INTERNET-DRAFT
Intended Status: Experimental
Expires: September 27, 2015

Jack Shaio
Caleb Lo
Don Broderick
The MITRE Corporation
March 26, 2015

A Common Layer 3 Interface for MANET
draft-shaio-manet-common-l3-interface-03

Abstract

This paper presents an approach that allows an algorithm to choose IP routing peers intelligently among the nodes in a MANET but does not involve any modifications to existing IP routing protocols. In addition, our approach works as a pure interface between (any) MANET radio terminal and any IP router, so nodes using our interface interoperate with nodes that do not use this interface or with nodes using different algorithms to select routing peers. This interface was prototyped and this paper includes test results for two different mobility patterns on a mobile network using OLSR for MANET routing and OSPFv2 for IP routing.

Status of this Memo

Table of Contents

1 Introduction

Many approaches to integrating a MANET network with IP routing have focused on modifying a standard IP routing protocol, usually OSPF, to reduce the number of IP routing peers selected. Key to these modifications is an algorithm that will choose intelligently, among all potential IP routing peers in the MANET, those that provide the "best" set of IP routing adjacencies. This prevents a routing protocol like OSPF from establishing too many routing adjacencies to MANET nodes that are intermittently unreachable as they move, leading to constant IP route changes. The key issue is to choose, from all MANET nodes that OSPF would accept as routing peers, a smaller set that provides good IP connectivity without too much route instability as nodes and links in the MANET change state. Examples of these approaches are in [RFC5614], [RFC5820] and [HEND1], with comparisons in [HEND2].

This paper presents an alternate approach to choosing IP routing peers among the MANET nodes [CL3]; it still allows an algorithm to choose routing peers intelligently but does not involve any modifications to existing IP routing protocols. In addition, our approach works as a pure interface between a MANET radio terminal and an attached IP router comprising a single mobile node, so nodes using our interface interoperate with nodes that do not use this interface or with nodes using different algorithms to select routing peers.

We prototyped this interface and this paper includes test results for two different mobility patterns on a mobile network using OLSR for MANET routing and OSPFv2 for IP routing. These results show our approach was effective on the cases we tested.

Since the Common Layer 3 interface approach presented here uses unmodified IP routers, it allows a wide range of commercially available IP routers to be used for the IP component of a MANET node. Our approach is easily adaptable to any SNMP-manageable MANET routing protocol. It can also be adapted, with more work, to MANET terminals

that are not SNMP manageable but have some other interface to export their knowledge of the MANET. The key advantage of this approach is that it does not depend either on the specific IP nor MANET routing protocol, so these can be changed independently of each other and integrated with this Common Layer 3 interface.

1.1 Terminology

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 [RFC2119].

2 Assumptions and Reference Model

The mobile network consists of mobile nodes, each consisting of a MANET terminal and an IP routing component. Users of the mobile network attach to the IP routing component, possibly via a local network.

For our prototype we used OSPFv2 [OSPFv2] as the IP routing component, so the remainder of the paper is tailored to OSPF but a different routing protocol could have been used with only minor adjustments.

2.1 Definitions and Acronyms

The following definitions and acronyms are used in the remainder:

MANET Terminal: The terminal comprising only MANET routing and radio access to the MANET, but not an IP routing component.

Mobile Node: A node consisting of a MANET terminal, an IP router and, optionally, a Common Layer 3 interface joining the two.

CL3: The Common Layer 3 interface described in this paper.

CL3 Node: A mobile node that has a Common Layer 3 interface.

Non-CL3 Node: A mobile node that does not have a Common Layer 3 interface.

Stray OSPF HELLO: An incoming OSPF HELLO message from a node not selected by the receiving CL3 node as a routing peer.

2.2 Reference Diagram

The remainder of this paper divides a mobile node into the three components shown in the diagram below.

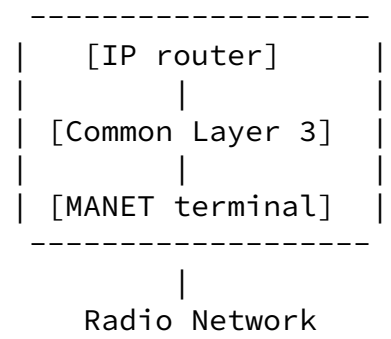

The IP router is running standard, widely available protocols without

any MANET-specific extensions. The router must have PPP and an IP routing protocol that is SNMP manageable (for example OSPFv2 used in our tests). As described below, Common Layer 3 will setup PPP sessions between itself and the router; for this it is convenient to have VLANs and PPPoE available as well.

The Common Layer 3 interface is a separate logical entity through which all packets between the router and the MANET terminal pass.

The MANET terminal runs the MANET routing protocol, which is assumed to be SNMP manageable. In our tests this was OLSR, which is a link state MANET protocol and its MIB provided a topology table.

3 Design Objectives

CL3 is a separate logical component that bridges the gap between the mobile network formed by the MANET nodes and the IP routing adjacencies created by its attached IP router. CL3 learns about the MANET network (e.g. nodes present, link qualities, MANET topology) from its attached MANET terminal, for example by making SNMP queries to the MANET routing MIB. CL3 uses this knowledge as input to its internal "network abstraction" algorithm which selects the set of IP

routing peers for the router to use. CL3 mechanisms described below will make the router use this set of routing peers, and no others, using standard features such as PPP and SNMP; no router modifications are required. There were four design objectives for CL3.

The first design objective was to use only standard and widely-available IP features on the IP router, motivated by the desire to be able to choose the IP component of the mobile node from the largest range of commercially-available routers. Our implementation of CL3 used OSPFv2 but a similar approach could instead have used OSPFv3, IS-IS, RIP or BGP. Other capabilities we used were PPPoE, VLANs and SNMP access to the table of OSPF neighbors, all widely-available features.

An example of a standard but not widely-available feature that might have been very useful is the access node control protocol [ANCP-USAGE]. This protocol allows a broadband access IP router to learn about the state of DSL lines from the DSLAM device that terminates them. Although a similar capability might have been useful in this context, this protocol is only available on a very limited set of IP routers and we avoided it.

The second design objective was to be decoupled from the MANET routing protocol as there are many of these and the field continues to evolve. All CL3 requires is that they provide some information about the MANET topology via SNMP or a comparable protocol; of course

CL3 has to be adapted to the specific MIB or protocol but otherwise does not depend on special protocol messages to/from the MANET terminal. In our tests we used Optimized Link State Routing [OLSR], a link state routing protocol that includes a topology table in its MIB. This additional information is used by our network abstraction algorithm but CL3 can easily be modified to use a different algorithm suitable for the many distance-vector MANET routing protocols in use.

The third design objective was that CL3 should not be a protocol, in other words, a CL3 node does not exchange messages with other CL3 peers. CL3 uses only standard mechanisms, SNMP and PPP primarily, to communicate with its attached IP router on one side and its attached MANET terminal on the other. This allows CL3 nodes to interoperate with non-CL3 nodes in the same MANET network. CL3 nodes can therefore be added or upgraded incrementally in the MANET network.

The fourth design objective was to allow the network abstraction algorithms in different CL3 nodes to work independently of each other and not require their calculations to result in the same set of routing peers. This means that if the network abstraction algorithm at node A selects node B as a routing peer, it is not necessary for the algorithm at node B to select node A in order for a routing adjacency to be formed between A and B. This design goal allows CL3 nodes with different network abstraction algorithms to interoperate.

In our prototype we used OSPF as the IP routing protocol and we met design goals 3 and 4 by passing "Stray OSPF HELLOs" (incoming OSPF HELLO messages not sent by a selected routing peer) to the network abstraction algorithm for consideration. The algorithm is not required to accept the sending node as a new routing peer but it may decide to do so. This allows a routing adjacency to be formed between nodes A and B, even when node A selected node B as a routing peer but B is not a CL3 node or the network abstraction algorithm at node B did not originally select node A.

These four design objectives were met in the Common Layer 3 interface architecture described below. We implemented a prototype of this architecture and present test results on mobile 40 node networks in Section 6.

4 IP Unicast Forwarding in the Common Layer 3 Interface

The function of each CL3 subcomponent can be derived from the way IP unicast packet forwarding is done at the CL3 interface between the IP router and the MANET terminal. The extension to IP multicast forwarding is described in Section 7.

Every node A on the MANET reachable from a given CL3 node B is a

potential IP routing peer for the attached router. Now nodes A and B might never establish a routing adjacency. For example, if OSPF is the routing protocol and the nodes have different HELLO timer settings, this difference in settings between nodes cannot be discovered by CL3 before the nodes are selected as routing peers. It can only be discovered after the routing adjacency fails to reach FULL state (in the OSPF case).

If a remote node is selected as a routing peer (and the routing adjacency forms), it will then be an IP next-hop for some IP routes in the router's forwarding table. The problem for CL3 is to identify the IP next-hop chosen by its attached IP router for each IP packet it forwards on the CL3 interface towards the MANET.

CL3 solved this problem by assigning a separate PPP session for each remote node selected as a routing peer. The endpoints of this PPP session are the IP router and CL3; note that these PPP sessions do not extend across the MANET to the remote node (unlike the approach in [RFC5578]). At the router these PPP interfaces are unnumbered and OSPF is enabled on them; as point to point interfaces their subnet mask is not checked when setting up OSPF adjacencies, an important benefit for a large MANET that cannot efficiently be in a single IP subnet.

The CL3 rules for IP unicast packet forwarding are:

1. Outgoing (IP router to MANET): Every packet received on a PPP session uses the remote MANET node associated with that session as its IP next-hop. CL3 encapsulates the IP packet into one or more MANET packets, as appropriate, and forwards these to the attached MANET terminal for transmission. The MANET destination is the remote node associated with the PPP session; the MANET source is the MANET address of the attached MANET terminal.

2. Incoming (MANET to node): The MANET source address of the incoming packet is checked by CL3. If it matches a remote node associated with a PPP session, the packet is forwarded to the router on that PPP session. [There is a slight modification to this rule for IP multicast to allow for IGMP and PIM snooping, see Section 7.]

   The packet is dropped if its MANET source address does not match any PPP session. However, if the packet contained an OSPF HELLO, this information is passed to the CL3 network abstraction algorithm before dropping the packet. This gives the network abstraction algorithm an opportunity to decide if that remote node should be used as a routing peer; if so, the remote node is associated with a new PPP session and further

forwarding of packets from that node proceeds as described

above, without generating exceptions for OSPF HELLO packets.

The special handling of OSPF HELLO packets that do not match a PPP session allows CL3 nodes to setup adjacencies to nodes that selected them, even when their own network abstraction algorithm did not select those nodes when it analyzed the MANET topology. It means that network abstraction algorithms running on CL3 nodes do not have to produce the same result and therefore can be modified independently. This meets the fourth CL3 design objective in Section 3. It also allows non-CL3 nodes to setup routing adjacencies with CL3 nodes, meeting the third design objective in Section 3.

Once the PPP interface is setup and associated with a specific remote node in the MANET, the router's OSPF configuration for that interface will cause it to send out OSPF HELLOs, which will be forwarded as described above to that remote node. If the remote peer is also configured to use OSPF on its MANET interface, it will send OSPF packets and CL3, based on their MANET source address, will forward them on the PPP interface to the router. Thus the router will have discovered a potential routing peer on the MANET using the standard OSPF HELLO mechanism, although the remote peer was selected for it by the CL3 network abstraction algorithm.

OSPF processing of HELLO packets will now determine if an adjacency can be setup between the two nodes; incompatible configurations of Area IDs or Network Masks, for example, could prevent this. CL3 uses SNMP to monitor the state of all the routing adjacencies it has selected. If any of these is not in FULL state after some timeout period, for example due to incompatible configuration or due to the RouterDeadInterval expiring, CL3 will free the PPP session. This allows CL3 nodes to work independently of each other, to interoperate with non-CL3 nodes and does not require centrally coordinated OSPF configurations in order to prevent leakage of resources such as PPP sessions.

5 Common Layer 3 Interface components

The description of IP unicast forwarding described above leads to a natural division of the Common Layer 3 interface into four components with well-defined functions:

Network Observer: Obtains Information about the MANET network. At a minimum includes all the nodes in the MANET reachable from the MANET terminal. Each of these nodes is a potential routing peer.

This information is obtained primarily from SNMP queries to the MANET terminal; of course different terminals and MANET routing

protocols will have different MIBs with different contents. Our prototype used OLSR as the MANET protocol and its MIB included a topology table describing the nodes and links in the MANET. Additional information could be obtained from other sources, for example an XML file mapping node IDs to capabilities, such as being an airborne node or being a satellite entry point, that could be useful in evaluating the node's suitability as a routing peer.

Network Abstraction: This is the component that takes as input the information on the MANET obtained by the network observer, processes it with its network abstraction algorithm and outputs the set of nodes in the MANET that should be used as IP routing peers by the attached IP router. As the data from network observer varies over time, network abstraction may decide to add new routing peers, possibly dropping existing ones to make room for them. It may also consider the network history of a remote node (has it been consistently reachable or just intermittently?) and apply approaches based on decision theory.

The network abstraction algorithm is key to the performance of CL3; we tested 15 different variations. Different CL3 nodes can use different algorithms and still setup routing adjacencies between them because of the forwarding rules for incoming stray OSPF HELLO messages.

Router Control: This component takes as input the list of routing peers selected by network abstraction and configures the IP router to use them, as described in Section 4 using PPP sessions. It also notifies the forwarding component of the binding between the PPP session and the MANET address of the routing peer.

To speed our prototype development, we used our router's broadband access features to create PPP interfaces automatically and apply a configuration template to them (which includes enabling OSPF) but clearly these could be avoided and interfaces configured directly by router control using SNMP or CLI scripts.

Router control monitors periodically, using [OSPFv2-MIB] in our case, the state of each OSPF adjacency it has tried to setup and tears down the PPP session if it is not in FULL state after an initial timeout period. This prevents incompatible OSPF configurations such as different Area IDs or Hello timers, from consuming a PPP session.

Of course, the SNMP monitoring mechanism also detects when an

adjacency has been torn down by OSPF, as will happen if the peer is no longer reachable on the MANET and RouterDeadInterval

seconds have passed without a HELLO message from it. This is what we used in our prototype but clearly a more effective approach is for CL3 to setup a BFD [BFD] session to its attached router and proxy for the routing peer. CL3 can deduce if the routing peer is up or down from the MANET information obtained by network observer, which comes directly from the MANET terminal and its MANET routing protocol.

Forwarding: This component implements the IP forwarding rules described in Section 4. It passes stray HELLO information to network abstraction and is told of routing peers to add or drop by router control.

6 CL3 Tests and Results

A CL3 prototype based on the ideas presented above was developed and tested on an emulated radio environment using a variety of network abstraction algorithms.

Our testbed used EMANE, a modular open source radio emulation package for the PHY and MAC layers that can be extended with models of different radio waveforms, including commonly used commercial waveforms and specialized military waveforms; see [EMANE] for details. We modified the open source EMANE software to accept mobility input from a script, allowing very long duration tests. The mobility input describes the strength and loss properties of the radio links at different points in time and must be precomputed based on the radio propagation model for the MANET, the terrain obstructions and node movements. With this information, EMANE can deliver incoming packets to one or more receivers and generate the correct levels of packet loss and delay. The EMANE model we used is a multi-hop radio network.

The MANET terminals and CL3 run on a dedicated Linux/Xen machine where each virtual machine (VM) corresponds to a MANET terminal plus the CL3 interface. Each VM has two VLANs, one to its attached router and another to its corresponding radio module on the EMANE machine. PPPoE runs on the router VLAN. Great care was taken to prevent two VMs on the same host from communicating directly via kernel routing

instead of going through the EMANE radio network. We did this by blackholing outgoing packets to the other MANET nodes but listening on the EMANE VLAN with a raw socket executable that sent the packets to the EMANE interface directly.

One router was configured into multiple VRFs, one for each node. We do not use any mechanism, including BGP, to share routes between these VRFs. By running OSPF on the PPP interfaces setup by CL3, VRFs can setup routing adjacencies with each other but all data packets

between them pass through their corresponding VM and its emulated radio module on the EMANE machine.

Packets sent over the EMANE radio emulation are encapsulated as IP in IP: the outer header has IP addresses known only to OLSR and the EMANE environment while the inner header has IP addresses known to the attached IP routers but not to the EMANE machine. This gave us an easy way to apply IP policy tools to the packet flows on the emulated radio network.

6.1 Mobility Patterns

We used two mobility patterns based on a mix of ground and airborne nodes. The airborne nodes move at higher speeds, which affects their radio propagation properties, and are within line of sight of each other. The speeds chosen are compatible with a mix of ground vehicles and helicopters.

The tethered mobility pattern has four groups of nine ground nodes, where each group patrols a different local area. A separate airborne node orbits the local area patrolled by each group so it can be viewed as being tethered to its group of ground nodes. The airborne nodes are always within line of sight of each other while at times ground nodes in one group may not be in line of sight of some ground nodes in other groups; this affects the radio links between them.

The orbiting mobility pattern has four groups of eight ground nodes each, with each group patrolling a different local area. However, there are also eight airborne nodes orbiting around the entire region, so at different times an airborne node will be above a different group of ground nodes, unlike the tethered mobility

pattern.

These two mobility patterns are more representative of a coordinated deployment of mobile nodes than the random waypoint model often used for MANET simulations.

6.2 Test Results

A baseline system was used to compare against CL3; this system did not use the CL3 network abstraction algorithms but it did monitor the state of its OSPF adjacencies and replaced the non-FULL ones with others. It also accepted stray OSPF HELLO packets just as CL3. The baseline system and CL3 were tested on both mobility patterns.

For measurements, software emulated users attached to each of the IP routers in the CL3 nodes. Every 60 seconds each user selected a

random set of 10 users and sent 10 pings to each. Our software collected the packet delivery ratio and at the end of the test calculated the minimum and average packet delivery ratio as well as its variance. It also identified "failed nodes" for each user, defined as the nodes that did not receive any pings at all from it. A node can be a failed node for one specific user but still receive pings from different users and this was indeed observed.

The test results are in the table below:

	Baseline Tethered	CL3 Tethered	Baseline Orbiting	CL3 Orbiting
#failed nodes:	36	0	33	0
min packet delivery:	0	0.593	0	0.620
avg packet delivery:	0.1	0.781	0.175	0.794

As can be seen, without CL3 the baseline system provided very poor connectivity despite its constant monitoring of OSPF adjacencies and replacement of failed ones. The most important metric in this test is the number of failed nodes: 33 or more out of a total of 40 nodes.

CL3 had no failed nodes in either test, its minimum packet delivery ratio was near 60% and its average packet delivery ratio close to 80%. The same network abstraction algorithm, with the same algorithm

parameters, was used for both mobility patterns. The variance in packet delivery between nodes was 0.20, and so the results are sampled from a near-uniform distribution.

The key advantage of CL3 over the baseline system is that its network abstraction algorithm can take a global look at all potential routing peers and select them so that they are well distributed across the MANET, rather than modifying the initial set of routing peers haphazardly, as the adjacencies fail.

7 IP Multicast Forwarding in CL3

IP multicast forwarding in CL3 adds a slight modification to the incoming forwarding rule in Section 4. The problem is that at the IP layer the MANET appears as a mesh of point to point links so IP multicast packets to a group of nodes in the MANET will be replicated, at the IP layer, once for each receiver. If the MANET technology allows some form of multicast, or is a broadcast network, this will be highly inefficient.

Assume that CL3 learns, either from the MANET terminal or from configuration data, that the MANET has a multicast address M that can

reach a set of nodes S. It can then setup a PPP session to the router and associate it with the MANET address M. The outgoing forwarding rule in Section 4 is not changed. The incoming rule is modified to:

2. Incoming (MANET to node): The MANET source address of the incoming packet is checked by CL3. If the packet does not contain a PIM or IGMP packet and its source address matches a remote node associated with a PPP session, the packet is forwarded to the router on that PPP session.

   If the packet does contain PIM or IGMP and there is a PPP session bound to a MANET multicast address that reaches the MANET source address, the packet is forwarded to the router on that PPP session. Otherwise, it is forwarded to the router on the PPP session bound to the source MANET address of the packet or dropped if there is no such PPP session.

As remote nodes join multicast groups or send PIM-JOIN requests

upstream, the CL3 nodes will be forwarding these to the router on the PPP sessions associated with the MANET multicast address. When that router needs to forward an IP multicast packet, it uses the PPP session bound to the MANET multicast address which is used by CL3 as the MANET destination address of the packet. Then the MANET multicast delivers the packet to the set of receivers by transmitting it only once over the air.

This is the reason for terminating the PPP sessions at the CL3 interface instead of extending them to the routing peer, as done in [RFC5578].

8 Summary

The Common Layer 3 interface described in this paper showed good packet delivery during our prototype tests. Yet it achieved these results using unmodified, widely-available, IP routers and did not depend on any modifications to the MANET routing protocol, only on SNMP access to it.

Use of CL3 will allow the IP routing components of a MANET to be fully independent of the MANET routing protocol. Although MANET protocols continue to evolve, CL3 enabled nodes will be able to adapt to them without requiring change to their IP routing component. They will also be able to change their IP routing protocol, for example moving from OSPFv2 (IPv4) to OSPFv3 (IPv6) without requiring changes to the MANET. This flexibility and modularity are the greatest advantages of the Common Layer 3 interface described here.

9 Security Considerations

CL3 is an interface between a MANET terminal on one side and an IP routing component on the other. These comprise a single mobile node in the MANET. Provided access to the MANET is restricted, for example by encrypting all packets entering it, the only entities that will send data seen by CL3 will be entities with access to the MANET. Therefore the security of a node with CL3 is as strong as the security of the original MANET network.

An entity with access to the MANET can send OSPF HELLO packets to a

CL3 node; these packets will be examined by CL3. Depending on the network abstraction algorithm used by CL3, the packets may cause it to setup an OSPF routing adjacency with the CL3 node. This does not pose a security risk provided only trusted entities have access to the MANET.

10 IANA Considerations

CL3 does not define any code points requiring assigned numbers.

11 References

11.1 Normative References

[CL3] Shaio, J., "Common Layer 3 Interface for Mobile Networks", MITRE Technical Report MTR090194, September 2009.

11.2 Informative References

[OSPFv2] J. Moy, "OSPF Version 2", RFC 2328, April 1998

[OLSR] T. Clausen, P. Jacquet, eds., "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003

[OSPFv2-MIB] D. Joyal, P. Galecki, S. Giacalone, eds., "OSPF Version 2 Management Information Base", RFC 4750, December 2006

[RFC5614] R. Ogier, P. Spagnolo, "Mobile Ad-Hoc Network (MANET) Extension of OSPF Using Connected Dominated Set (CDS) Flooding", RFC 5614, August 2009

[RFC5820] A. Roy, M. Chandra, "Extensions to OSPF to Support Mobile Ad-Hoc Networking", RFC 5620, March 2010

[ANCP-USAGE] S. Ooghe, N. Voigt, M. Platnic, T. Haag, S. Wadhwa, "Framework and Requirements for an Access Control Mechanism in Broadband Multi-Service Networks", RFC 5851, May 2010

[BFD] D. Katz, D. Ward, "General Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010

[RFC5578] B. Berry, et.al., "PPP over Ethernet (PPPoE) Extensions for Credit Flow Control and Link Metrics", RFC 5578, February 2010

[RFC6320] J. Moisand, N. Voigt, T. Taylor, T. Haag, S. Wadhwa, "Protocol for Access Control Mechanism in Broadband Networks"", RFC 6320, October 2011

[EMANE] Cengen Labs, "Extendable Mobile Ad-hoc Network Emulator", http://labs.cengen.com/emane/, June 2012.

[HEND1] Henderson, T., Spagnolo, P., Kim, J., "A Wireless Interface Type for OSPF", IEEE Military Communications Conference MILCOM, vol. 2, pages 1256-1261, IEEE, October 2003.

[HEND2] Henderson, T., Spagnolo, P., Pei, G., "Evaluation of OSPF MANET Extensions", Boeing Technical Report: D950-10897-11, http://hipserver.mct.phantomworks.org/ietf/ospf/reports, July 2005.

[MITRE] MITRE, Common Layer 3 Prototype Source Code, https://sourceforge.net/projects/commonlayer3/files/

Authors' Addresses

Jack Shaio
MITRE Corporation
202 Burlington Road
Bedford, MA 01730
EMail: jshaio@yahoo.com

Caleb Lo
MITRE Corporation
202 Burlington Road
Bedford, MA 01730
EMail: clo03@alumni.caltech.edu

Don Broderick
MITRE Corporation
202 Burlington Road
Bedford, MA 01730
EMail: donb@mitre.org