IETF MANET Working Group   Mario Gerla
INTERNET-DRAFT   Xiaoyan Hong
Expiration: May 17, 2001   Li Ma
                           University of California, Los Angeles
                                                      Guangyu Pei
                                           Rockwell Science Center
                                              November 17, 2000

Landmark Routing Protocol (LANMAR) for Large Scale Ad Hoc Networks

<draft-ietf-manet-lanmar-00.txt>

Status of This Memo

Abstract

   The Landmark Routing Protocol (LANMAR) utilizes the concept of
   "landmark" for scalable routing in large, mobile ad hoc networks.
   It relies on the notion of group mobility: i.e., a logical group
   (for example a team of coworkers at a convention) moves in a

coordinated fashion. The network address of a node is <Group ID,
Host ID>. A landmark is dynamically elected in each group. The
route to a landmark is propagated throughout the network using a
Distance Vector mechanism. Separately, each node in the network
uses a "scoped" routing algorithm (e.g., FSR) to learn about routes
within a given (max number of hops) scope. To route a packet to a
destination outside its scope, a node will direct the packet to the
landmark corresponding to the group ID of such destination. Once
the packet is within the scope of the landmark, it will typically
be routed directly to the destination. Remote groups of nodes are
"summarized" by the corresponding landmarks. The solution to
drifters (i.e., nodes outside of the scope of their landmark) is
also handled by LANMAR. Landmark dynamic election enables LANMAR
to cope with mobile environments. By using a "scoped" routing
scheme, we dramatically reduce routing table size and update
overhead in large nets. LANMAR is well suited to provide an
efficient and scalable routing solution in large, mobile, ad hoc
environments in which group behavior applies and high mobility
renders traditional routing schemes inefficient.

                              Contents

1. Introduction

   This document describes the Landmark Routing Protocol (LANMAR) [1,2]
   developed by the Wireless Adaptive Mobility (WAM) Laboratory [4] at
   University of California, Los Angeles.

   The concept of landmark routing was first introduced in wired area
   networks [6]. The original scheme required predefined multi-level
   hierarchical addressing. The hierarchical address of each node
   reflects its position within the hierarchy and helps to find a route
   to it. Each node knows the routes to all the nodes within its
   hierarchical partition. Moreover, each node knows the routes to
   various "landmarks" at different hierarchical levels. Packet
   forwarding is consistent with the landmark hierarchy and the path
   is gradually refined from the top level hierarchy to lower levels
   as a packet approaches its destination.

   LANMAR borrows the concept of landmark and extends it to the
   wireless ad hoc environment. LANMAR scheme does not require
   predefined hierarchical address, but it uses the notion of landmarks
   to keep track of logical subnets in which the members have a
   commonality of interests and are likely to move as a "group"
   (e.g., brigade in the battlefield, colleagues in the same
   organization, a group of students from same class, or a team of
   co-workers at a convention and a tank battalion in the battlefield).

Each such logical group has an elected landmark. For each group,
underlining "scoped" routing algorithm will provide a one-level
scope. The routing update packets are restricted only within the
scope. Accurate routing information for nodes within scope is
maintained. The routing information to remote nodes (nodes outside
the node's scope) is "summarized" by the corresponding landmarks.
Thus, the LANMAR scheme largely reduces the routing table size and
the routing update traffic overhead. It greatly improves
scalability.

In addition, in order to recover from landmark failures,
a "landmark" node is elected in each subnet. Landmark election
provides a flexible way for the LANMAR protocol to cope with a
dynamic and mobile network. The protocol also provides a solution
for drifters (Nodes that are outside the fisheye scopes of the
landmarks of their logical groups). Extra storage, processing and
line overhead will be incurred for landmark election and drifter
bookkeeping. However, the design of the algorithms provides
solutions without compromising scalability. For example, the
routing overhead for handling drifters is typically small if the
fraction of drifting nodes is small. More analysis is given in
Section 7.

The LANMAR runs on top of a proactive routing protocol. It also
requires that the underlining routing protocol support the scoped
subnetworking. Fisheye State Routing Protocol (FSR) [7,8] is
such a protocol that supports LANMAR. In FSR, the link state
protocol is used within the scope. However, the fisheye scope
technique can also be applied to a distance vector type protocol,
such as DSDV [3], in which the hop distance can be used to
bind the scope for routing message updating. The main advantage of
LANMAR is that the routing table includes only the nodes within the
scope and the landmark nodes. This feature greatly improves
scalability by reducing routing table size and update traffic O/H.

Thus the Landmark Routing Protocol provides an efficient and
scalable routing solution for a mobile, ad hoc environment while
keeping line and storage overhead (O/H) low. Moreover, the
election provides a much needed recovery from landmark failures.

2. Terminology

2.1. General Terms

This section defines terminology used in LANMAR.

   node

A MANET router that implements Landmark Routing Protocol.

neighbor

Nodes that are within the radio transmission range.

scope

A zone that is bounded by hop distances.

host protocol

A proactive protocol underlines the Landmark Routing Protocol.

subnet

Logical groups of nodes that present similar motion behavior.

group

This is used interchangeably with subnet.

2.2. Specification Language

The keywords "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 [5].

3. Protocol Applicability

3.1. Networking Context

LANMAR is best suited for large scale mobile ad hoc wireless networks. The landmark scheme on top of "scoped" routing algorithm has large advantages in reducing routing update packet size and keeping progressive accurate routes to remote nodes. It achieves high data packet to routing packet ratio. Moreover, the fact that the route error is weighted by distance obviously reduces the sensitivity to network size.

LANMAR is also suited for high mobility ad hoc wireless networks. This is because in a mobility environment, a change on a link far away from the source does not necessarily cause a change in the routing table at the source and that all the information about remote nodes is summarized by landmarks.

[3.2](#). Protocol Characteristics and Mechanisms

* Does the protocol provide support for unidirectional links?(if so, how?)

     Yes. Since LANMAR is on top of a host protocol, if the host
     protocol supports unidirectional links, e.g., the FSR protocol,
     LANMAR can use the unidirectional link states as well.

* Does the protocol require the use of tunneling? (if so, how?)
     No.

* Does the protocol require using some form of source routing? (if so, how?)

     No.

* Does the protocol require the use of periodic messaging? (if so, how?)

     Yes. The LANMAR periodically broadcasts landmark information
     to its neighbors. But the updates are piggybacked in the host
     routing update messages.

* Does the protocol require the use of reliable or sequenced packet delivery? (if so, how?)

     No. As the packets are sent periodically, they need not be
     sent reliably.

* Does the protocol provide support for routing through a multi-technology routing fabric? (if so, how?)

     Yes. It is assumed that each node's network interface is
     assigned a unique IP address.

* Does the protocol provide support for multiple hosts per router? (if so, how?)

     Yes. The router that has multiple hosts can use network ID of
     these hosts as the address to participate LANMAR.

* Does the protocol support the IP addressing architecture? (if so, how?)

     Yes. Each node is assumed to have a unique IP address (or
     set of unique IP addresses in the case of multiple interfaces).
     The LANMAR references all nodes/interfaces by their IP address.

This version of the LANMAR also supports IP network addressing (network prefixes) for routers that provide access to a network of non-router hosts.

* Does the protocol require link or neighbor status sensing (if so, how?)

No.

* Does the protocol have dependence on a central entity? (if so, how?)

No.

* Does the protocol function reactively? (if so, how?)

No.

* Does the protocol function proactively? (if so, how?)

Yes. The LANMAR proactively maintains landmark information at each node.

* Does the protocol provide loop-free routing? (if so, how?)

Yes. As the protocol uses routing paths learned from the host protocol for in-scope destinations, if the host protocol provides loop-free routing, so does LANMAR. For out-scope destinations, only routes to landmarks are used. Because these routes are DSDV, it is loop free. When a packet is approaching the destination, in-scope routes are used again.

* Does the protocol provide for sleep period operation?(if so, how?)

Yes. However, this requires TDMA MAC layer support. The router can be scheduled to sleep during idle periods.

* Does the protocol provide some form of security? (if so, how?)

Yes. When a node broadcasts routing update message, only entries of in-scope nodes and landmarks are included. This will prevent other remote nodes from being heard.

* Does the protocol provide support for utilizing multi-channel, link-layer technologies? (if so, how?)

Yes. In fact, the multi-channel can be used to separate routing messages from user data packets.

4. Protocol Overview

4.1. Protocol Descriptions

As mentioned in Section 1, the landmark concept we adopt here uses the notion of logical subnets in which the members have a commonality of interests and are likely to move as a "group". Each logical subnet has one node serving as a "landmark" of that subnet. The protocol requires that the landmark of each subnet have the knowledge of all the members in its group. The LANMAR protocol also uses a scope at each node. The size of the scope is a parameter measured in hop distance. It is chosen in such a way that if a node is at the center of a subnet, the scope will cover the majority of the subnet members. If the shape of a subnet is likely to be a cycle, the center node's scope will cover all the members of the subnet. If this center node is elected as a landmark, it fulfills the requirement of the protocol. If a landmark does not locate at the center, there will be some members drifting off its scope. The landmark will keep records for these drifters.

The LANMAR relies on an underlining proactive protocol with the ability of providing "scoped" routing. The LANMAR itself is a distance vector type protocol. Each node maintains a distance vector for landmarks of all the subnets and a distance vector for drifters in its subnet. The distance vectors for landmarks and drifters are exchanged through piggybacking in the periodical routing update packets (link state or distance vector) of the underling routing protocol. The size of the landmark distance vector equals to the number of logical subnets and thus landmark nodes. Both sizes of the two distance vectors are small.

The routing update scheme uses the "scoped" routing algorithm. The main idea is summarized below. Each node broadcasts routing information periodically to its immediate neighbors. In each update packet, the node sends routing table entries within its scope and also piggybacks the landmark and drifter distance vectors. Sequence numbers are used in LANMAR update packets. Each node advances its sequence number before sending an update packet. Through the exchange process, the table entries with larger sequence numbers replace the ones with smaller sequence numbers.

Let's take, for example, the FSR as our host protocol. For nodes within the scope, the updating is in a certain frequency. But for nodes beyond the scope, the update frequency is reduced to zero; Only the update frequency of the landmark nodes remains

unaltered. As a result, each node maintains accurate routing
information for in-scope nodes and keep routing directions to the
landmark nodes for out-scope nodes, or say, for remote groups.
A packet directed to a remote destination initially aims at the
landmark of that remote group; as it gets closer to the landmark,
it may eventually switch to the accurate route to the destination
provided by in-scope nodes of the destination.

4.2. Landmark Election

Dynamic election/re-election of landmark node is essential for
LANMAR to work in a wireless mobile environment. Basically, each
node tracks other nodes of its group in its scope and computes
"weight", e.g. the number of the nodes it has found. At the
beginning of the LANMAR, no landmark exists. Protocol LANMAR
only uses the host protocol functionality. As host routing
computation progresses, one of the nodes will learn (from
the host protocol's routing table) that more than a certain number
of group members (say, T) are in its scope. It then proclaims
itself as a landmark for this group and adds itself to the landmark
distance vector. The node broadcasts the landmark information to
the neighbors jointly with the host update packets.

When more than one node declares itself as a landmark in the same
group, as the landmark information floods out, each node will
perform a winner competition procedure. Only one landmark for each
group will survive and it will be elected. To avoid flapping
between landmarks (very possible in a mobile situation), we use
hysteresis in the replacement of an existing landmark. I.e., the
old Landmark is replaced by the new one only if its weight is, say
less than 1/2 of the weight of the current election winner. Once
ousted, the old leader needs the full weight superiority to be
reinstated.

This procedure is carried out periodically in the background (low
overhead, anyway). At steady state, a landmark propagates its
presence to all other nodes like a sink in DSDV. It is extremely
simple and it converges (by definition). In a mobile environment,
an elected landmark may eventually lose its role. The role shifting
is a frequent event. In a transient period, there exist several
landmarks in a single group. The transient period may be actually
the norm at high mobility. This transient behavior can be
drastically reduced by using hysteresis.

When a landmark dies, its neighbors will detect the silence after
a given timeout. The neighbors of the same group will then take
the responsibilities as landmarks and broadcast new landmark
information. A new round of landmark election then floods over

the entire network.

4.3. Drifters

Typically, all members in a logical subnet are within the scope of
the landmark, thus the landmark has a route to all members. It may
happen, however, that some of the members "drift off" the scope,
for example, a tank in a battalion may become stranded or lost.
To keep track of such "drifters", i.e., to make the route to them
known to the landmark, the following modification to the routing
table exchange is necessary. Each node, say i, on the shortest
path between a landmark L and a drifter l associated with that
landmark keeps a distance vector entry to l. Note that if l is
within the scope of i, this entry is already included in the
routing table of node i. When i transmits its distance vector to
neighbor, say j, then j will retain the entry to member l only if
$d(j,l)$ is smaller than the scope or $d(j,L)$ is smaller than $d(i,L)$.
The latter condition occurs if j is on the shortest path from i
(and therefore from l) to L. This way, a path is maintained from
the landmark to each of its members, including drifters.

The occurrences of drifters are dynamic in a mobile network. In
order to timely remove the staled drifter information, the time
when a node detects itself a drifter is propagated with the
distance vector.

5. Protocol Specifications

This section discusses the operation of LANMAR routing protocol.
The sending and receiving of landmark updates are in the proactive
nature. The routing packets are processed separately from
ordinary data packets.

5.1. Data Structures

Each node has a unique "logical" identifier defined by a subnet
field and a host field. The host field is unique in the subnet and
might in fact coincide with the physical address. The "logical"
identifier can also be an IP address when the subnet address can
logically group the nodes. Moreover, each node keeps a landmark
status as part of identifier. As LANMAR runs on top of a host
routing protocol, it shares the underlining data structures.
Particularly, LANMAR requires the underlining data structures to
associate a landmark status with each node's identifier. In
addition, LANMAR keeps a drifter list and a landmark distance
vector.

5.1.1. Landmark Status

Each node has not only a "logical" identifier, which basically is
its address, but also a landmark status. The status includes
a flag which indicates whether the node is a landmark or not and a
weight (the number of group members the node detects within its
scope). Anywhere in the tables of the host protocol, when a
node's address is kept, a landmark status is recorded. The status
is piggybacked in the periodical routing update packets of the
host protocols together with the node address. There are two
fields for a status:
- Landmark flag
- Number of group members in its scope

5.1.2. Landmark Distance Vector

Landmark distance vector (LMDV) gives the next hop information to
all landmarks in the network. Every subnet has an entry in LMDV.
Initially, the entries in the underlining routing table, which
point to landmarks, are copied to LMDV. The LMDV entries are
updated when landmark update message is received or underlining
topology table is changed. LMDV functions as a part of the
routing table. It has the following fields:
- Landmark status
- Next hop address
- Distance

5.1.3. Drifter List

The drifter list of LANMAR provides the next hop information to
forward the packets to the drifters known to the current node.
The entries are updated explicitly using landmark update message.
It works as a part of routing table. The drifter list (DFDV) has
following fields:
- Destination drifter address
- Next hop address
- Last declaration time

5.2. LANMAR Update Message Format

The messages necessary for LANMAR protocol are piggybacked in the
host periodic update packets. The format given below is not
exactly how the bits appear in the host update packets. Where to
put these fields in the host update packets is an implementation
issue. The necessary fields include current node's landmark
status (assuming that the node's address can be obtained from IP
header), its landmark distance vector LMDV and the drifter list
DFDV. The next two subsections will describe how these

information is processed.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Landmark Flag |   N_members   |    Reserved   |  N_landmarks  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Landmark Address 1                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Next Hop  Address 1                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  N_members 1  |  Distance 1   |        ...                    :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                          . . .                                :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                          . . .                | N_drifters    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Drifter Address 1                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Next Hop  Address 1                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Distance 1   |            Timestamp 1        |    ...        :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                          . . .                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

5.2.1. Description of the fields

Landmark Flag and N_members

   These two fields are the landmark status. N_members is the
   number of group members in the node's scope.

Reserved

   The bits are set to '0' and are ignored on reception.

N_landmarks

   The number of entries of the landmark distance vector.

Landmark Address 1, Next Hop Address 1, N_members 1 and Distance 1

   The first entry in the landmark distance vector.
   Landmark Address 1 and N_members 1 are the destination landmark
   status.
   Next Hop Address 1 and Distance 1 is the next hop and distance

to the landmark.

These fields are repeated N_landmarks times for each entry in
landmark distance vector.

N_drifters

The number of entries of the drifter list.

Drifter Address 1, Next Hop Address 1, Distance 1, Timestamp 1

The first entry in the drifter list.
Next Hop Address 1 and Distance 1 are the next hop and distance
to the Drifter Address 1.
Timestamp 1 is the time when the node becomes a drifter.

These fields are repeated N_drifters times for each entry in
the drifter list.

The length of the message (including the host update fields) is
limited to the maximum IP packet size. In that case, multiple
packets may be required to broadcast all the topology entries for
host protocol. The landmark distance vector and drifter list will
be piggybacked in each packet with the same sequence number.

5.3. Processing Landmark Updates

Landmark update information is a part of the LANMAR periodic
routing update message, which is carried in host updating packets.
The update information includes sender's landmark status and LMDV.
Landmark update message is used for landmark election and helps
nodes build their LMDV.

5.3.1. Originating a Landmark in a Subnet

Every time a node detects a topology change, it recalculates the
number of group members in its scope. The new number of group
members is recorded in its landmark status and replaces the old
values in all the tables containing the landmark status. If this
number is greater than a threshold T, the node qualifies as a
landmark only when there is no landmark for the group according to
its LMDV, or it wins the election when competing with the existing
landmark. The landmark distance vector is updated accordingly.
If the host protocol has new shortest paths to landmarks, the new
paths are copied from the host routing table to the LMDV.
The landmark status and the LMDV will be broadcast to neighbors
with the next host update packet.

5.3.2. Receiving Landmark Updates

   When a node receives a landmark update message, it compares its
   LMDV entries with the incoming LMDV entries for each subnet. New
   landmark enrties will be copied. Competing landmarks will
   be solved through a winner competition algorithm. The winner will
   be elected as the landmark for the subnet. The node's landmark
   status and LMDV will be updated according to the outcomes of the
   competition. The distance information can be obtained either from
   the incoming LMDV entry or from host routing table. The updated
   LMDV and the node's landmark status will be propagated jointly with
   the host routing update packets.

5.3.3. Winner Competition

   When more than one node declares itself as a landmark in the same
   group, a simple solution is to let the node with the largest
   number of group members win the election and in case of tie,
   lowest ID breaks the tie. The other competing nodes defer.
   However, this method is likely to cause the oscillation of
   landmark roles between nodes.

   To use hysteresis in replacing an existing landmark, let us assume
   the competing node's number of members is M, the existing
   landmark's number of members is N and a factor value S. When M is
   greater than N*S, then the competing node replaces the existing
   landmark. Or, when N reduces to a value smaller than a threshold T,
   then it gives up the landmark role. A tie occurs when M falls
   within an interval [N*1/S, N*S], then the node with larger member
   number wins the election. If a tie occurs again with equal member
   number, i.e., M equals to N, it is broken using lowest ID. A tie
   can always be broken using lowest ID as the address is used as ID
   and it is unique.

5.4. Processing Drifter Updates

   Drifter update information is a part of the LANMAR periodical
   routing update message, which is carried in host updating packets.
   The update information is the drifter list (DFDV) of the sender.
   The computation of the DFDV at each node includes checking the node
   itself to see whether it is a drifter and recording paths to other
   drifters.

5.4.1. Originating a Drifter Entry

   By checking the distance to the landmark of its group, each node
   easily knows whether it has become a drifter. If the distance is
   larger than the scope, the node will put itself into its drifter

list and record the time instance. This drifter information
will be sent back to the landmark hop by hop along the shortest
path to it which can be learned from the LMDV. For each drifter,
only the node on its shortest path to the landmark needs to receive
its information, so before the entry is broadcast, the next hop to
landmark is attached with its entry. The DFDV will be propagated
with the next host routing update packet.

5.4.2. Receiving Drifter Updates

Upon receiving an update packet, the DFDV part is retrieved and
processed. If an entry of incoming DFDV indicates that the current
node is its next hop to the landmark, i.e., the current node is on
the drifter's shortest path to the landmark, the current node will
insert or update its drifter list. The node sending the update
packet is recorded as the next hop to the drifter. The reverse path
to the drifter is thus built up. The procedure ends when the
landmark receives the drifter entry. The updated DFDV will be
propagated with the next host routing update packet.

6. Data Packet Forwarding

Data packets are relayed hop by hop. The host protocol routing
table, drifter list and landmark distance vector are looked up
sequentially for the destination entry. If the destination is
within a node's scope, the entry can be found directly in the
routing table and the packet is forwarded to the next hop node.
Otherwise, the drifter list DFDV is searched for the destination.
If the entry is found, the packet is forwarded using the next hop
address from DFDV. If not, the logical subnet field of the
destination is retrieved and the LMDV entry of the landmark
corresponding to the destination's logical subnet is searched.
The data packet is then routed towards the landmark using the next
hop address from LMDV. The packet, however, is not necessary to
pass through the landmark. Rather, once the packet gets within the
scope of the destination on its way towards the landmark, it is
routed to the destination directly.

7. Discussion about Storage and Processing Overhead for Drifters

The routing storage and processing overhead introduced by the
distance vector extension to handle drifters is typically small
if the fraction of drifting nodes is small. Consider a network
with N nodes and L landmarks, and assume that a fraction F of the
members of each logical subnet have drifted. In the worst case,
the path length from landmark to drifter is the square root of N

(assuming a grid topology). Thus, sqrt(N) is the bound on the
number of extra routing entries required at the nodes along the
path to the drifter. The total number of extra routing entries is
sqrt(N)*L*(F*N/L) where N/L is the average logical group size.
Thus, the extra storage per node is F*sqrt(N). Now, let us assume
that the number of nodes in the scope = # of landmarks = logical
group size = sqrt(N). Then, the basic routing table overhead per
node (excluding drifters) is 3*sqrt(N). Thus, the extra overhead
caused by drifters is F/3. If 20% of the nodes in a group are
outside of the landmark scope, i.e., have drifted, the extra
routing O/H required to keep track of them is only 7%.

Acknowledgments

References

[1] G. Pei, M. Gerla and X. Hong, "LANMAR: Landmark Routing for
Large Scale Wireless Ad Hoc Networks with Group Mobility",
Proceedings of IEEE/ACM MobiHOC 2000, Boston, MA, Aug. 2000.

[2] M. Gerla, X. Hong, G. Pei, "Landmark Routing for Large Ad Hoc
Wireless Networks", Proceedings of IEEE GLOBECOM 2000,
San Francisco, CA, Nov. 2000.

[3] C.E. Perkins and P. Bhagwat, "Highly Dynamic Destination-
Sequenced Distance-Vector Routing (DSDV) for Mobile Computers,"
In Proceedings of ACM SIGCOMM'94, London, UK, Sep. 1994,
pp. 234-244.

[4] UCLA Wireless Adaptive Mobility (WAM) Laboratory.
    http://www.cs.ucla.edu/NRL/wireless

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

[6] P. F. Tsuchiya, "The Landmark Hierarchy: a new hierarchy for
routing in very large networks", Computer Communication Review,
vol.18, no.4, Aug. 1988, pp. 35-42.

[7] G. Pei, M. Gerla, and T.-W. Chen, "Fisheye State Routing:
A Routing Scheme for Ad Hoc Wireless Networks", Proceedings of
ICC 2000, New Orleans, LA, Jun. 2000.

    [8] G. Pei, M. Gerla, and T.-W. Chen, "Fisheye State Routing in
    Mobile Ad Hoc Networks", Proceedings of Workshop on Wireless
    Networks and Mobile Computing, Taipei, Taiwan, Apr. 2000.

Chair's Address

    The MANET Working Group can be contacted via its current chairs:

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

        Phone: +1 301 405-6630
        Email: corson@isr.umd.edu

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

        Phone: +1 202 767-2001
        Email: macker@itd.nrl.navy.mil

Authors' Addresses

    Questions about this document can also be directed to the authors:

        Mario Gerla
        3732F Boelter Hall
        Computer Science Department
        University of California
        Los Angeles, CA 90095-1596
        USA

        Phone: +1 310 825-4367
        Fax: +1 310 825-7578
        Email: gerla@cs.ucla.edu

        Xiaoyan Hong
        3820 Boelter Hall
        Computer Science Department

          University of California
          Los Angeles, CA 90095-1596
          USA

          Phone: +1 310 825-4623
          Fax: +1 310 825-7578
          Email: hxy@cs.ucla.edu

          Li Ma
          3803 Boelter Hall
          Computer Science Department
          University of California
          Los Angeles, CA 90095-1596
          USA

          Phone: +1 310 825-1888
          Fax: +1 310 825-7578
          Email: mary@cs.ucla.edu

          Guangyu Pei
          Rockwell Science Center
          1049 Camino Dos Rios
          P.O. Box 1085
          Thousand Oaks, CA 91358-0085
          USA

          Phone: +1 805 373-4639
          Fax: +1 805 373-4383
          Email: gpei@rsc.rockwell.com