Mobile Ad Hoc Networking and the IETF

Joseph P. Macker[a] — *macker@itd.nrl.navy.mil*

M. Scott Corson[b] — *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 summarizes the proceedings of the last MANET WG meeting, presents some issues currently under discussion on the MANET mailing list (manet@itd.nrl.navy.mil)*, and gives some rationale behind the architectural design approach being promoted within the group.*

I. Proceedings of Los Angeles MANET Meeting

The MANET WG met for two hours, and the meeting began with five presentations on different topics.

A. IMEP

Phil Papadopoulos from Oak Ridge National Laboratory gave an update regarding the Internet MANET Encapsulation Protocol (IMEP), and described a group coloring mechanism for providing one-hop, point-to-multipoint, selective repeat-based reliability for broadcast traffic. The mechanism uses a combination of sequence number and color to manage a one-hop delivery neighborhood, and allows for implicit and explicit delivery.

There is no working group consensus regarding the IMEP draft—consensus only includes its authors at this phase. The intent of IMEP is to provide a network-layer support sublayer for potentially multiple MANET routing and network control policies. It provides multiple support functions, none of which are mandatory, but all of which should be useful to some upper layer protocol. The protocol's functionality and packet framing structure are extensible. It is intended to become a working group document without an explicit author list—only a list of contributors and a general editor. Most of what was presented in the talk is not in the current draft, but will appear in a revised version shortly.

B. ZRP

Marc Pearlman from Cornell gave an update regarding recent work on the Zone Routing Protocol (ZRP) and presented some simulation performance results. Several mechanisms were described including query detection (a method to terminate requests that arrive at previously queried nodes where intermediate nodes "listen in" and may terminate requests), and "selective bordercasting" which computes the set of inner nodes required to reach a given set of outer nodes covered (can be viewed as a multihop version of multipoint relaying, only the computation is NP-hard). Marc also presented a loose description of the call-to-mobility ratio, and a view on its usefulness in MANETs.

C. TCP/MAC Interaction

Mario Gerla from UCLA presented recent results in contention-based MAC layer/TCP interaction to give the group some insight as to how TCP might perform in MANETs. Tests were run on a 7-node tandem network using both CSMA/CA and FAMA with error-free channels. An interesting result is that TCP was seen to perform better with fixed (as opposed to adaptive) window sizes. The well-known TCP behavior of being unfair to long connections competing with shorter connections was shown to exist in a dramatic fashion within MANETs using contention-based MAC schemes. Future work at UCLA will examine approaches such as the Snoop protocol from Berkeley, ECN and fair queuing at the link layer to improve TCP performance.

D. MPR

Amir Qayyum from INRIA described the concept of Multi-Point Relaying—a simple mechanism for efficient flooding in MANETs (already adopted by the HIPERLAN standard)—which is proposed as an optional mode in IMEP for use by any upper layer protocol requiring efficient flooding support. Essentially, MPR is a distributed algorithm wherein nodes periodically exchange their one-hop neighbors sets thereby enabling each other to compute their corresponding two-hop neighborhoods. Nodes then select a subset of their one-hop neighbors to serve as Multipoint Relays which forward their traffic to the two-hop nodes. This enables efficient implementation of flooding in MANETs and should be used when the bandwidth savings accrued from flooding overhead reduction (i.e. when sufficient flooding-based traffic exists) justifies the cost of the periodic message exchanges required to support MPR.

E. DSR

Dave Johnson from Carnegie Mellon University presented the Dynamic Source Routing (DSR) draft. DSR is a protocol for performing on-demand source routing. The basic protocol is very general, with multiple options intended to improve performance. The protocol is described as being suitable for small networks (5 to 10 hops typical path lengths). He also presented the recent modification of DSR to support multiple interfaces per router. The addressing scheme does not follow traditional IP addressing practice—which assigns an IP address to each interface—but, rather, assigns a single IP address to a router with each of its interfaces indexed by an 8-bit number. The approach also uses IPv6-style header options and requires IPv4 kernel modification for use in current host systems acting as MANET routers. He also briefly mentioned work underway in NS-2 to support MANET routing studies that should be available for group use by summer.

F. General Discussion

During the general discussion that followed the presentations, there was group consensus that the principal reason for using IP-layer routing in MANETs was to support an IP routing topology or "fabric" consisting of multiple, heterogeneous physical layer topologies, or perhaps even multiple topologies of the same technology using different frequencys (see Figure 1a).

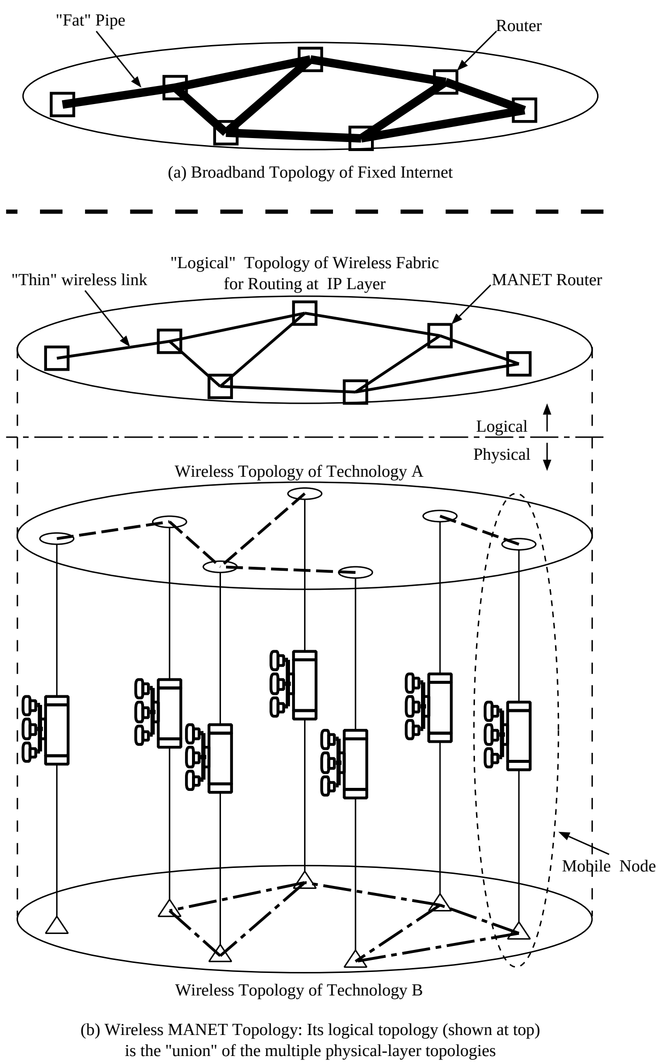

Figure 1: Fixed Internet and MANET Topologies

Discussions then followed regarding node architectures, subnet address advertisement and transport/network layer interaction. There was no consensus reached on any of these subjects before the meeting had to be closed due to a lack of time. However, the general thought regarding TCP/network interaction was the TCP modifications should be a the option of last resort, and that MANET technology should be designed to support TCP as well as possible. It was suggested that these subjects be raised on the mailing list.

II. MANET Architectural Philosophy

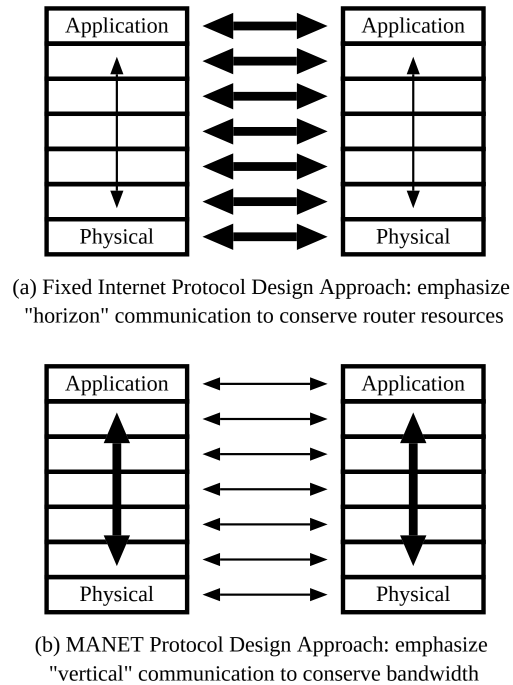

Figure 2: Fixed and MANET Protocol Design Philosophies

The traditional, fixed Internet is a network with a multi-hop topology as shown in Figure 1a. So too is the *logical* topology of a MANET as seen in Figure 1b, which can essentially be viewed as a "mobile Internet"—only in microcosm—where MANET nodes can be viewed as "mobile subnets". Both networks are resource-constrained, but the constraints *differ* in the two environments. The fixed Internet has a relatively static topology, and the bandwidth available between adjacent routers is very large, and in many networks unused (e.g. dark fiber). The principle constraints are *processing* and *storage* capacity within the routers themselves. In contrast, MANET topologies may be very *dynamic*, and it is the *bandwidth* available between routers that is very limited. This will likely always remain so relative to fixed networks. Additional, some MANET environments may involve routers which are *energy-constrained* as well. For MANETs without energy constraints, bandwidth is the most precious resource, and the routers may have a large amount of processing and storage resources to apply to each bit transmitted relative to fixed networks. For energy-constrained MANETs, energy considerations may dominate, and low power processing techniques and networking protocols would be required. It should be noted that, typically, *transmission* is the most energy consumptive activity that low-power wireless routers must perform. Thus, it can be seen that low-power design may also lead to approaches that are parsimonious in their use of bandwidth as well.

The resource constaints of the fixed Internet (what we term a "bandwidth abundant" environment) have naturally led to a protocol design approach that favors expenditure of bandwidth while minimizing—to the greatest extent possible—the need for processing or storage in routers. This design approach relies on "horizontal" peer-to-peer communication between peer protocol layers on neighboring routers (as shown in Figure 2a, while minimizing the amount of "vertical" intralayer communication within the protocol stack on a given router. This is sometimes referred to as the principle of "strict protocol layer separation". This approach has the added benefit in that it minimizes the degree of fate sharing between adjacent protocol layers, and keeps things simpler regarding protocol design.

The resource constraints in MANETs (essentially the inverse of those in the fixed Internet) argue for a different design philosophy, viz. one which minimizes horizontal communication (this expends bandwidth) and increases vertical communication within the protocol stack (see Figure 2b). Protocol stacks designed in this fashion become more logically "tightly-coupled", with increased two-way vertical communication sufficient to permit upper layer protocols to bind more tightly with lower layer protocols, thereby removing inefficienies that might result in additional horizontal communication. In this approach, the practice of designing a protocol to run atop *any* lower level protocol (i.e. sometimes referred to as "protocol independence") is likely not possible. Upper layer protocols will likely become dependent on lower layer protocols for protocol-specific functionality.

Of course, this latter design approach emphasizing tighter integration runs counter to that of the existing Internet, and the extent to which it can be realized may largely be dictated by economics, simplicity, and interoperability with the existing Internet protocols. Engineering trade-offs must be made. For example, modification of transport functionality such as TCP—while possibly desirable—may not be feasible if interoperability with the existing network is desired. In this case, the proposed design approach may only be feasible for integrating the lowest three layers (yet to be designed) in support of TCP (already fielded). This still leaves the possibility that future transport and application-level protocols can be designed in an integrated fashion—possibly incorporating Application Layer Framing (ALF) concepts—thereby achieving greater efficiency.

III. Conclusions

From the preceding, 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.

MC^{2}R is now retroactively listed in INSPEC

Produced by the Institution of Electrical Engineers, INSPEC is the leading English-language information service providing access to the world's scientific and technical literature in physics, electrical engineering, electronics, communications, control engineering, computers and computing, and information technology. It contains almost 6 million references cataloged since 1969 and is updated weekly.

More information on INSPEC is available from:
http://www.iee.org.uk/publish/inspec/inspec.html