

TECHNICAL RESEARCH REPORT

A Framework for Scalable Hierarchical Routing in Mobile Ad Hoc Networks

*by Karthikeyan Chandrashekar, Raquel Morera,
Anthony McAuley, John Baras*

TR 2004-49



ISR develops, applies and teaches advanced methodologies of design and analysis to solve complex, hierarchical, heterogeneous and dynamic problems of engineering technology and systems for industry and government.

ISR is a permanent institute of the University of Maryland, within the Glenn L. Martin Institute of Technology/A. James Clark School of Engineering. It is a National Science Foundation Engineering Research Center.

Web site <http://www.isr.umd.edu>

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE A Framework for Scalable Hierarchical Routing in Mobile Ad Hoc Networks				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland, Institute for Systems Research, College Park, MD, 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

A Framework for Scalable Hierarchical Routing in Mobile Ad Hoc Networks

Karthikeyan Chandrashekar*, Raquel Morera^Ψ, Tony McAuley^Ψ, John Baras*

* Institute for Systems Research, University of Maryland, College Park, MD, USA

^Ψ Telcordia Technologies, One Telcordia Drive, Piscataway, NJ, USA

Abstract—The theoretical performance advantages of dividing a network into independent routing domains is well known; however, the actual benefits are hard to quantify and are often not sufficient to outweigh the added complexity. Justification of domains is especially hard in mobile ad hoc networks (MANETS), because reconfiguration overhead increases and use of single interface routers. Nevertheless, we believe that with the right domain configuration and inter-domain routing protocol we can get better performance using hierarchy than flat routing, especially in heterogeneous and dynamic networks. This paper proposes a framework for scalable routing in MANETs based on auto-configured optimized routing domains and an enhanced inter-domain routing scheme. To minimize overall overhead, the inter-domain routing protocol exploits existing messages needed for domain maintenance. The framework allows different routing protocols to run in each domain. OPNET simulations show the benefits of the proposed approach using OLSR for intra-domain routing. Results show significant reduction in protocol overhead, increased route stability and increased route availability in a dynamic heterogeneous network.

Index Terms—scalable routing, network management, inter-domain, MANET routing,

I. INTRODUCTION

Mobile Wireless Ad Hoc Networks (MANETS) are a set of connected wireless nodes configured to form an infrastructure-less network. MANETS are extremely important to those applications where there is a need to rapidly deploy a network without any pre-existing infrastructure, i.e. future battlefield networks, sensor networks and emergency networks. Unlike fixed infrastructure networks, MANET's core network can be very dynamic: routers and servers can leave and join the network, networks can split and merge, etc. These dynamics put networking functions under a lot of strain, especially for large heterogeneous networks, such as that proposed for Future Combat Systems (FCS).

Many of the MANET routing protocols claim to provide good performance and low overhead in large ad hoc networks. There are several factors that affect the performance and scalability of MANET routing protocols like mobility characteristics, traffic flow etc. Most MANET flat routing protocols perform well under specific conditions but do not

scale well in general e.g. [2]. Even approaches to scalable routing that use a multiple routing schemes to improve performance [3-7] do not perform well under all conditions. This is because no single routing protocol can perform well enough under all conditions. It is therefore imperative to be able to choose routing protocols based on the current prevailing network conditions.

In this paper we propose a framework to achieve scalable MANET routing that exploits hierarchical domains that support heterogeneity in intra domain routing protocol. Dividing the network into independent domains [1] helps both with scalability and heterogeneity. In each domain, networking functions operate with more homogenous and limited number of nodes and inter-domain communication is carried out through border nodes. For inter domain routing we propose a scheme that exploits the domain maintenance protocol [11] to reduce overhead and improve the stability (in maintaining shortest paths). We show the benefits of the proposed approach by evaluating protocol overhead, route stability and data delivery in a dynamic heterogeneous network.

The paper is organized as follows, in section II we describe our framework for scalable routing, and in section III we describe the inter-domain routing protocol. Section IV presents the OPNET simulation results and section V concludes the paper.

II. FRAMEWORK FOR SCALABLE ROUTING IN MANETS

To provide scalability to the routing functions and support heterogeneity, we divide the network into independent routing domains. Routing domains are a generalization of the notion of clusters used in the literature. Domain generation and maintenance algorithms ensure that the network is split appropriately and the domains are valid and stable. Intra-domain routing protocols create and maintain local scope routes. Each domain can run the routing protocol that best suits the characteristics of the nodes in that domain. This feature is important, as the nodes should be able to change network functions based on the current state of the network. The inter-domain routing scheme ensures that routes connecting independent domains are formed and maintained. The inter-domain routing proposed is independent of the characteristics of the nodes (fast moving, stationary etc) and the intra-domain routing scheme selected for each domain.

Figure 1 below shows a simple illustration of the above idea where a single connected network is split into two domains.

* Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technology Alliance (CTA) Program, Cooperative Agreement DAAD19-2-01-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

Nodes and domains are configured once the domains are formed. Each domain is configured with a different IP address mask. The domain with IP addresses 192.0.0.X runs OLSR [8] and the domain configured with IP addresses 192.0.1.X runs TORA [9]. In the figure, there are three border routers in each domain. Every node can belong to only one domain. We emphasize here that the border router is a single interface node that can, however, communicate with multiple domains at a time, i.e. the border nodes of each domain are within range of border nodes from other domains and are able to communicate with them as they share the same physical medium even though they have been configured as separate domains.




Figure 1: Multi-domain framework

A. Dividing the network into domains

Several approaches have been proposed for domain/cluster formation in ad hoc networks. Several metrics are considered in forming the domains e.g. hop-count restricted domains, stable domains, balanced domains etc. Both centralized and distributed algorithms have been proposed for domain formation. Once nodes are configured within a routing domain, they will know the routing protocol to run (best suited to the characteristics of nodes and links in the domain) and their IP address. We will not discuss how domains are generated as it is out of the scope of the paper.

B. The Beacon Protocol

Since the network is dynamic we require a domain maintenance protocol to maintain domains as nodes move and the topology changes. The beacon protocol [11] achieves this function. In every domain there is a beacon node responsible for periodically broadcasting the beacon message. The beacon message contains the minimum information so nodes can detect one of the following: a) connected to their current domain; b) split from the network; c) split from their original domain, but can join a new domain; d) connected to their original domain, but can join a neighboring domain that suits better the nodes characteristics. The domain Identifier ID (DID) field is mandatory in every beacon message. Upon

reception of the beacon message, nodes check whether it corresponds to a beacon message from their current domain. If so, nodes forward the beacon message to their neighbors. Contrary, if a node receives a beacon message from a domain different to its current domain, it first evaluates the information contained in the beacon message to decide whether it must join the new domain or remain in its current domain. The priority field in the beacon message contains the information needed by the nodes to decide whether to join a new domain or stay in their current domain. Any combination of the following metrics: beacon age, node degree, number of nodes in current domain, lowest ID can be encoded in the priority field for nodes to determine if a domain change is required. Nodes capable of receiving beacons from multiple domains become border nodes for their domains. The border routers also curtail the spread of the beacon messages thereby defining the boundary of the domains. Each node in the network belongs to a particular domain and the border routers are capable of communicating across domains even though they have only one IP interface.

C. Intra-domain Routing

Domains delimit local routing scope. Domains forming different subnets automatically restrict routing within the subnet. Border routers drop packets from other domains, restricting any broadcast message within the domain boundaries. Nodes within the same domain must run the same routing protocol. Any of the existing MANET routing protocols can be selected for intra-domain routing. The choice of the protocol clearly depends on the characteristics of the domain. If the domain is stable i.e. if the nodes within the domain do not experience many link changes then this domain can run a proactive protocol with a small update frequency or a reactive protocol. On the other hand a proactive protocol with frequent updates is more suitable for volatile domains. The type of traffic and size of the domain can also influence the choice of the routing protocol. The key feature of the domain based routing framework is that we are able to run the routing protocol best suited for the characteristics of the domain.

D. Inter-domain Routing

In the Internet, the Border Gateway Protocol [10] performs the inter-domain routing functions. A BGP like protocol in our context would require the external domain border routers to exchange route tables and then the border routers within a single domain will have to exchange this information. All this places considerable dependency on the border routers and may not be the best approach for dynamic networks where the border routers are continuously changing.

In MANETs, the Zone Routing Protocol (ZRP) [4, 5] and hierarchical approaches Cluster Based Routing Protocol (CBRP) [3] and Landmark routing (LANMAR) [6][7] have been proposed as frameworks for scalable routing. ZRP uses a proactive protocol within the local zone and the inter-domain routing is a combination of a reactive routing protocol and a border-cast protocol; this approach is heavily dependent on the border routers. LANMAR uses Fisheye routing within the

local scope defined based on node mobility characteristics. Cluster heads within each group are called landmarks, which become the Landmarks in the inter-domain routing. Landmark-based hierarchical addressing allowing packets to be routed based on the landmark (group) as well as the host ID.

Achieving a good performance inter-domain routing protocol is challenging, especially in MANETs where even the border nodes may consist of single interface nodes with limited bandwidth and energy. The inter-domain routing protocol has to take into account the following:

- Stability of the inter-domain routes, minimize oscillation amongst different routes
- Minimize the non-optimality of the routes
- Minimize overhead in the network
- Minimize sensitivity to border router mobility (or dependency on the mobility of border routers)
- Support heterogeneity of domains, i.e. be independent of the routing protocol running in each domain

With all this considerations in mind, we propose the following inter-domain routing protocol.

III. INTER-DOMAIN ROUTING

We previously described that the routes within a domain are discovered and maintained by the local routing protocol. The domains themselves are maintained by the beacon protocol. In this section we propose an inter-domain routing scheme that uses the domain maintenance protocol to discover routes to external domains. We describe first a simple inter-domain routing mechanism based only on the beacon protocol and then an enhanced scheme that includes border router information.

A. Beacon Based Inter-Domain Routing

In each domain, the beacon node broadcasts the beacon message to all nodes in the domain. The beacon message contains the domain ID so nodes can identify whether the beacon message corresponds to the domain they belong to. If rather than stopping the propagation of the beacon message at the border routers, we allowed beacon messages to cross-domain boundaries, the entire network would know how many domains are in the network and what is their domain ID. Not only this, but nodes would also know in what direction domains are located. The downside of this approach is obviously a larger overhead in the network.

As the goal of inter-domain routing is to ensure that every node in the network learns the existence of all subnets (prefixes) and therefore all destinations in the network, the inter-domain routing protocol can use the beacon messages also as routing messages. This is possible domain ID a node belongs to is included as part of the node's address. Packets outside the domain are routed based on this domain ID. We assume then, that this domain ID can be part of the IP address prefix so all nodes in the same domain share a common IP address prefix. Therefore, we extended the beacon message to

include the subnet information of the domain. We also add a flag to the beacon message to indicate whether the message is intended for domain maintenance purposes or inter-domain routing purposes. This flag is set to DOMAIN when the beacon message propagates within the domain boundaries and is then used to maintain domains. Border nodes, however, change the flag to ROUTING and forward the message outside the domain boundaries rather than dropping it (which is the original operation mode of the beacon protocol). Thus the beacon message now propagates throughout the network instead of being restricted within the domain. Nodes receiving beacon messages with the function flag set to ROUTING no longer consider these beacons for domain changes. Instead, nodes store the subnet information (or domain ID) contained in the beacon message in their forwarding tables. The subnet mask (or domain ID) is stored as the destination address and the address of the node from which the beacon is received as the next-hop to that destination. This way, every node in the network knows the next-hop node to all destination subnets in the network. As the beacon message is a broadcast message, a node may receive multiple copies of the same beacon message. The sequence number in the beacon message is used to discard subsequent copies of the same beacon message. This route to the destination domain is always fastest (shortest path and less congested). However, routes to farther domains can oscillate. Clearly the accuracy of the routes depends on the update frequency of the beacon message.




Figure 2: Beacon based inter-domain route discovery

Figure 2 describes the inter-domain routing protocol. Node 192.0.0.1 receives the beacon from domain 1 (subnet 192.0.1.X) from node 192.0.0.2. Thus node 191.0.0.1 stores 192.0.0.2 as the next-hop for the subnet destination 192.0.1.X.

1) Data forwarding

Routes to destinations within the domain are discovered by the intra-domain routing protocol. In the case that the intra-domain routing is reactive, the source nodes know if the unknown destination belongs to its domain as the beacon advertises the subnet information to all nodes; in proactive

routing all nodes within the domain are already known. If the destination does not belong to the domain, nodes send packets to the next-hop towards the destination subnet. This information is stored in the node's forwarding table and is obtained from the beacon protocol. Hop by hop, packets finally reach the destination domain. Once packets reach the destination domain, packets are routed to the destination node using the intra-domain routing protocol. It is clear that the accuracy of the inter-domain routes depends on the beacon frequency. It is possible that the next-hop information is outdated between successive beacon updates due to link failures, mobility etc. (need to identify what outdated means) Even in this case the route will still be valid albeit not the shortest path. The intra-domain routing will find the path to a node registered as a next-hop in the forwarding table but which is actually more than 1-hop away.

This scheme is simple but the accuracy of the routes depends on the beacon update frequency. Link changes at any of the intermediate node pairs between a source and a destination subnet results in a change in the path between the source and the destination. To avoid this we can either refresh the routes frequently i.e. increase beacon frequency or reduce the number of nodes in the path between the source and the destination. Clearly, the first solution increases the overhead as the beacon message is propagated throughout the network.

B. Border router and beacon based inter-domain routing

We note that the border routers are like gateway nodes, able to communicate with multiple domains/subnets. Border routers can be the representative nodes describing the path from the source to the destination i.e., the list of next-hop nodes is now a set of border routers that need to be traversed to reach the destination from the source. In order to implement this enhancement we introduce a `LAST_BR` field in the beacon message. Initially, when a beacon is generated this field is invalid. When a border router retransmits the beacon it stores its address in the `LAST_BR` field. A node that receives the beacon checks if a valid `LAST_BR` field exists, if so, the node stores the `LAST_BR` address as the next-hop to the destination domain, else the node stores the address of the node from which received the beacon as the next-hop.

1) Data forwarding

As before destinations within the domain are handled by the intra-domain routing protocol. For an out of domain destination the next-hop address is that of a border router. The intra-domain routing is now responsible for finding the shortest path to the border router. This process continues until the destination domain is reached.

Clearly with this enhancement the number of link changes is limited to the set of neighboring border routers and hence we expect that near-shortest-path routes will be maintained between the source and the destination. In this approach routes to other domain are more stable, but the protocol is sensitive to border router mobility.

IV. SIMULATION ENVIRONMENT AND RESULTS

We set up a simple domain based network to emphasize the benefits of domain based routing in terms of the overhead and convergence properties of the routing. We then evaluate our inter-domain routing scheme in terms of data delivery and overhead.

A. Simulation Environment

OPNET network simulator is used to evaluate our domain based routing framework. The purpose of this evaluation is to study the benefits of the domain framework. We compare the performance of the network when running flat routing (i.e. without any domains) and with domains. The comparison metrics are routing overhead, convergence time and data delivery. To this end, we design a simplistic scenario with 26 nodes as shown in Figure 3; the network is split into four domains. Three of these are fixed and domain (4) consists of moving nodes. The nodes in each domain are automatically configured to belong to different subnets. The beacon protocol is used to maintain domains and the beacon nodes or domain heads are pre-configured, in this scenario nodes 1,8,13 & 15 are the beacon nodes for the 4 domains. The beacon frequency is 5 seconds. Nodes in the moving domain (4) move according to the Billiard mobility model which is the random direction model with reflections at the boundary. The node speed varied from 0 – 10 m/s. Sources in the network generate UDP traffic 10 kbps flows with the packet size being 300 bytes. There are 3 sources in the network and they are chosen such that there is a source from each of the 3 static domains. The 3 corresponding destinations are chosen from the moving domain (4). For simplicity we run OLSR in all domains i.e. in the non-domain case OLSR is used as the flat routing protocol across the network and in the domain case OLSR is run within the local scope defined by the domains. The inter-domain routing protocol implemented is the beacon based inter-domain routing protocol. 802.11b is used as the MAC layer and simulation is run a duration of 30 minutes.




Figure 3: Simulation scenario

B. Results

Figure 4 shows the comparison of the total routing traffic (OLSR traffic) received by all nodes in the network for the single domain case (flat routing) and for the case where the network is split into different domains. For each case, there are two graphs, one when all nodes are static, and another one when some nodes in the network are mobile.

The routing overhead of OLSR is due to the exchange of the HELLO, TC (Topology Control) and the HNA (Host-Network Address) messages. As expected, we observe that the overhead due to the exchange of routing messages is lower in the domain based network than the non-domain network. In the domain based network routing messages are contained within the domain whereas in the non-domain network the routing messages are propagated throughout the network. As OLSR is a proactive protocol, route updates are initiated periodically and the routing overhead is independent on link changes. This is also shown in the figure.




Figure 4: Routing overhead

However, this is not a fair comparison between the two approaches, since when the network is split into domains, inter-domain routing adds to the OLSR overhead. This overhead must also include the beacon messages.

Figure 5 shows the total traffic load at the MAC layer, this is an indication of the total control protocol overhead as there is no traffic in this scenario. In the non-domain case the load represents the total routing traffic along with headers etc at the MAC layer. In the domain case MAC load represents the routing traffic plus the beacon protocol messages along with headers. Clearly, we see that the domain based network has lesser overhead than that of the flat, non-domain network, despite the overhead of the beacon messages. This shows that even though we need a domain maintenance protocol it is beneficial to split the network into domains to limit the routing overhead. From the graph we observe that mobility

has no impact in these conclusions. However, we observe that the curves have some fluctuations induced by MAC layer collisions (lost packets). Some topologies are more prone to MAC layer collisions than others.




Figure 5: Total overhead

Convergence time must also be considered when evaluating scalability and stability of routing protocols. Figure 6 shows the convergence properties of the OLSR routing in the non-domain and the domain case. Convergence activity is an event that results in the addition, update or deletion of an entry in the routing table. The routing is said to have converged if there is no convergence activity for a fixed duration T . We would like to point out that convergence activity does not mean that there are unknown routes still being discovered, it just implies that the routes are being updated as better paths are being discovered. In the graph a transition from 0-1 indicates the start of a new convergence period and the transition from 1-0 the end of the period, the corresponding abscissa value indicated the duration of the convergence period. We see that the non-domain static scenario has several convergence windows of significant duration (order of minutes) whereas the domain static network has convergence windows that converge quickly (order of seconds). This indicates that localizing the routing information can improve the convergence properties. We see that mobility exaggerates the situation with the domain based network also experiencing significant convergence activity. Also both networks have continuous convergence activity after some time which is indicated by the lack of lines on the graph (at time 12m for the non-domain case and at time 16m for the domain based network). Mobility can significantly affect the convergence properties of the routing and again we see that the domain based network is better than the non-domain network as the convergence activity is now restricted to the local routing scope.




Figure 6: Convergence activity




Figure 7: Packets dropped due to MAC collisions

A better understanding of the behavior of the convergence activity can be obtained by looking at the number of packets dropped due to collisions at the MAC layer. Figure 7 shows that a significant number of packets are dropped due to collisions in the non-domain network when compared to the domain based network. This is due to the extra traffic generated in the network, result of the excess propagation of the routing messages in the non-domain network. Furthermore we see that mobility also affects the number of collisions as the 1-hop neighborhood is now continuously changing. Localizing the routing within domains also reduces the impact of the MAC layer performance.

Figure 8 below that the data sent and received in the various scenarios. The domain based static network receives all the traffic sent by the sources. The non-domain static network experiences slight losses due to MAC collisions. In the mobile scenarios (remember the destinations are in the mobile domain) there is a significant loss in data received; this is due to a combination of lack of routes to destination as well as MAC layer collisions. We see that the domain based mobile network performs significantly better than the non-domain mobile network. The partitioning of networks into domains also improves data delivery.




Figure 8: Data traffic sent and received

V. CONCLUSIONS

The domain based framework for routing in ad hoc networks can provide scalable routing. Moreover, the network supports diversity in choosing the intra-domain routing protocols. We see that the domain framework enhances routing by reducing overhead and enhancing data delivery. The inter-domain routing protocol is independent of the intra-domain routing and is based on the beacon protocol. We see also suggest an enhancement to the inter-domain routing protocol that incorporates the border router set to improve path stability. Results show that the total overhead in a domain based network including the cost for maintaining domains is still less than that of non-domain or flat networks. Even for a simple network split into 4 domains we can reduce the routing overhead by 25%. Work is currently in progress to set up experiments with multiple routing protocols running in different domains. We are also studying the stability of critical nodes like domain heads and the border routers. More results regarding this appear in the final version of the paper.

REFERENCES

- [1] R. Morera & A. McAuley, "Flexible Autoconfigured Domains for more Scalable, Efficient and Robust Battlefield Networks", In *proceedings of IEEE MILCOM*, OCT 2002
- [2] S.R. Das, C.E. Perkins and E. M. Royer, "Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks", In *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, Mar. 2000, pp. 3-12.
- [3] M. Jiang, J. Y. Li, and Y. C. Tay. Cluster based routing protocol (CBRP) functional specification. *IETF Internet draft*, Aug. 1999. <http://www.ietf.org/ietf/draft-ietf-manetcbrp-spec-01.txt>.
- [4] Z.J. Haas, "A New Routing Protocol for the Reconfigurable Wireless Networks," In *Proceedings of IEEE ICUPC'97*, San Diego, CA, Oct.1997, pp. 562-566.
- [5] Z.J. Haas and M. R. Pearlman "Determining the Optimal Configuration for the Zone Routing Protocol", In *IEEE Journal on Selected Areas in Communications*, Aug. 1999, pp. 1395-1414.
- [6] G. Pei, M. Gerla and X. Hong, "LANMAR: Landmark Routing for Large Scale Wireless Ad Hoc Networks with Group Mobility," in *Proceedings of IEEE/ACM MobiHOC 2000*, Boston, MA, Aug. 2000, pp. 11-18.
- [7] M. Gerla, X. Hong, G. Pei, "Landmark Routing for Large Ad Hoc Wireless Networks", in *Proceedings of IEEE GLOBECOM 2000*, San Francisco, CA, Nov. 2000.
- [8] T. Clausen and P. Jacquet "Optimized Link State Routing Protocol (OLSR)." RFC 3626, IETF Network Working Group, October 2003.
- [9] V.D. Park and M.S. Corson, "A Highly Adaptive Distributed RoutingAlgorithm for Mobile Wireless Networks," In *Proceedings of IEEE INFOCOM'97*, Kobe, Japan, Apr. 1997, pp. 1405-1413.
- [10] Rekhter, Y., and T., Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, T.J. Watson Research Center, IBM Corp., Cisco Systems, March 1995.
- [11] R. Morera, et al, "Robust Router Reconfiguration in Large Dynamic Networks". CTA 2003

ⁱⁱ The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied of the Army Research Laboratory or the U.S. Government