Mobile Ad Hoc Networking and the IETF

Joseph P. Macker[a] **M. Scott Corson**[b]
macker@itd.nrl.navy.mil *corson@isr.umd.edu*
[a] Information Technology Division, Naval Research Laboratory, Washington, DC, USA
[b] Institute for Systems Research, University of Maryland, College Park, MD, USA

This article provides a snapshot of work underway within the Mobile Ad hoc NETworks (MANET) Working Group of the Internet Engineering Task Force (IETF). The article presents some issues currently under discussion on the MANET mailing list, and also summarizes the proceedings of the last MANET WG meeting.

I. Architectural Issues from the Mailing List

A. Scope of the Work

A thread of discussion continuously reappearing in different forms on the mailing list *(manet@itd.nrl.navy.mil)* concerns the scope of the MANET effort. There is a common misperception that the working group is attempting to solve *everything*, i.e. all problems relating to addressing, routing, security, interoperability with the existing fixed Internet, etc.—it is not. Its goal is to standardize routing technology, and whatever additional technology is necessary to support IP-level routing in a MANET. It is this latter point that perhaps causes confusion. As there currently exists no Internet standards for MANETs, and since this is a very different networking paradigm, it is easy to see that much work is needed, and hard to see how anything can be accomplished without solving many of the problems at hand. However, many of the necessary solutions concerning issues such as interoperability and security can and will be developed in concert with and in the context of other working groups—some of which exist and some of which may yet to be formed. While not an explicit goal, it is in the interests of this group to help formulate the MANET interoperability architecture necessary now (IPv4), and in the future (IPv6), and to help coordinate its integration within the IETF. As representatives of a new technology, it is also incumbent on this group's participants to be ambassadors for MANET technology within the IETF, and to seek out the appropriate working groups as necessary to facilitate integration of MANET technology into the greater Internet.

B. Nature of the Work

Another reappearing thread concerns the nature of the group's effort and what it is trying to accomplish. This is a *standards* effort. As such, it is trying to strike a *balance* between a closed and open system architecture. Many points raised on the mailing list concern system performance issues that certainly can be most efficiently realized by a customized, closed system design (a.k.a. "stove-pipe" systems), but this runs counter to the notion of producing an open, standards-based system. A primary aspect of MANET is to facilitate development of an *open*, *extensible architecture* for IP-based mobile ad hoc networking. Architecting such a system requires trading off performance to gain flexibility and interoperability. A well-architected system possesses minimal performance impact and maximum flexibility. The MANET WG hopes to help realize such an architecture by developing routing and related supporting technologies. For its architecture to be open and flexible yet high-performance, it needs to capture its functional integrity in the well-known protocol and interface interrelationships of a vertically-integrated protocol stack. Vertical integration is achieved through usage of "extended interfaces" which increase the degree of interaction between adjacent protocol layers.

The Internet MANET Encapsulation Protocol (IMEP) is one step in that direction. It interposes a translation layer between the IP layer and the routing and other Upper Layer Protocols (ULP). This layer logically unifies a collection of underlying physical layer technologies into a single IP routing fabric. Through this unification also comes a partial, although not complete, homogenization of the view of the underlying heterogeneous technologies from the perspective of the routing and other ULPs. In other words, each physical layer technology (e.g. a WAVELAN radio) will have a unique set of performance characteristics (range, power, bandwidth, error rate, etc.) and a different device driver interface to control its operations. Typically, this interface information is not made available to the internetwork layer, and the device is only identified by an IP address. However, this level of information hiding is high, and does not sufficiently expose the physical-layer characteristics to the routing level to permit sophisticated forwarding policies in a multi-technology, multihop wireless routing fabric. Through IMEP, the interface is not only identified by an IP address—it can be augmented with a standardized set of characteristics which are initialized by each particular underlying technology. Thus, from the routing policy's perspective, the interfaces have the same characteristics (homogeneous), but have different values or settings for these characterisitics (heterogeneous). Based on these values, forwarding and load balancing policies can make decisions that reflect the underlying layer. Such an extension of the IP interface to include link-layer characteristics facilitates support for topology and context-aware routing, and for prioritized or reservation-based, QoS-aware networking. This latter capability has yet not been specified in the IMEP draft, but incorporation of a flexible layer like IMEP into the MANET architecture, where such functionality can be added, is necessary for enabling future work in this direction.

While many aspects of MANET technology are still in the research phase, the WG is not a research group. Thus only technologies which have emerged from the research world and which (i) have the appropriate architectural characteristics (e.g. distributed, bandwidth-efficient, etc.), (ii) have been proven formally correct and (ii) have been shown via simulation to have a performance benefit that warrants their overhead and implementation cost should be proposed for standardization. For example, for a "load balancing" policy for packet forwarding to be standardized, it should be shown theoretically and practically to be stable, and to improve overall system performance. The latter point means that the performance improvement derived from a load balancing policy must be balanced against its cost, and implemented only if there is a net gain. Many concepts and ideas that have been floated on the mailing list do not yet meet these criteria.

II. DC MANET WG Meeting

The MANET Working Group last met in Washington, DC in December 1997. Portions of this section are summarized from the meeting's minutes taken by Eric Guttman of Sun Microsystems. The meeting consisted primarily of presentations of proposed approaches for MANET routing and support technology, with additional presentations focusing on the use of simulations to aid in protocol development and comparison. The meeting closed with an open discussion of MANET issues.

A. Routing and Support Technologies

Prior to the meeting, three Internet-Drafts (I-D) were authored presenting approaches for MANET routing. These are the Temporally-Ordered Routing Algorithm (TORA), the Ad hoc On-demand Distance Vector (AODV) routing algorithm and the Zone Routing Protocol (ZRP). Additionally, a draft was authored on an underlying routing support protocol, termed the Internet MANET Encapsulation Protocol (IMEP), that is proposed for use in MANETs. The following briefly describes the proposed protocols.

1. *IMEP*

The IMEP is a multipurpose network-layer protocol designed to support the operation of many routing algorithms or other ULPs intended for use in MANETs. The protocol sits in the stack between these protocols and the IP layer. The protocol currently incorporates mechanisms for supporting link status sensing, control message aggregation and encapsulation, broadcast reliability, network-layer address resolution and provides hooks for interrouter security authentication procedures. Future capabilities are likely to be added. Many of these features are optional, and can be turned on or off as necessary depending on the underlying physical layer network technology. A key concept supported by the protocol is the notion of an IP routing "fabric", as discussed in the first article of this series, which presents a partially homogenized view of a heterogeneous collection of underlying physical-layer technologies and topologies.

2. *TORA*

The TORA is neither a link-state, path-finding nor distance-vector routing algorithm. It is one of a class of "link-reversal" algorithms designed to enable loop-free routing. TORA uses a metric referred to as the "height" of the node to assign a direction to links for forwarding packets to a given destination. The node heights can be totally-ordered lexicographically, and thus define a directed acyclic graph rooted at the destination at all times. TORA implements three principle functions: creating routes, maintaining routes and erasing routes. Creating routes is performed "on demand" using a query/reply process. Maintaining routes is triggered when a node loses its last downstream link, and either has one or more upstream links or needs a route itself. The algorithm reorients the directed acyclic graph such that all paths lead to the destination. TORA does not compute the shortest path: paths may be suboptimal in this respect. When initially constructed, the paths are close to optimal and tend to "loosen" as it reacts to topological changes. A secondary mechanism, one not tied to the rate of topological change, is used to reoptimize routes. Network partitions are detected when a distributed link reversal wave reflects off the network's boundaries and returns to the node which originated it. A node that detects a partition initiates the process of erasing the invalid routes.

3. *AODV*

AODV is described as a loop-free routing algorithm that is free of the "counting to infinity" problem. It is essentially a distance vector algorithm which incorporates priniciples from on-demand source routing to reduce the number of actively maintained routes. It uses sequence numbers to achieve loop freedom and maintains one route per destination. The source sequence number is provided for a reverse route. The destination sequence number is used for the forward route. A route discovery request will not obtain older routes than the one which is implied knowledge of the requester (a request includes the sequence number of the destination already known to the requester). Previously broadcasted messages will also not be rebroadcasted through the use of a unique id. In AODV, nodes notify all neighbors when link breakage is detected. It is claimed there is little overhead for AODV and that it is simple to implement.

4. *ZRP*

The ZRP is a scheme which combines the use of traditionally-proactive, flat routing algorithms (e.g. distance vector) with reactive, on-demand source routing. Each node considers itself to be in the center of a circle (or "zone") with some radius (measured in hops). Nodes are assumed to know the identity of all nodes within their zone and, consequently, all nodes on the border of the zone. Flat routing is used between nodes in a zone (intrazone routing) and source routing is used between zones (interzone routing). The latter is accomplished through a process termed "bordercasting." With bordercasting, when a destination node is not within the local zone, a node transmits or bordercasts a query to all nodes on its zones border asking them if they know where the destination is located. If so, they reply. If not, they continue the bordercast flood. The intent of bordercasting is to produce a more efficient flooding process than pure flooding, and to produce source routes consisting of sequences of border nodes, i.e. routes which are shorter than pure source routes (which would otherwise contain the identities of all intervening nodes).

B. Simulation Efforts and Technologies

There were also two invited presentations on simulation technologies and approaches for MANETs.

1. NS2 Simulator

The presentation gave an overview of how the NS2 simulator is structured and how it might be useful to the MANET WG. The simulator simulates traffic, multiple protocols and scaling. The simulator consists of a C++ engine combined with Object-oriented TCL (OTCL). Protocols, links, nodes, packets and other elements are modeled. Objects are implemented in C++ and connected and controlled using OTCL. It is a fine grained simulation (e.g., links may require 3 objects). Routing protocols may be constructed from a system of objects. Error generators can be added to a topology to simulate motion using random distributions. Queuing is supported with various different algorithms. Visualization is possible using the Network AniMation (NAM) tool. There is currently no help desk—one must get help via the mailing list. There is a user community that is beginning to solve problems independently. The simulator was developed by researchers principally involved in transport layer protocol design and, as such, has a wealth of detailed models of TCP implementations. It also supports other protocols like RTP, RTCP and SRM.

Regarding its suitability for usage within the MANET WG, what is lacking is the general abstraction of a "wireless broadcast channel". There are researchers doing mobility simulation with a NS derivative. They can be reached via the mailing list. It is likely they are developing such an abstraction as several wireless multiple access protocols have recently been added for wireless LAN research.

2. Opnet-based Modeling Efforts

Recent work at Mitre involved the usage of the commercially-available simulation package Opnet to model MANETs. This work was funded by the US Army, and concerned evaluation of wireless networking protocols for their usage in future Army networks. As such, the simulation model considered mixes of traffic classes, patterns and service requirements typical of military systems along with suitable models and distributions of nodes, terrain and mobility factors. The physical and lower layer protocols were approximated as noise and background interference is hard to simulate. It was concluded that physical layer overhead is important to consider in terms of framing and transmission overhead. The evaluation framework consider the following factors for networks of 300 and 1000 nodes: traffic mixes, types; connectivity (dense/sparse); source and destination addressing; mobility factors; traffic loads; network sizes; and routing protocols. Some conclusions regarding network performance is that TCP doesn't work well in MANETs. By the end of the year, the project at MITRE will let folks have access to their model. They will start work on layouts for terrains, using averages and histograms for their simulators. Mixtures of voice, data with various distributions of range will also be included.

C. Working Group Discussion

A period of open discussion was held at the end of the DC meeting to review working group status and discuss issues. One of the action items from the Munich IETF MANET meeting was to develop a longer WG schedule given the relatively short milestone schedule of the present charter. A candidate revised charter schedule was reviewed in DC during this meeting. The charter was reviewed in DC during this meeting and there was consensus among the group that it appeared sensible.

A rough outline of the agenda for the WG follows:

- By Summer 1998, initial MANET working proposals and prototypes should exist and initial performance results should be presented.
- By March 1999, any required modifications to protocols should be achieved and additional performance analysis/comparison should be completed. Progress should continue during this timeframe on the further development and enhancement of prototype protocol capabilities and I-D specifications.
- The goal is that roughly by December 1999 we would have a standard track protocol(s) to advance as a proposed standard.

Now that the working group has some intitial I-Ds to work with, it was agreed that appropriate revision and group input on existing material needs to continue.

The following decisions were made:

- Obtain final group input and wrap up (as informational documents) the MANET Issues and Terminology drafts in the February 1998 timeframe.
- Further discuss requirements for various protocol proposals and present a taxonomy of approaches at the LA meeting in March-April.

Some topics for discussion at the upcoming LA IETF include:

- Link layer interface issues—how much more technical research is required here before sufficient understanding exists so that standardized interfaces to the link layer can be incorporated into the routing standards process. How far can and should MANET investigate such issues? Is this more appropriate for another WG or forum?
- Security considerations—MANET requirements and issues need further discussion.
- Taxonomy and discussion of present draft proposals.

A number of other issues and comments were raised during open discussion:

- Implementation status and reviews should be done as they are ready.
- Qualitative and quantitative comparison of protocols: We need common models so that we can compare and analyze results achieved by the different design efforts. The

group will continue to push for the development and use of common models as appropriate.

It was noted that while common simulation and traffic models would be useful, their *mandatory* use across the board may be impractical. This would be an ideal way to make comparisons between the protocols. However, simply producing and defining some models and techniques would be valuable output from the MANET WG.

- What are the expected interactions between other Internet protocols? There should be focus on DNS, Security, etc. The MANET WG should focus on the routing implications of such interactions.

To clear up some of the complex requirements and issues involved with mobile routing, it was suggested that applicability statements be included when and where appropriate by authors of candidate I-D specifications. These statements should include notes on protocol assumptions and requirements, and should also include a description of network domain(s) for which the approach is best suited towards as well as any known limitations. This will help clear up the difficult job of analyzing the appropriateness of protocols for mobile routing where it is likely that no "one size fits all" solution exists.

Several open issues were discussed for future consideration within the working group. The requirement and assumptions behind the use broadcast technology by protocols was brought up. Some protocols are not well suited for use in the absense of such link-layer technology. It was recommended by the chairman that applicability statements will have to be written on link-layer assumptions and requirements of routing proposals. It was also noted that with some wireless media, you cannot eavesdrop on communications. This may additionally affect the performance or assumptions of particular protocol candidates.

The group discussed the general topic of supporting heterogenous link layers—in the case of MANET, this means a heterogenous multihop fabric. It was pointed out that this is a major goal of routing at the IP layer and should be preserved within MANET unless strong arguments are given to the contrary.

An additional point was raised that emphasis so far has focused largely on potential radio routing technologies. There are also "non-radio" wireless technologies, e.g. diffused IR, that can also be used to form ad hoc networks. This raised the question again as to whether the group should write about "channel access" and "specific link layers" in future Internet drafts. It was cautioned that if care is not taken, this can lead to less general designs at the Internet layer. But there was consensus that this warrants more exploration and that such technologies form a potentially valuable role if appropriate standarized interfaces to various broad classes of link-layer technologies are defined.

It was generally agreed that MANET security issues need to be further discussed and explored. Should the WG consider denial-of-service attacks? It was suggested that a more detailed security context for MANET approaches be developed—with explicit work items which can be solved and with their assumptions clearly stated.

Towards the end of discussion, the topic of optimizing link-layer interaction was again discussed by the group. With power and range constraints, specialized link-layer information can be very useful to upper layers making routing and forwarding decisions. These issues need more exploration and discussion within the group as they form a central trade-off between creating more generalized mobile routing protocols versus highly specialized protocols designed for specific link layers. It is possible to pursue a compromise in this design space as protocol candidates, such as IMEP, can serve as a common interface mechanism between different routing/forwarding policies and link-layer interfaces.

III. Conclusions

From the summary of the mailing list and working group meeting discussions, it is clear that there are many issues facing the MANET WG. Its charter is to develop IP routing technology, but this must be done with an eye towards developing an architecture well-suited for supporting the future standards efforts, and of achieving interoperability with the current and future Internet.

Book Review (cont.)

(Continued from page 6)

In CDPD, mobile end systems are assigned fixed addresses but the subnetwork point of attachment can change. The two protocols that are defined to support network layer mobility for CDPD are: (a) MNLP (mobile network location protocol) and (b) MNRP (mobile network registration protocol). These protocols relies on the concept of a home Mobile Data Intermediate Systems (MD-IS).

Issues related to CDPF deployment and operation are reported in Chapter 11. Although this chapter sounds interesting, its contents are lacking and it may perhaps be made more practical, i.e., to report practical deployment issues and cases. Presentation of unexpected performance results or technical difficulties faced may probably attract readership.

Finally, instead of a conclusion, the authors presented the last chapter entitled: "Evolving CDPD Issues of the Future". Although this may form a section of the conclusion chapter, it should never be used as a substitute. A conclusion chapter is definitely necessary to recollect the main contents of the book, to re-emphasize important points, to highlight the authors' view of CDPD's future and any directions on future wireless data networks.

Overall, this book is "really" a book on CDPD architecture, interfaces, operations and standards. Readers who are interested to find out more about CDPD should consult this book. However, this book does not reveal commercial specifications of existing CDPD systems and hence the authors have provided a directory of companies and organizations (forums, vendors, carriers, service provides, etc) in the appendix for consultation by readers.