taken place.

Another issue arises when several routers meet which have no connectivity to any ABR. In such a case, the algorithm as described so far requires them all to join all of their various areas, which at best is a great deal of overhead. A better approach would be to specify a special "joining area". This should be a stub area, to limit the advertisement of summaries into it. Routers issue hellos in this area if and only if either

- the router has no route to an ABR in any other area

or

- the router has a route to an ABR in another area, and
- it is receiving hellos from at least one router in the "joining area" which has no ABR connectivity to the backbone.

In the latter case, the router advertises itself as an ABR into the stub "joining area" (but not the other area), as will be discussed.

Consider the examples in figures 3a, 3b, 3c, and 3d.

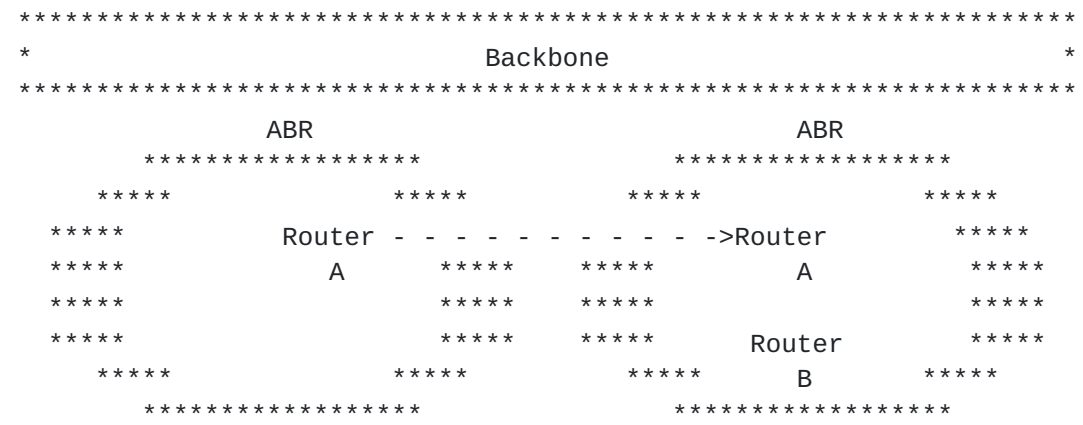

Figure 3a: Router changing areas through a non-connectivity zone

In Figure 3a, we see one fairly common situation: Router A leaves its area, has no connectivity, and then finds another area. It has the choice of joining the new area, meeting the devices in the area through the "joining area", or having no connectivity. The downside of using the "joining area" in this case is that it requires extra overhead. It should join the new area. As router A hears router B's hello multicasts (which indicate that it has connectivity to an ABR), it creates an OSPF interface in that area based on the values advertised in B's hello message including the area ID. In this new area, it will download the link state database, calculate its OSPF routing, and join the area. Even though its old area is summarizing prefixes into the backbone, and therefore into the new area, its own prefix being advertised into the new area will be a longer prefix match, and will therefore take precedence, even in the old area.

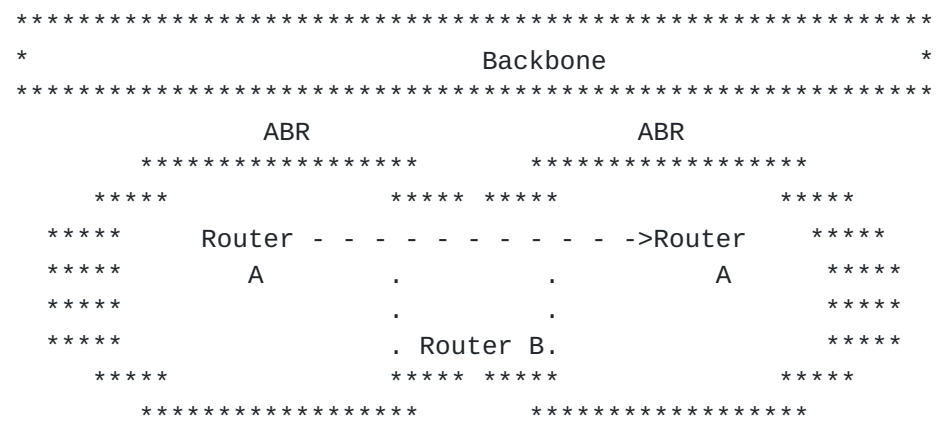

Figure 3b: Router changing areas through an overlapping region

Figure 3b is like 3a, with the exception that the two areas overlap. In this case, Router A has a choice as it moves from one area to another, as does Router B. The simplest choice for each to make in the region of overlap is to minimize their own level of change - they remain solely in their own "home" areas, and communicate via the backbone. However, as Router A moves, it finds that it eventually loses connectivity to its ABR, and therefore to the backbone. To communicate globally, it must therefore join the new area, which in essence reduces to the case in figure 3a.

Similarly, if the router is not in its "home" area but has connectivity to an ABR in the area it has roamed to, it has no reason to change areas other than rejoining its home area. It should stay in the area it has roamed to until that no longer works.

When the router comes into contact with a device with ABR connectivity in its "home" area, the same thing happens but with a different bias. Router A prefers its "home" area over all others due to the global optimization that summarization affords. Therefore, when it hears such a hello in its home area, it joins that area even if it has ABR connectivity in another area, and then leaves the other area. In such a case, for robustness, it does not actively leave the other area until connectivity to an ABR has been established in its own area.

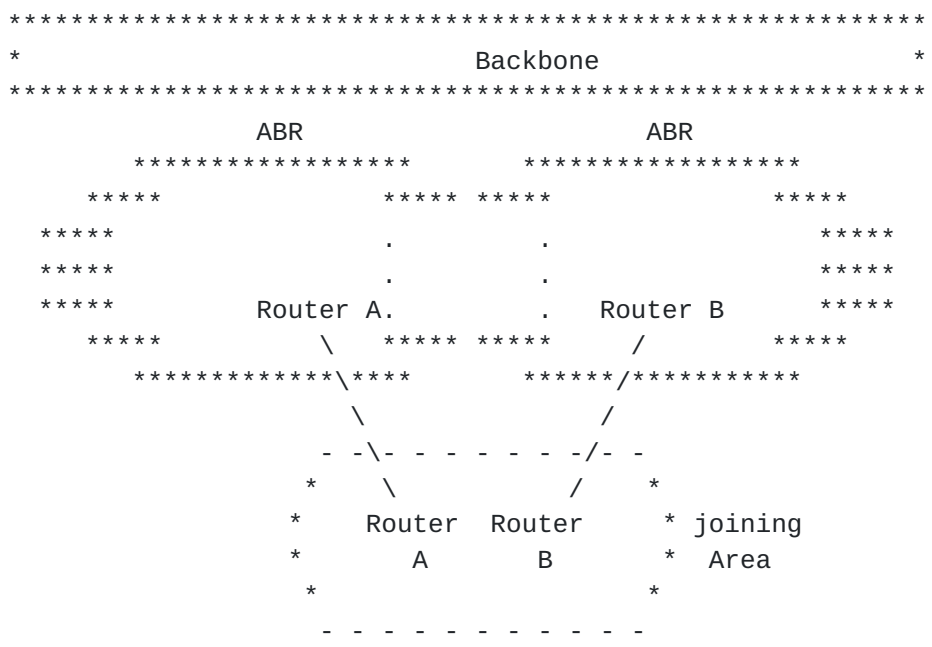

Figure 3c: Routers meeting apart from their "homes" in the "joining area"

In figure 3c, two routers leave their regions of connectivity, as Router A did in figure 3a. However, rather than finding each other's areas, they find each other as entities isolated from the backbone. As soon as they lose ABR connectivity, they started issuing hello messages in the "joining area", and now neighbor there. This affords them local connectivity (with each other).

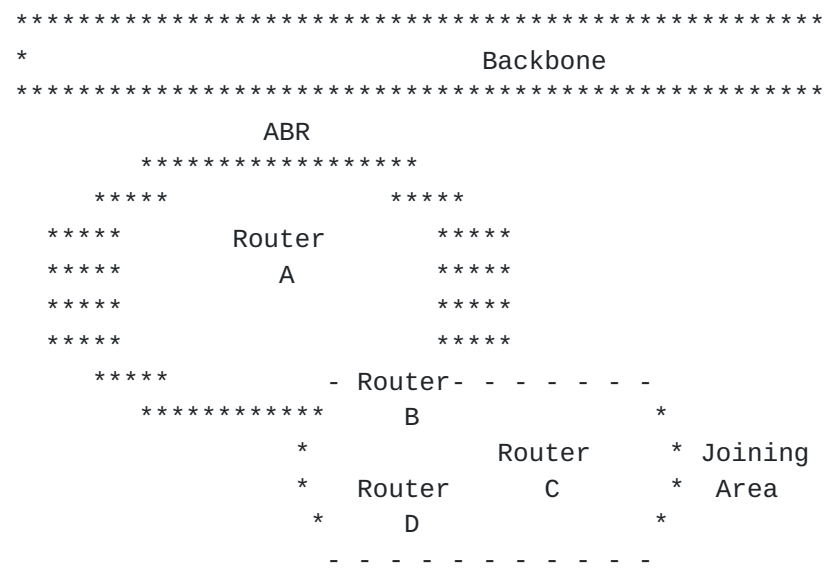

Figure 3d: Routers in the "joining area" meet another area

In figure 3d, a device (Router B) in the "joining area" hears hellos from a device in another area which has ABR connectivity. Its first instinct is, of course, to join that area, either because it is its home area, or because it is an area with ABR connectivity. However, doing so precipitously leaves the other routers in the "joining area" without connectivity. Therefore, the router does not actively leave the "joining area", but participates in a controlled switchover, leaving only when its services are no longer needed.

Once it has synchronized with the other area, it starts issuing hellos in the new area that indicate that it is connected to an ABR. Other routers in the "joining area" will hear these, and by the same logic, synchronize with it in the new area. In short order, the "joining area" will collapse into the new area, lattices, prefixes, and all, and the last router out will turn off the lights.

4.4.2 MANET Interface Type

Section 4.1details the behavior and issues of either the point to multipoint interface or a multicast interface. In context, it seems that MANET calls for an interface type which

- Is multicast capable, and uses multicast for link state flooding.
- Does not elect a designated router.
- Enables a router that becomes active with a large number of communicating routers simultaneously to synchronize its LSA

database with them serially (or at least one of them first) on a unicast basis, on the assumption that they are likely to already be synchronized among themselves.

- Results in a set of point to point relationships with its neighbors being advertised in its router links LSA.

- Repeats a new LSA in a multicast on the interface it was received on, both to implicitly acknowledge its receipt and to propagate it to neighbors who may not have received the initial multicast.

This is very similar to the point to multipoint interface type, with the exception of the final bullet. The implication is that it need not respond to a multicast announcement with a unicast acknowledgement; the multicast retransmission implicitly acknowledges the update. However, a unicast retransmission of the update needs to result in a unicast acknowledgement. Thus, an LSA update requires each router in the area to make a single multicast transmission (ie, there area linear distribution effects), potentially with some level of unicast retransmission. There is one side-effect of this behavior that bears investigation: sending a multicast which requires its receivers to each potentially send a message has correlated transmissions as a necessary result. In a CSMA environment, the implication is that they are likely to collide, resulting in a high rate of loss.

Many of the MANET routing protocols find ways to not have correlated reasons to transmit, by not acknowledging, or by including the acknowledgements in uncorrelated messages. On multiplexed interfaces, OSPF is often implemented with some form of randomized delay or link layer serialization prior to acknowledgement, to limit this effect. There is an issue in randomized delays, however, in a radio environment: for the randomization to have the necessary effect, the distribution of the messages must be uniform, and the interval must be long enough that any two transmitters have a very low probability of collision. As a result, the period over which the transmissions occur must be a multiple of the message duration and the number of relevant routers, minimally on the order of two to four times that product. This tends to result in an arbitrary extension of the network convergence interval.

One could also imagine solving that using link layer disciplines similar to that used in LDDI, wherein each router generates a link layer sequence number and transmissions are made in that order. The routing protocol could carry in its "hello" message a "transmission sequence number", which is essentially a random number that the routers verify is different among neighboring routers. In LDDI, the link protocol assigns each device a time slot by giving it a sequence

number in a range 1..N. System number 1 gets the first turn; it either leaves a minimum-sized message duration idle, or transmits a message. System 2 listens, and when System 1's message duration or transmission is done, leaves a minimum-sized interval or transmits a message. The process repeats through system N, who will have waited through N-M idle time slots and seen M messages. System N then sends its own message, if it has one, followed by a short message as though from system 0, triggering the start of a new round. Through this scheme, they effectively pass a token for access to the otherwise-CSMA link, but do so without a "token" message.

In a radio network, when acknowledging a multicast message, this could be emulated if every router organized a sequence number among its neighbors. This would be a random integer not duplicated among neighboring routers, an a relatively small range such as 1..twice the number of neighboring routers. It is transmitted in the "hello" message, and any router receiving its own number from a neighboring system is obligated to change its number. When a multicast link state advertisement is received, or similar message which the router realizes that it must acknowledge and is likely to collide with others while doing, it schedules the message sequence*interval time units in the future, to limit the probability of collision; with CSMA techniques, there is an improved chance of collision avoidance. "Interval", of course, must be defined; one might expect it to be the MANET interface's MTU in bits divided by its link speed, perhaps with some randomization.

A simple way to generate the sequence number would be to sort the addresses listed in the combined hellos of a set of neighbors. If a system has N neighbors, and constructs the union of the link addresses or router ids that they are advertising, and sorts them numerically as unsigned numbers, any given system's sequence number may be its own index in that array. While every system will have a slightly different view of the array, the approach at least has the possibility of distributing the traffic.

There are good reasons to put this behavior in the link layer protocol, as the designers of LDDI did. However, if the MANET routing protocol is the only protocol that has a message-burst issue, one could also argue for making it a configurable feature if the routing protocol.

4.4.3 Metric issues: selecting a path with adequate link quality

OSPF leaves the design of the routing metric to the administration, with only the proviso that it will use the route that minimizes the sum of the metrics en route, and they must fit in a specified range.

One example might be a hierarchical function

metric = f(policy)*2^j + g(quality)*2^k + h(throughput)

where

- The OSPF metric is in 1..2^16-1, and specifies "goodness" of the link. The "best" links have small numbers.

- "f(Policy)" is a variable with four to eight values, and indicates the device's willingness to act as a router. A device with a stable power source, for example, may be more willing than a battery-powered device, and a device with a recently charged battery may be more willing than a device whose battery is depleted. Preferred links have small numbers.

- "g(Quality)" is a measure of link quality to the neighbor, with a small number of values indicating "good" to "poor" quality - on the order of two to four numeric values plus "not reachable". This might derive from the signal to noise ratio, the correction rate of the convolutional decoder, the bit error rate, or similar measures. The measure it is derived from should be filtered using a technique such as a median value filter feeding into an exponentially weighted moving average, with the result being compared to thresholds to determine the value to advertise. High quality relationships have small numbers.

- "h(Throughput)" is a measure of the capacity of the link to move data, perhaps an 12 bit integer. Since radios vary in their effective transmission rate, both by design and environment, this may have to be variable. If it is, changes should be similarly filtered. The fastest links have small numbers

The reason for such a complex metric is that mobile ad hoc networks have a more complex environment than wired networks. As mentioned in Section 2, signal quality can vary on the same neighbor relationship in the absence of motion, and neighbor relationships can be very dynamic. The metric should enable path optimization where it can, but focus on measured link quality and communication policy where it must.

The downside of this approach is that the utility of links can change rapidly and dramatically, changing the routing. The changes in routing, of course, are to work around problems in the network, and without the changes, communication is hindered. However, oscillation is itself problematic (as the BBN SPF experience demonstrated), and to be minimized.

One issue that remains, with this metric as proposed, is that it describes what is being received from a neighbor, while OSPF metrics typically describe what can be sent to an interface or a neighbor. Ideally, this should be advertised at the link layer, so that the network layer protocol need not change. Otherwise, the simplest way to describe this will be to have the metric advertised by the routing peer be used in the SPF calculation, which at best is a cumbersome modification to that algorithm.

In any event, it seems best if the metric, or at least the "goodness" and "speed" components, is considered a vale read from or presented through a link layer API. The ideal API would enable reading the value on demand, and either present the new value when significant events happen (such as a major change in the metric) or trigger its being read.

4.4.4 Scaling Properties

The scaling properties of a manet routing protocol are a subject of frequent concern. It is to ameliorate these issues that OLSR developed the concept of a Multi-Point Relay, or MPR. In essence, if a few devices can stand in routing as interchange points and the rest can adopt the role of a host, the scaling properties of a network are improved.

Without further modification, OSPF cannot readily develop that role. What it can do, however, is limit the neighbor relationships. If a router discovers that it has more neighbors than some threshold, perhaps the number of Router IDs than will fit in a Hello message, one option is to send only a subset of those Router IDs. The router might choose, for example, to neighbor only with those routers whose metrics are in the lowest third of the range, or only with those most "willing" to connect.

4.4.5 IPv4 routing using IPv6

As discussed in the IP Version 6 Addressing Architecture [5][29], there are two ways to carry IPv4 addresses within IPv6 addresses. One, written "::a.b.c.d", describes IPv4 addresses whose end system supports both IPv4 and IPv6 and need not be translated either way. The other, written "::FFFF:a.b.c.d", describes an address used by an IPv4-only host.

If the latter is carried as an IPv6 address in OSPF V3, the end system can (at least in theory) be relied on to send it as IPv4 messages; in the event that a host sends an IPv6 message to it, a translating gateway can translate the messages. However, OSPF V3 does not natively install IPv4 routes, depending instead on an IPv4

routing protocol to do this.

It seems that it would be wise to implement a configuration option that would import IPv4 interface prefixes and advertise them in IPv6 routing, and would generate an IPv4 route table from IPv6 routes in these cases. This would facilitate IPv4 connectivity in an IPv6 routing infrastructure.

5. Application and Transport Protocols

A related set of issues has been reviewed by various researchers at the transport layer and its counterpart in applications. Wireless networks are notably subject to loss due to issues in physics and timing issues among devices that cannot "hear" each other but are attempting to communicate with other devices, which can. The temptation is to change TCP (which is widely used in Internet communications) or to absorb the transport into the application layer in a special manner. One example of such an application protocol runs on UDP and places a sequence number in each message. The entire file is transferred in serial order, with the receiver acknowledging received messages. Transmission of unacknowledged packets is repeated at intervals until all packets are acknowledged. Such procedures avoid the vagaries of TCP's congestion avoidance behavior, but are obviously ill-suited to a larger internet.

Section 2.1 stated, however, that interoperation with the larger internet is indeed important. Manets occur as stand-alone networks, but they also occur as stubs off larger backbones, and as transit systems between small edge networks. Further, commercially available applications are designed to use TCP. In order to use those applications in manet environments, we face a choice: we can convince the manufacturer to rehost the application for our amusement, or we can adapt the environment to support the application. Both avenues have difficulties.

One particularly promising development is found in a combination of TCP Congestion Control [8] with "New Reno" Fast Recovery [9], Selective Acknowledgment [1][13], and Explicit Congestion Notification [14]. If congestion is explicitly signaled and managed, then lost data may be more aggressively retransmitted, and still remain interoperable with more reliable parts of the Internet.

6. Conclusions

The discussion yields no strong conclusions at this time. A number of protocols have been considered in research, with the result that we have learned quite a bit about these networks. Some of that learning has been relearning lessons already learned in the Internet itself, with the resultant reinvention of related solutions, or remembering the reasons they were invented.

6.1 Selecting a protocol

It is not obvious that a single protocol is an adequate solution for all MANET problems. As in wired networks, there is room for creativity, and for difference of opinion. To give an idea of the classes of applications and aspects of solutions that drive them, consider these three cases:

6.1.1 Military Communications

A SEAL team and a battalion of Army Rangers land in some random mountainous country in southern Asia. Rather than depend on local infrastructure, they will use satellites and may install temporary fixed infrastructure if the plan of battle warrants it.

In contexts like this, proactive manet protocols (TBRPF, OLSR, a modified OSPF, etc) make the most sense. The important issue is not the high level of dynamism. It is the fact that there is a continuous lower-rate change, the fact that there is no external infrastructure to depend on, and the fact that applications in the network need the network to be stable and working when they decide the time has come to use it.

6.1.2 Automotive Networks

Consider a service in which automobiles can "talk to the road" to obtain maps, weather, congestion-and-accident-ahead updates, and so on. It might also be used to check passing hotel availability, and to talk car to car, for example to facilitate navigation in a caravan. In Japan, this is done with wideband CDMA in late-model automobiles.

It does not make a lot of sense to view "every road in the world" as a single manet; more likely, sections of road, geographic regions, administrations running the roads, or political regions will be manets. It is also not obvious what the "backbones" would be - freeways, perhaps, but in much of the world distinguished systems are not obvious. Protocols like AODV have strong appeal in such environments. They would, however, need some provision for crossing

administrative boundaries in the same fashion, perhaps by having nodes that pass an administrative boundary give it some routing information as they pass.

Mobile IP has been proposed as a means of managing administrative boundaries. However, it seems to not make operational sense to this author. The issue is the rate of change. A vehicle mostly needs to "talk" to the next street lamp and the car a half a mile away, especially when the car takes an unexpected turn. The radio in the car halfway in between, which I have been traveling with for the past ten minutes, is a more stable relay system than the street lamps I'm passing or any ISP link. There is no reason to be constantly updating my home ISP every time I pass another street lamp, and no reason to form a new IPv6 address, either.

6.1.3 Classic Sensor Network

There's a brushfire in your favorite place to not have a brushfire, and we want to manage it. So we have a plane drop "golf balls" all over it - organic styrofoam-like stuff protecting the level of electronics one finds in a wrist watch these days, which has just enough smarts and capability to GPS where it is located, to periodically say "I'm [here], and I'm still here", and to forward similar messages from other devices. These talk to a central PC, which keeps track of which "golf balls" are still up and which have failed. When a device fails, the central system dispatches someone to wonder why.

This seems to be a good application for a source routed protocol like DSR. The only routing any given system needs is its current route to the central system. Exploration surges will be huge at times, especially at first, but having once subsided will likely be of limited impact. If the central system needs a route back, it has plenty of capacity in which to store it.

6.2 Commercial Considerations

Commercially, if I had to hang my hat on two solutions, one reactive and one proactive, at this time I would go with AODV and some extensions to OSPF V3. The reasons for those choices are:

- AODV is well advanced and has good capabilities for its target networks.

- AODV is publicly documented, as opposed to being partially documented in corporate research reports or intellectual property declarations.

- AODV has preliminary work on QoS Routing and IPv6 routing in place; for other protocols, these are futures.
- OSPF has significant market traction.
- OSPF can be extended to include MANET-type interfaces, incorporating lessons from OLSR, TBRPF, and so on.
- If this is done, OSPF can span wired and wireless networks.

7. Security Considerations

In routing, beside the fundamental faults of undebugged code, there are three primary threats.

- o A network may require privacy in order to not give away important information.
- o A network device may be improperly configured, so that the information it exchanges is incorrect, or its presence otherwise disrupts the network.
- o A rogue system may mimic, inject, or remove routes in order to disrupt traffic flow.

OSPF performs simple neighboring parameter verification to detect and avoid misconfigured neighbors, and several protocols mention IPSEC for authentication or privacy. Besides that, no proposed manet routing protocol explicitly addresses any of these issues.

In MANET networks where link privacy is a significant consideration, it is logical to presume physical or link layer encryption. IPSEC encryption could be used, but a radio listener who read the IP headers could deduce much of the routing information. This could be of military value, for example; if I know that a large number of communicating systems are reached via one system and a small number by another, and can determine which is which via radio direction finding, I may be able to locate a force concentration or detect the splitting off of a sortie. I may also be able to locate a single point of failure, a system that is temporarily critical to battlefield communications, and know to target it.

The deployment of any form of link or IPSEC encryption, however, requires some form of key distribution. This is a problem which has not been solved at this writing.

Neighbor authentication and privacy techniques do not, however, place a signature on an LSA, such as is suggested in OSPF with Digital Signatures [3], or otherwise address the issues raised in Routing Policy System Security [11]. Thus, these protocols do not secure the network against a rogue system once its neighbors decide to trust it.

8. Acknowledgements

The author acknowledges the inputs of many in this document, most especially Joe Macker, Scott Corson, Abhay Roy, Alexander Zinin, Elizabeth Belding-Royce, and Charlie Perkins. Joe commented in detail and contributed text to some sections of the document. Brainstorming with Abhay was particularly useful in working out the details of OSPF V3 routing exchanges.

References

[1] Mathis, M., Mahdavi, J., Floyd, S. and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, October 1996.

[2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[3] Murphy, S., Badger, M. and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

[4] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[5] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

[6] Rajagopalan, B., Nair, R., Sandick, H. and E. Crawley, "A Framework for QoS-based Routing in the Internet", RFC 2386, August 1998.

[7] Corson, M. and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, January 1999.

[8] Allman, M., Paxson, V. and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.

[9] Floyd, S. and T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm", RFC 2582, April 1999.

[10] Apostolopoulos, G., Kamat, S., Williams, D., Guerin, R., Orda, A. and T. Przygienda, "QoS Routing Mechanisms and OSPF Extensions", RFC 2676, August 1999.

[11] Villamizar, C., Alaettinoglu, C., Meyer, D. and S. Murphy, "Routing Policy System Security", RFC 2725, December 1999.

[12] Coltun, R., Ferguson, D. and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

[13] Floyd, S., Mahdavi, J., Mathis, M. and M. Podolsky, "An Extension to the Selective Acknowledgement (SACK) Option for TCP", RFC 2883, July 2000.

[14] Ramakrishnan, K., Floyd, S. and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.

[15] Das, S., Perkins, C. and E. Royer, "Ad Hoc On Demand Distance Vector (AODV) Routing", draft-ietf-manet-aodv-10 (work in progress), January 2002.

[16] Perkins, C., "IP Flooding in Ad hoc Mobile Networks", draft-ietf-manet-bcast-00 (work in progress), November 2001.

[17] Johnson, D., Maltz, D., Hu, Y. and J. Jetcheva, "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks", draft-ietf-manet-dsr-07 (work in progress), February 2002.

[18] Gerla, M., "Fisheye State Routing Protocol (FSR) for Ad Hoc Networks", draft-ietf-manet-fsr-02 (work in progress), January 2002.

[19] Gerla, M., "Landmark Routing Protocol (LANMAR) for Large Scale Ad Hoc Networks", draft-ietf-manet-lanmar-03 (work in progress), January 2002.

[20] Jacquet, P. and T. Clausen, "Optimized Link State Routing Protocol", draft-ietf-manet-olsr-05 (work in progress), October 2001.

[21] Lewis, M., Templin, F., Bellur, B. and R. Ogier, "Topology Broadcast based on Reverse-Path Forwarding (TBRPF)", draft-ietf-manet-tbrpf-04 (work in progress), January 2002.

[22] Johnson, D. and J. Jetcheva, "The Adaptive Demand-Driven Multicast Routing Protocol for Mobile Ad Hoc Networks (ADMR)", draft-jetcheva-manet-admr-00 (work in progress), August 2001.

[23] Labiod, H. and H. Moustafa, "The Source Routing-based Multicast Protocol for Mobile Ad Hoc Networks (SRMP)", draft-labiod-manet-srmp-00 (work in progress), November 2001.

[24] Perkins, C. and E. Belding-Royer, "Quality of Service for Ad hoc On-Demand Distance Vector Routing", draft-perkins-manet-aodvqos-00 (work in progress), November 2001.

[25] Perkins, C., "IP Address Autoconfiguration for Ad Hoc Networks", draft-perkins-manet-autoconf-01 (work in progress), November 2001.

[26] Belding-Royer, E., "Global Connectivity for IPv4 Mobile Ad hoc Networks", draft-royer-manet-globalv4-00 (work in progress), November 2001.

[27] Wakikawa, R., "Global Connectivity for IPv6 Mobile Ad Hoc

Networks", draft-wakikawa-manet-globalv6-00 (work in progress), November 2001.

[28] Zitterbart, M. and K. Weniger, "IPv6 Stateless Address Autoconfiguration for Hierarchical Mobile Ad Hoc Networks", draft-weniger-manet-addressautoconf-ipv6-00 (work in progress), February 2002.

[29] Hinden, B. and S. Deering, "IP Version 6 Addressing Architecture", draft-ietf-ipngwg-addr-arch-v3-07 (work in progress), November 2001.

[30] Yi, Y., "Passive Clustering in Ad Hoc Networks (PC)", draft-yi-manet-pc-00 (work in progress), November 2001.

Author's Address

Fred Baker
Cisco Systems
1121 Via Del Rey
Santa Barbara, CA 93117
US

Phone: +1-408-526-4257
Fax: +1-413-473-2403

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.