              Mobile Ad hoc Networking (MANET):
   Routing Protocol Performance Issues and Evaluation Considerations

Status of this Memo

Copyright Notice

Abstract

   This memo first describes the characteristics of Mobile Ad hoc
   Networks (MANETs), and their idiosyncrasies with respect to
   traditional, hardwired packet networks. It then discusses the effect
   these differences have on the design and evaluation of network
   control protocols with an emphasis on routing performance evaluation
   considerations.

1. Introduction

   With recent performance advancements in computer and wireless
   communications technologies, advanced mobile wireless computing is
   expected to see increasingly widespread use and application, much of
   which will involve the use of the Internet Protocol (IP) suite. The
   vision of mobile ad hoc networking is to support robust and efficient
   operation in mobile wireless networks by incorporating routing
   functionality into mobile nodes. Such networks are envisioned to
   have dynamic, sometimes rapidly-changing, random, multihop topologies
   which are likely composed of relatively bandwidth-constrained
   wireless links.

   Within the Internet community, routing support for mobile hosts is
   presently being formulated as "mobile IP" technology. This is a
   technology to support nomadic host "roaming", where a roaming host
   may be connected through various means to the Internet other than its
   well known fixed-address domain space. The host may be directly
   physically connected to the fixed network on a foreign subnet, or be

connected via a wireless link, dial-up line, etc. Supporting this
form of host mobility (or nomadicity) requires address management,
protocol interoperability enhancements and the like, but core network
functions such as hop-by-hop routing still presently rely upon pre-
existing routing protocols operating within the fixed network. In
contrast, the goal of mobile ad hoc networking is to extend mobility
into the realm of autonomous, mobile, wireless domains, where a set
of nodes--which may be combined routers and hosts--themselves form
the network routing infrastructure in an ad hoc fashion.

2. Applications

The technology of Mobile Ad hoc Networking is somewhat synonymous
with Mobile Packet Radio Networking (a term coined via during early
military research in the 70's and 80's), Mobile Mesh Networking (a
term that appeared in an article in The Economist regarding the
structure of future military networks) and Mobile, Multihop, Wireless
Networking (perhaps the most accurate term, although a bit
cumbersome).

There is current and future need for dynamic ad hoc networking
technology. The emerging field of mobile and nomadic computing, with
its current emphasis on mobile IP operation, should gradually broaden
and require highly-adaptive mobile networking technology to
effectively manage multihop, ad hoc network clusters which can
operate autonomously or, more than likely, be attached at some
point(s) to the fixed Internet.

Some applications of MANET technology could include industrial and
commercial applications involving cooperative mobile data exchange.
In addition, mesh-based mobile networks can be operated as robust,
inexpensive alternatives or enhancements to cell-based mobile network
infrastructures. There are also existing and future military
networking requirements for robust, IP-compliant data services within
mobile wireless communication networks [1]--many of these networks
consist of highly-dynamic autonomous topology segments. Also, the
developing technologies of "wearable" computing and communications
may provide applications for MANET technology. When properly combined
with satellite-based information delivery, MANET technology can
provide an extremely flexible method for establishing communications
for fire/safety/rescue operations or other scenarios requiring
rapidly-deployable communications with survivable, efficient dynamic
networking. There are likely other applications for MANET technology
which are not presently realized or envisioned by the authors. It
is, simply put, improved IP-based networking technology for dynamic,
autonomous wireless networks.

3. Characteristics of MANETs

   A MANET consists of mobile platforms (e.g., a router with multiple
   hosts and wireless communications devices)--herein simply referred to
   as "nodes"--which are free to move about arbitrarily. The nodes may
   be located in or on airplanes, ships, trucks, cars, perhaps even on
   people or very small devices, and there may be multiple hosts per
   router. A MANET is an autonomous system of mobile nodes. The system
   may operate in isolation, or may have gateways to and interface with
   a fixed network. In the latter operational mode, it is typically
   envisioned to operate as a "stub" network connecting to a fixed
   internetwork. Stub networks carry traffic originating at and/or
   destined for internal nodes, but do not permit exogenous traffic to
   "transit" through the stub network.

   MANET nodes are equipped with wireless transmitters and receivers
   using antennas which may be omnidirectional (broadcast), highly-
   directional (point-to-point), possibly steerable, or some combination
   thereof. At a given point in time, depending on the nodes' positions
   and their transmitter and receiver coverage patterns, transmission
   power levels and co-channel interference levels, a wireless
   connectivity in the form of a random, multihop graph or "ad hoc"
   network exists between the nodes. This ad hoc topology may change
   with time as the nodes move or adjust their transmission and
   reception parameters.

   MANETs have several salient characteristics:

      1) Dynamic topologies: Nodes are free to move arbitrarily; thus,
      the network topology--which is typically multihop--may change
      randomly and rapidly at unpredictable times, and may consist of
      both bidirectional and unidirectional links.

      2) Bandwidth-constrained, variable capacity links: Wireless links
      will continue to have significantly lower capacity than their
      hardwired counterparts. In addition, the realized throughput of
      wireless communications--after accounting for the effects of
      multiple access, fading, noise, and interference conditions,
      etc.--is often much less than a radio's maximum transmission rate.

      One effect of the relatively low to moderate link capacities is
      that congestion is typically the norm rather than the exception,
      i.e. aggregate application demand will likely approach or exceed
      network capacity frequently. As the mobile network is often simply
      an extension of the fixed network infrastructure, mobile ad hoc
      users will demand similar services. These demands will continue to
      increase as multimedia computing and collaborative networking
      applications rise.

3) Energy-constrained operation: Some or all of the nodes in a
MANET may rely on batteries or other exhaustible means for their
energy. For these nodes, the most important system design criteria
for optimization may be energy conservation.

4) Limited physical security: Mobile wireless networks are
generally more prone to physical security threats than are fixed-
cable nets. The increased possibility of eavesdropping, spoofing,
and denial-of-service attacks should be carefully considered.
Existing link security techniques are often applied within
wireless networks to reduce security threats. As a benefit, the
decentralized nature of network control in MANETs provides
additional robustness against the single points of failure of more
centralized approaches.

In addition, some envisioned networks (e.g. mobile military networks
or highway networks) may be relatively large (e.g. tens or hundreds
of nodes per routing area). The need for scalability is not unique
to MANETS. However, in light of the preceding characteristics, the
mechanisms required to achieve scalability likely are.

These characteristics create a set of underlying assumptions and
performance concerns for protocol design which extend beyond those
guiding the design of routing within the higher-speed, semi-static
topology of the fixed Internet.

4. Goals of IETF Mobile Ad Hoc Network (manet) Working Group

The intent of the newly formed IETF manet working group is to develop
a peer-to-peer mobile routing capability in a purely mobile, wireless
domain. This capability will exist beyond the fixed network (as
supported by traditional IP networking) and beyond the one-hop fringe
of the fixed network.

The near-term goal of the manet working group is to standardize one
(or more) intra-domain unicast routing protocol(s), and related
network-layer support technology which:

   * provides for effective operation over a wide range of mobile
   networking "contexts" (a context is a set of characteristics
   describing a mobile network and its environment);

   * supports traditional, connectionless IP service;

   * reacts efficiently to topological changes and traffic demands
   while maintaining effective routing in a mobile networking
   context.

The working group will also consider issues pertaining to addressing,
security, and interaction/interfacing with lower and upper layer
protocols. In the longer term, the group may look at the issues of
layering more advanced mobility services on top of the initial
unicast routing developed. These longer term issues will likely
include investigating multicast and QoS extensions for a dynamic,
mobile area.

5. IP-Layer Mobile Routing

An improved mobile routing capability at the IP layer can provide a
benefit similar to the intention of the original Internet, viz. "an
interoperable internetworking capability over a heterogeneous
networking infrastructure". In this case, the infrastructure is
wireless, rather than hardwired, consisting of multiple wireless
technologies, channel access protocols, etc. Improved IP routing and
related networking services provide the glue to preserve the
integrity of the mobile internetwork segment in this more dynamic
environment.

In other words, a real benefit to using IP-level routing in a MANET
is to provide network-level consistency for multihop networks
composed of nodes using a *mixture* of physical-layer media; i.e. a
mixture of what are commonly thought of as subnet technologies. A
MANET node principally consists of a router, which may be physically
attached to multiple IP hosts (or IP-addressable devices), which has
potentially *multiple* wireless interfaces--each interface using a
different wireless technology. Thus, a MANET node with interfaces
using technologies A and B can communicate with any other MANET node
possessing an interface with technology A or B. The multihop
connectivity of technology A forms a physical-layer multihop
topology, the multihop connectivity of technology B forms *another*
physical-layer topology (which may differ from that of A's topology),
and the *union* of these topologies forms another topology (in graph
theoretic terms--a multigraph), termed the "IP routing fabric", of
the MANET. MANET nodes making routing decisions using the IP fabric
can intercommunicate using either or both physical-layer topologies
simultaneously. As new physical-layer technologies are developed,
new device drivers can be written and another physical-layer multihop
topology can be seamlessly added to the IP fabric. Likewise, older
technologies can easily be dropped. Such is the functionality and
architectural flexibility that IP-layer routing can support, which
brings with it hardware economies of scale.

The concept of a "node identifier" (separate and apart from the
concept of an "interface identifier") is crucial to supporting the
multigraph topology of the routing fabric. It is what *unifies* a set
of wireless interfaces and identifies them as belonging to the same

   mobile platform. This approach permits maximum flexibility in
   address assignment. Node identifiers are used at the IP layer for
   routing computations.

5.1. Interaction with Standard IP Routing

   In the near term, it is currently envisioned that MANETs will
   function as *stub* networks, meaning that all traffic carried by
   MANET nodes will either be sourced or sinked within the MANET.
   Because of bandwidth and possibly power constraints, MANETs are not
   presently envisioned to function as *transit* networks carrying
   traffic which enters and then leaves the MANET (although this
   restriction may be removed by subsequent technology advances). This
   substantially reduces the amount of route advertisement required for
   interoperation with the existing fixed Internet. For stub operation,
   routing interoperability in the near term may be achieved using some
   combination of mechanisms such as MANET-based anycast and mobile IP.
   Future interoperability may be achieved using mechanisms other than
   mobile IP.

   Interaction with Standard IP Routing will be greatly facilitated by
   usage of a common MANET addressing approach by all MANET routing
   protocols. Development of such an approach is underway which permits
   routing through a multi-technology fabric, permits multiple hosts per
   router and ensures long-term interoperability through adherence to
   the IP addressing architecture. Supporting these features appears
   only to require identifying host and router interfaces with IP
   addresses, identifying a router with a separate Router ID, and
   permitting routers to have multiple wired and wireless interfaces.

6. MANET Routing Protocol Performance Issues

   To judge the merit of a routing protocol, one needs metrics--both
   qualitative and quantitative--with which to measure its suitability
   and performance. These metrics should be *independent* of any given
   routing protocol.

   The following is a list of desirable qualitative properties of MANET
   routing protocols:

      1) Distributed operation: This is an essential property, but it
      should be stated nonetheless.

      2) Loop-freedom: Not required per se in light of certain
      quantitative measures (i.e. performance criteria), but generally
      desirable to avoid problems such as worst-case phenomena, e.g. a
      small fraction of packets spinning around in the network for
      arbitrary time periods. Ad hoc solutions such as TTL values can

bound the problem, but a more structured and well-formed approach
is generally desirable as it usually leads to better overall
performance.

3) Demand-based operation: Instead of assuming an uniform traffic
distribution within the network (and maintaining routing between
all nodes at all times), let the routing algorithm adapt to the
traffic pattern on a demand or need basis. If this is done
intelligently, it can utilize network energy and bandwidth
resources more efficiently, at the cost of increased route
discovery delay.

4) Proactive operation: The flip-side of demand-based operation.
In certain contexts, the additional latency demand-based operation
incurs may be unacceptable. If bandwidth and energy resources
permit, proactive operation is desirable in these contexts.

5) Security: Without some form of network-level or link-layer
security, a MANET routing protocol is vulnerable to many forms of
attack. It may be relatively simple to snoop network traffic,
replay transmissions, manipulate packet headers, and redirect
routing messages, within a wireless network without appropriate
security provisions. While these concerns exist within wired
infrastructures and routing protocols as well, maintaining the
"physical" security of of the transmission media is harder in
practice with MANETs. Sufficient security protection to prohibit
disruption of modification of protocol operation is desired. This
may be somewhat orthogonal to any particular routing protocol
approach, e.g. through the application of IP Security techniques.

6) "Sleep" period operation: As a result of energy conservation,
or some other need to be inactive, nodes of a MANET may stop
transmitting and/or receiving (even receiving requires power) for
arbitrary time periods. A routing protocol should be able to
accommodate such sleep periods without overly adverse
consequences. This property may require close coupling with the
link-layer protocol through a standardized interface.

7) Unidirectional link support: Bidirectional links are typically
assumed in the design of routing algorithms, and many algorithms
are incapable of functioning properly over unidirectional links.
Nevertheless, unidirectional links can and do occur in wireless
networks. Oftentimes, a sufficient number of duplex links exist so
that usage of unidirectional links is of limited added value.
However, in situations where a pair of unidirectional links (in
opposite directions) form the only bidirectional connection
between two ad hoc regions, the ability to make use of them is
valuable.

The following is a list of quantitative metrics that can be used to assess the performance of any routing protocol.

   1) End-to-end data throughput and delay: Statistical measures of data routing performance (e.g., means, variances, distributions) are important. These are the measures of a routing policy's effectiveness--how well it does its job--as measured from the *external* perspective of other policies that make use of routing.

   2) Route Acquisition Time: A particular form of *external* end-to-end delay measurement--of particular concern with "on demand" routing algorithms--is the time required to establish route(s) when requested.

   3) Percentage Out-of-Order Delivery: An external measure of connectionless routing performance of particular interest to transport layer protocols such as TCP which prefer in-order delivery.

   4) Efficiency: If data routing effectiveness is the external measure of a policy's performance, efficiency is the *internal* measure of its effectiveness. To achieve a given level of data routing performance, two different policies can expend differing amounts of overhead, depending on their internal efficiency. Protocol efficiency may or may not directly affect data routing performance. If control and data traffic must share the same channel, and the channel's capacity is limited, then excessive control traffic often impacts data routing performance.

   It is useful to track several ratios that illuminate the *internal* efficiency of a protocol in doing its job (there may be others that the authors have not considered):

      * Average number of data bits transmitted/data bit delivered--this can be thought of as a measure of the bit efficiency of delivering data within the network. Indirectly, it also gives the average hop count taken by data packets.

      * Average number of control bits transmitted/data bit delivered--this measures the bit efficiency of the protocol in expending control overhead to delivery data. Note that this should include not only the bits in the routing control packets, but also the bits in the header of the data packets. In other words, anything that is not data is control overhead, and should be counted in the control portion of the algorithm.

   * Average number of control and data packets transmitted/data
   packet delivered--rather than measuring pure algorithmic
   efficiency in terms of bit count, this measure tries to capture
   a protocol's channel access efficiency, as the cost of channel
   access is high in contention-based link layers.

Also, we must consider the networking *context* in which a protocol's
performance is measured. Essential parameters that should be varied
include:

   1) Network size--measured in the number of nodes

   2) Network connectivity--the average degree of a node (i.e. the
   average number of neighbors of a node)

   3) Topological rate of change--the speed with which a network's
   topology is changing

   4) Link capacity--effective link speed measured in bits/second,
   after accounting for losses due to multiple access, coding,
   framing, etc.

   5) Fraction of unidirectional links--how effectively does a
   protocol perform as a function of the presence of unidirectional
   links?

   6) Traffic patterns--how effective is a protocol in adapting to
   non-uniform or bursty traffic patterns?

   7) Mobility--when, and under what circumstances, is temporal and
   spatial topological correlation relevant to the performance of a
   routing protocol? In these cases, what is the most appropriate
   model for simulating node mobility in a MANET?

   8) Fraction and frequency of sleeping nodes--how does a protocol
   perform in the presence of sleeping and awakening nodes?

A MANET protocol should function effectively over a wide range of
networking contexts--from small, collaborative, ad hoc groups to
larger mobile, multihop networks. The preceding discussion of
characteristics and evaluation metrics somewhat differentiate MANETs
from traditional, hardwired, multihop networks. The wireless
networking environment is one of scarcity rather than abundance,
wherein bandwidth is relatively limited, and energy may be as well.

In summary, the networking opportunities for MANETs are intriguing
and the engineering tradeoffs are many and challenging. A diverse
set of performance issues requires new protocols for network control.

A question which arises is "how should the *goodness* of a policy be
measured?". To help answer that, we proposed here an outline of
protocol evaluation issues that highlight performance metrics that
can help promote meaningful comparisons and assessments of protocol
performance. It should be recognized that a routing protocol tends
to be well-suited for particular network contexts, and less well-
suited for others. In putting forth a description of a protocol, both
its *advantages* and *limitations* should be mentioned so that the
appropriate networking context(s) for its usage can be identified.
These attributes of a protocol can typically be expressed
qualitatively, e.g., whether the protocol can or cannot support
shortest-path routing. Qualitative descriptions of this nature
permit broad classification of protocols, and form a basis for more
detailed *quantitative* assessments of protocol performance. In
future documents, the group may put forth candidate recommendations
regarding protocol design for MANETs. The metrics and the philosophy
presented within this document are expected to continue to evolve as
MANET technology and related efforts mature.

7. Security Considerations

   Mobile wireless networks are generally more prone to physical
   security threats than are fixed, hardwired networks. Existing link-
   level security techniques (e.g. encryption) are often applied within
   wireless networks to reduce these threats. Absent link-level
   encryption, at the network layer, the most pressing issue is one of
   inter-router authentication prior to the exchange of network control
   information. Several levels of authentication ranging from no
   security (always an option) and simple shared-key approaches, to full
   public key infrastructure-based authentication mechanisms will be
   explored by the group. As an adjunct to the working groups efforts,
   several optional authentication modes may be standardized for use in
   MANETs.

8. References

   [1] Adamson, B., "Tactical Radio Frequency Communication Requirements
       for IPng", RFC 1677, August 1994.

Authors' Addresses

   M. Scott Corson
   Institute for Systems Research
   University of Maryland
   College Park, MD 20742

   Phone: (301) 405-6630
   EMail: corson@isr.umd.edu

   Joseph Macker
   Information Technology Division
   Naval Research Laboratory
   Washington, DC 20375

   Phone: (202) 767-2001
   EMail: macker@itd.nrl.navy.mil

Full Copyright Statement