

Number 708



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

IDRM: Inter-Domain Routing Protocol for Mobile Ad Hoc Networks

Chi-Kin Chau, Jon Crowcroft, Kang-Won Lee,
Starsky H.Y. Wong

January 2008

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2008 Chi-Kin Chau, Jon Crowcroft, Kang-Won Lee,
Starsky H.Y. Wong

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

IDRM: Inter-Domain Routing Protocol for Mobile Ad Hoc Networks

Chi-Kin Chau, Jon Crowcroft, Kang-Won Lee, Starsky H. Y. Wong *

Abstract

Inter-domain routing is an important component to allow interoperation among heterogeneous network domains operated by different organizations. Although inter-domain routing has been extensively studied in the Internet, it remains relatively unexplored in the Mobile Ad Hoc Networks (MANETs) space. In MANETs, the inter-domain routing problem is challenged by: (1) dynamic network topology, and (2) diverse intra-domain ad hoc routing protocols. In this paper, we propose a networking protocol called IDRM (Inter-Domain Routing Protocol for MANETs) to enable interoperation among MANETs. IDRM can handle the dynamic nature of MANETs and support policy-based routing similarly to BGP. We first discuss the design challenges for inter-domain routing in MANETs, and then present the design of IDRM with illustrative examples. Finally, we present a simulation-based study to understand the operational effectiveness of inter-domain routing and show that the overhead of IDRM is moderate.

1 Introduction

Mobile ad hoc networks (MANETs) can enable effective communications in dynamic operation environments including a coalition military operation, emergency operation for disaster recovery, and on-the-fly team formation for a common mission, such as search and rescue. In these situations, multiple groups and organizations may need to come together, communicate, and collaborate to achieve a common goal. For example, in a disaster recovery scenario, the local police force may need to coordinate with fire fighters, military forces, and medical crews by sharing information and communicating with each other regardless of the particular networking technologies that each group uses. Such application scenarios call for development of a technology to enable end-to-end communications over heterogeneous MANETs governed by distinct administrative domains.

Facilitating interoperation among multiple MANETs pre-sents a significant challenge at multiple levels, from physical to application layers. As a first step towards a full MANET interoperation solution, in this paper, we investigate the important problem of inter-domain routing in MANETs, and propose a novel protocol to support that. In the

*Authors are ordered alphabetically. Chi-Kin Chau and Jon Crowcroft are with Computer Laboratory, University of Cambridge, UK. Kang-Won Lee and Starsky H. Y. Wong are with IBM T. J. Watson Research Center, Hawthorne, USA. Email: {chi-kin.chau,jon.crowcroft}@cl.cam.ac.uk, {kangwon,hwong}@us.ibm.com

Internet, the Border Gateway Protocol (BGP) [9] provides a standard mechanism for inter-domain routing among heterogeneous domains, called autonomous systems (AS). The principle of BGP is to enable *opaque* interoperation, where each domain has the administrative control over its intra-domain routing protocol and inter-domain routing policy, which is not known (or opaque) to the other domains. Despite several reported peculiarities [7], BGP has been rather effective in the wired world. Unlike in the static Internet, inter-domain routing is a relatively new problem in MANETs with significant challenges. We identify two major challenges. First, in MANETs, the network connectivity changes dynamically, thus an inter-domain routing protocol must be able to cope with such changes as network partitions/merges and connectivity changes. Second, MANET environment has spawned out a new breed of routing protocols [1] that are specialized for dynamic networks, and they require special handling to participate in inter-domain routing.

In this paper, we propose a novel networking protocol, called IDRM (**I**nter-**D**omain **R**outing Protocol for **M**ANETs) to enable inter-domain routing for MANETs. The salient features of IDRM are as follows. First, IDRM requires no surrender of the administrative control from each domain. Thus each domain can specify inter-domain routing policies in the spirit of the policy-based routing as supported by BGP in the Internet. Second, IDRM has been designed to effectively address the two main challenges identified above. Particularly, it employs a proactive routing for inter-domain gateway communication to readily detect any topology changes, and adapt to those changes. Third, it supports each domain to participate in the inter-domain routing operation without any changes to their native intra-domain routing protocols. We evaluate the effectiveness of the inter-domain routing in MANETs in various operation scenarios. We also present asymptotic analysis of the control overhead incurred by IDRM, and show that the overhead of IDRM is moderate.

It is important to note that our goal is *not* to extend the BGP framework for MANETs. BGP is a highly specialized protocol designed to cope with the scale and operational challenges of the Internet. Compared to the Internet, MANETs are considerably small yet highly dynamic environments. Instead, our approach is to borrow the core design principles of BGP that are valid in our context and take a clean slate approach to enable inter-domain routing in MANETs. Recognizing this area has not received much attention from the research community, our more ambitious goal is to start a discussion in this area by proposing a first cut solution that is both practical and amenable for accommodating future extensions.

Outline: Sections 2 and 3 discuss the related work and background. Section 4 presents the challenges of inter-domain routing in MANETs. Section 5 presents the design of IDRM. Section 6 evaluates the effectiveness and overhead of IDRM. Section 7 discusses several remaining design issues.

2 Related Work

In the literature, there are several proposals to enable interoperations among multiple wireless domains [5] [10]. Most of them focus on high level architectures and provide a sketch of required components (e.g., translation of different naming spaces, and different protocols). Plutarch is an architecture that translates address spaces and transport pro-

protocols among domains to support interoperation of heterogeneous networks [5]. TurfNet is another proposal for inter-domain networking without requiring global network addressing or a common network protocol [10]. While these related works have considered various issues regarding interoperation of multiple networks, none of them provided a specific solution for inter-domain routing between MANETs. Our work builds on these high level architecture and proposes a practical framework for inter-domain routing in MANETs that can support opaque interoperation.

In the wireless context, there have been proposals to take advantage of heterogeneous routing protocols to adapt to network dynamics and traffic characteristics. Hybrid routing protocols combine different routing protocols and adaptively use them to improve the performance. For example, SHARP [8] uses both proactive and reactive routing protocols to balance between the two and adapt the routing behaviour according to traffic patterns. The basic idea of SHARP is to create proactive routing zones around nodes where there are lots of data traffic, and use reactive routing in other areas. Although hybrid routing is similar to inter-domain routing in that it combines different routing protocols, its main goal is to improve the routing performance in a single domain via adaptation. On the contrary, this paper studies the problem of routing across several domains where each domain may employ any routing protocol.

Cluster-based networking in MANETs [3] is similar to inter-domain routing because it also concerns the routing between clusters of nodes. The idea of cluster-based networking is to form self-organizing clusters and a routing backbone among clusterheads. In this way, cluster-based networks can take the advantage of hierarchical routing and achieve a scalable routing solution in a single domain. Although cluster-based routing has a structural similarity to inter-domain routing, there are fundamental differences. The nature of inter-domain routing is on multiple heterogeneous domains with autonomous control; thus the hierarchy of the network is given. Also as we will see later in the paper, gateway nodes are fixed. On the other hand, a cluster-based routing is applicable in a single domain with a large number of participants, and the clusterheads must be elected among the participating nodes. Basically, cluster-based routing is a hierarchical routing concept applied to the MANET domain, whereas our work concerns enabling interoperation of multiple MANET domains.

Finally, although there have been some discussions on inter-domain routing for MANETs, they ignore the practical issues in inter-domain routing. For instance, in [11], the authors presented three protocols for inter-domain routing, and evaluated them using performance-oriented metrics such as packet delivery ratio and end-to-end delay. Their proposals were based on a strong assumption that all nodes can act as gateway nodes if needed. Also they did not consider issues such as policy-based interaction between domains. Our work differs from the related work in that we focus more on addressing practical issues in enabling opaque interoperation among heterogeneous MANET domains.

3 Background

Inter-domain routing should support the interoperation of networks governed by different administrative domains that usually employ different routing protocol designs, metrics and policies. This section reviews the main characteristics of the current inter-domain routing framework in the Internet.

At minimum, an inter-domain routing protocol should be able to handle the following two main problems:

(1) Heterogeneous Intra-domain Protocols:

The internal routing protocols within each domain can be any routing protocol (e.g., OSPF, IS-IS, RIP). While most of these routing protocols are based on a shortest path algorithm, in general, the routing metrics can be expressed as arbitrary preferences over forwarding paths, based on the performance (e.g., QoS-based routing) or security metrics. One of the functions of an inter-domain routing protocol is to hide the differences between intra-domain routing protocols.

(2) Heterogeneous Inter-domain Routing Policies:

Support for domain-level routing policies is an important component of inter-domain routing. For example, there are different business relations among domains, such as provider-customer relation where providers are liable to provide transit connections for their customers, and peer-peer relation where connections are offered only to facilitate mutual traffic flow. Also, there are security reasons that a domain is unwilling to traverse insecure domains based on local security settings. Some connections between specific pairs of gateways may be configured as backup, which will not be invoked except when no other path is available.

In BGP, the routes to an internal destination within the same domain are determined by an intra-domain routing protocol, whereas the routes to an external destination are determined by the inter-domain routing policies among domains. BGP relies on a path vector protocol for exchanging inter-domain level reachability information. One of the advantages of the path vector protocol is that it is easy to specify domain administrator's preferences in the route selection thereby enabling a policy-based routing.

In the Internet, each domain is assigned a unique identifier called Autonomous System (AS) Number. Each domain can filter and selectively announce the routes to a specific destination through path vector protocol that maintains a list of AS numbers of the traversed domains. Upon receiving the route announcements from neighbors, each domain decides the routes based on its inter-domain routing policy specified by network administrators. The external process of BGP operating across domains for exchange of route information is called the exterior Border Gateway Protocol, whereas the internal process within a domain to assist route decision is called the interior Border Gateway Protocol.

The design of BGP relies on several assumptions: (a) Gateways have complete knowledge of reachable destinations within the domain, independent of data transmission (i.e., intra-domain routing is proactive); (b) To reduce the size of routing tables, BGP gateways rely on IP prefix aggregation (i.e., each sub-network can be effectively represented by an IP-prefix); and (c) Loop detection in the inter-domain topology is done by identifying repeated AS numbers (i.e., a single AS number uniquely identifies a single domain). Although these assumptions are reasonable in wired networks, they may not be true in MANETs. We study some design challenges in MANET environments in the next section.

4 Challenges

MANETs are fundamentally different from wired networks, and the operating environments of mobile wireless devices have resulted in a set of new heterogeneous elements

to be considered in inter-domain routing. One major challenge to provide inter-domain routing to MANETs is the dynamic nature of network topology. Unlike the Internet, there are no designate network borders between MANETs because they can move around to connect via different gateways, and multiple networks may even geographically overlap some times.

Also, a domain of MANET can be partitioned into disjoint networks without direct intra-domain connectivity, and the intra-domain connectivity may only be maintained by traversing the nodes in other domains. These properties impair the direct application of a traditional inter-domain routing protocol to MANETs.




Figure 1: The MANET of domain *A* is partitioned due to mobility.




Figure 2: A similar setting in terms of topology change in BGP.

Consider Figure 1 consisting of three MANET domains. One might apply BGP to this scenario as in Figure 2. However, there are several issues that make BGP inapplicable. First, the path vector protocol in the BGP implicitly assumes the availability of the following functions:

(1) **Internal Gateway Detection:** The internal gateways within the same domain can detect the presence of each other so that they know whom to communicate with about the information of external routes.

(2) **Internal Network Knowledge:** The gateways know the reachable destinations and the internal routes to the destinations within the domain.

These functions are normally supported by the proactive intra-domain routing protocols (e.g., distance-vector and link-state routing protocols) through continual maintenance of network state information. However, we cannot always assume the availability of this information in MANETs that uses a reactive or hybrid routing protocol in their domains, as they do not necessarily provide these functions out of their regular operation. Without careful design, a direct application of a path vector protocol over MANETs with add-on processes to support these functions may be undesirable to MANETs with dynamic node mobility and scarce wireless communication bandwidth.

Second, in BGP every destination is identified by an IP address, which follows a certain network hierarchy. To announce the destinations in a domain, gateways will aggregate the IP addresses in the domain by suitable IP prefixes (e.g., 92.168.0.0/16). However, in MANETs, mobility and ad hoc deployment can create arbitrary network partition, unlike the perfect split of IP addresses as in Figure 2. Hence, IP prefixes may not suitably aggregate the IP addresses in partitioned MANETs. Thus we cannot use the prefix-based routing of BGP, and this may create a problem of providing scalable inter-domain routing tables.

Third, BGP relies on a path vector protocol that filters the paths consisting of repeated AS numbers to prevent looping. For example, in Figure 2, after topology change, the inter-domain level path from a source in AS 45 (92.168.1.0/8) to AS 2334 (112.18.0.0/16) is AS 45→AS 310→AS 45→AS 2334. This path will be filtered by the BGP path vector protocol, and hence it will prevent the nodes in AS 45 (92.168.1.0/8) from reaching AS 2334 (112.18.0.0/16).

These challenges led us to design a new inter-domain routing protocol for MANETs as we present in the next section.

5 IDRМ: Inter-Domain Routing Protocol for MANETs

We propose the IDRМ to enable opaque interoperation among multiple domains of MANETs. In this way, each MANET retains administrative control within its own domain while participating in collaboration. To enable inter-domain communications, IDRМ requires special nodes as *gateways*. The role of gateways is more than just handling inter-domain routing; they need to bridge any technical seam that may exist between MANETs at physical, MAC, network, and transport levels. For example, they may need to speak multiple radio technologies, understand different MAC layer interaction, or translate between different protocols. However, the main focus of this paper is limited to the inter-domain routing functions of the gateways.

IDRМ relies on a path vector protocol to support a policy-based routing, where each domain can define arbitrary policies on its preferences on domain-level paths. But unlike BGP, in order to balance the trade-off between overhead and performance, we carefully design IDRМ to proactively support the internal gateway detection and make use of partial internal network knowledge for detecting network partitioning and merging, but does not require complete knowledge of reachable destinations and internal routes. In this way, we can assure that the core inter-domain route information is shared among gateways and major network events such as network partition and merge can be detected timely. We call this scheme a *semi-proactive* path vector routing because at the inter-domain level information update is proactive but at the intra-domain level the routing can be reactive or hybrid. We discuss different design alternatives in Section 7.

5.1 Design of IDRМ

We represent each mobile node by a unique node ID in lower case (e.g., a_1, a_2), and each domain by a unique domain ID in upper case (e.g., A, B). We assume each node belongs

to only one domain. In each domain, there are a subset of nodes functioning as gateways. A pair of nodes are connected by a communication link, if one is within the transmission radius of another. A MANET is a connected directed graph with nodes belonging to the *same* domain.¹

For the simplicity of presentation, we assume the following throughout the paper (with a note that these assumptions are not fundamental to our design).

- The node IDs are unique throughout the entire MANETs. In practice, each MANET will typically belong to an organization, which will have an address space pre-assigned or allocated on-demand from a centralized entity.
- There is a naming service similar to DNS that translates a name to an ID (see [10] for example). When a sender wants to send a packet to a destination, it can obtain the ID of the destination.
- The communications between inter-domain gateways are bidirectional, and the gateways can support multiple radio access technologies to enable the communications among different domains.
- A non-gateway node is always willing to forward packets for intra-domain nodes, but not necessarily for other inter-domain nodes.

Now we explain the key design points of the IDRM. There are several issues that we need to handle: (1) partition and merge of domains, (2) membership announcement, (3) support for policy-based routing, and (4) data plane operations. The first two points are due to node mobility and dynamic topology, and the latter two are general issues with inter-domain routing with autonomy of each domain.

5.1.1 Handling Domain-level Topology Changes

As illustrated in Section 4, one of the key challenges for inter-domain routing in MANET is that the network topology may change dynamically. In particular, a single domain may be partitioned into multiple MANETs due to node mobility. Consider the case in Figure 1 when a domain A has split into two networks. In this case, the gateways in domain A first need to discover the partition. This is done via periodic internal gateway detection. In a domain where the intra-domain routing protocol is proactive, this event will be eventually detected via route updates. For a domain with a reactive intra-domain routing protocol, however, this event may not be detected for a long time. To handle this problem, in IDRM, the gateways maintain soft state by periodically sending beacons to each other. The period of beacon can be adaptively set based on the mobility of the nodes and the rate of topology change.

After detecting a partition, the gateways in the same partition should generate a new MANET ID so that the new partition can be uniquely identified. We want this computation can be performed independently at each gateway so that control traffic is minimized, yet all the gateways in the same partition to generate the same ID. At the

¹We use “MANET” to denote a physical network component, and use “domain” to denote a logical grouping throughout this paper. For example, there may be multiple MANETs of a single domain due to a network partition as illustrated in Figure 2.

same time, we also want the collision of IDs of different networks to be as low as possible. To achieve this goal, we use a pseudo random number generator to generate a new ID using the IDs of all the gateways in the network as input. For this, the gateways in the same network must discover all the intra-domain gateways in the partition. Then they use a simple hash function (e.g., MD5) to generate a random number, then prefix it by the domain ID to get a new MANET ID. We encode the domain ID in the new MANET to support a dynamic policy translation (discussed in 5.1.3). In the paper, for clarity, we use a more explicit ID that looks like $[A:a_1:a_2:\dots:a_n]$, where A is the original domain ID and a_1, \dots, a_n are the IDs of gateways in the partition. By dynamically assigning a new ID, we can prevent the path vector routing algorithm from mistakenly considering the route via partitioned networks as a loop.²

Conversely, when two or more partitioned MANETs come close and re-connected, this condition should be detected by the gateways and a new ID for the merged MANET should be generated. At a high level, this process is similar to the case of network partitioning - the gateways in each partition should keep sending out beacons and when they receive a beacon from a gateway that is not currently in the set of connected intra-domain gateways, then they detect that network merge has occurred. Once the merge is detected, the gateways collect the IDs of all the gateways in the merged MANET and generate a new MANET ID.

5.1.2 Membership Management and Announcement

In addition to MANET ID generation, the gateways also need to collect the IDs of all the nodes in the MANET for advertisement of the membership to other domains. Note that the partition is arbitrary, thus we cannot rely on IP prefix-based routing for partitioned domains. There are two possible approaches to deal with the situation. First, the gateways can coordinate and reassign the node IDs so that each partitioned domains can have unique prefix. However, this will incur significant management overhead (e.g., to generate unique node IDs, to update name-to-ID mapping) and thus will only be useful when the new network topology will remain unchanged for a long time.

Second, a more practical approach to handle topology changes is to let the gateways in partitioned networks advertise the membership information in the form of membership digest, and this digest is used for inter-domain routing. For a reasonable size MANET, we find that a plain membership digest just containing a set of node IDs (e.g., IP addresses) without any compression will suffice as opposed to a more scalable solution (e.g., based on a Bloom filter [2]). In particular, we find that for a network where each domain has less than 1000 members, it is better to use the plain membership list. We discuss more on this topic in Section 6.3. Obviously, the second approach can cope with network dynamics better and is more graceful when partitioned MANETs merge (by merging the membership digests). Hence, the second approach is employed in IDR.M.

Keeping track of the non-gateway membership in a domain poses a similar challenge to network partition detection, but for a larger number of nodes. As in partition detection,

²It is possible to extend this basic protocol to include a leader election process and let the leader of a domain coordinate intra-domain operations (e.g., hierarchical beacons among gateways, or MANET ID generation). But we do not discuss such schemes here for clarity of presentation.

proactive routing protocols naturally provide updates on the membership changes (i.e., addition or departure of nodes). In a reactive routing case, a gateway may have a stale view of its membership, and can only discover the membership change when it has data to transfer. Alternatively, it can periodically initiate a membership query to update its information. However, this will incur significant overhead, and should not be used periodically. Instead, when a node could not be reached, a gateway in a reactive domain may need to contact the gateways in other partitions to discover the node, which will in turn initiate membership query in those networks. If it finds a network with the destination, the data will be forwarded there, otherwise the destination node is considered disconnected.

5.1.3 Policy Support

Inter-domain routing policy is enforced in the same way as in BGP, by using path vector routing at AS-level. By exchanging route announcements in path vector protocol, inter-domain routing policies will be translated as the decisions of filtering and selecting routes at gateways.

For example, if a gateway a_1 in domain A is willing to provide transit service to a neighbouring domain B for a certain destination c_1 , then a_1 appends its MANET ID to the route announcement of the selected path to c_1 and announces it to a connected gateway b_1 in domain B . Upon receiving the route announcement, b_1 will decide if this path is more preferable than the current using path to c_1 based on its routing policy. If a new path is selected, b_1 will record the source of announcement as a_1 and distributes the announcement to other internal gateways in the MANET.

There are a variety of ways to specify routing policy rules. For example, in a next-hop-based policy specification, gateways will select paths only based on the next-hop domain in the route announcement (which may be based on commercial relations like customer, provider, or peer). In a path-based policy specification, a domain will specify a complete ordered preference of all the acyclic out-going domain-level paths. Paths with higher rank are more preferable, and low priority paths (e.g., backup paths) will be given a low rank. In a cost-based policy specification, a domain will assign a numerical cost to every other domain as a subjective evaluation of the performance. Gateways will select the paths with the minimum total cost of all the downstream domains.

In MANETs, a single domain may be partitioned and merged back dynamically, and the routing policies must be translated dynamically as the domain level topology changes. In particular, we require a mechanism to dynamically translate dynamic MANET IDs into appropriate static domain IDs that can be specified in the routing policies. In IDR, we encode the original domain ID in the dynamic MANET so that the policy translation is straightforward.

5.1.4 Data Plane Operations

The packet forwarding process in the data plane will make use of the routing information collected from the above control plane operations.

When a node sends data packets to an external destination (in another domain or in another partitioned network), it forwards the packets to one of the reachable intra-domain gateways. In a reactive domain, the sending node will first initiate a route discovery, and a

gateway node that has a route to the destination will respond. In a proactive domain, the sending node will have a list of intra-domain gateways, and select one of them based on its own preferences. Note that the list of reachable intra-domain gateways will be learned from the regular route updates. In either case, the gateway will first see if it is directly connected to the domain that contains the destination. If it is then it just forwards the packet; otherwise, it will forward the packets to a gateway connected to the destination domain (based on the inter-domain routing information).

For packets coming into the domain, the gateway performs a protocol translation and invokes the intra-domain routing protocol. In a reactive domain, the gateway will initiate a route discovery process if it does not already have the route in the cache. In a proactive domain, the gateway can determine if the destination is reachable from the local routing table. An illustration of data forwarding process is presented in Section 5.3.

If for some reason the destination cannot be reached (e.g., the node may have been disconnected from any domain) IDRM does not provide feedback for unreachable destination as it may generate unnecessary control traffic. Just like the routing in the Internet, the problem should be handled at a higher layer. Although we only discuss proactive and reactive routing protocols in this paper, it is not difficult to see that this framework can support other types of intra-domain routing protocols (e.g., geo-routing and hybrid routing). We do not present these cases due to space constraints.

5.2 Protocol Specification

This section describes the inter-domain routing protocol of IDRM in pseudo codes. We present three algorithms to be executed at each gateway. Algorithm 1 is a subroutine to generate a new route announcement. Algorithm 2 is a continual process of a gateway to handle the interaction between inter-domain gateways. Algorithm 3 is a continual process to manage the intra-domain membership.

For a gateway i in a domain A , let $G^{\text{intra}}(i)$ denote a set consisting of the intra-domain gateways to that i has connectivity, and $G^{\text{inter}}(i)$ denote a set consisting of the inter-domain gateways i is directly connected. Let $M(i)$ denote the set of intra-domain members to that i has connectivity.

Algorithm 1 Route Announcement Generation

```

if (any change in  $G^{\text{intra}}(i)$ ) then
  // generate a new MANET_ID
  MANET_ID  $\leftarrow f(A, G^{\text{intra}}(i))$ 
  // else MANET_ID does not change
end if
if (any change in  $M(i)$ ) then
  // generate a new membership digest
  MD  $\leftarrow b(M(i))$ 
  // else the membership digest does not change
end if
path  $\leftarrow \{\text{MANET\_ID}\}$ 
return a new route announcement [MD, path]

```

Algorithm 1 checks any change in the membership of $G^{\text{intra}}(i)$ and $M(i)$, and generate

a new route announcement if necessary. Here, the function f denotes a one-way hash function (e.g., MD5) to create a MANET_ID based on the original domain ID, and the set of gateways, and the function b generates a membership digest based on the $M(i)$ (see Section 5.1.2 for how to obtain $M(i)$). Function b can be configured to either return a plain membership list or a membership summary (e.g., using a Bloom filter). One practical concern is securing the route updates to prevent attacks on routing table. Securing the route updates and other aspects of path vector protocol in MANETs is a future research topic, which will be briefly discussed in Section 7.

Algorithm 2 Main Routine of the Gateway

```

while (true) do
  if (timer > announcement interval) then
    // generate a new route announcement
    call Algorithm 1
    send the route announcements to  $G^{\text{inter}}(i)$ 
  end if
  if (received a route announcement [MD, path]) then
    if (announcement from  $g^{\text{new}}$  not in  $G^{\text{inter}}(i)$ ) then
      // new connected inter-domain gateway found
       $G^{\text{inter}}(i) \leftarrow G^{\text{inter}}(i) \cup \{g^{\text{new}}\}$ 
    end if
    if ((no route to MD) OR (path  $\prec$  route to MD)) then
      // update the path vector
      insert [MD, path] at the top
      path  $\leftarrow$  append (MANET_ID, path)
      announce [MD, path] to  $G^{\text{inter}}(i) \cup G^{\text{intra}}(i)$ 
    end if
  end if
  increment timer and sleep
end while

```

Algorithm 2 presents the main function of a gateway participating in IDRM. The main routine consists of two parts. First, it periodically polls its domain status, generates a new route announcement, and broadcasts a route announcement to its neighbouring inter-domain gateways. Second, it wakes up when a new route announcement is received from one of its neighbours and process them. In the route announcement, **path** is an ordered list of MANET_IDs, i.e., $[\text{MANET_ID}_1, \dots, \text{MANET_ID}_n]$, which indicates the nodes in MD can be reached by traversing MANET_ID_1 , then $\text{MANET_ID}_2, \dots$, and finally MANET_ID_n . When it processes a route announcement, it first examines if the origin of the announcement is already in its list of neighbours. If it is a new neighbour it updates the list. Then it compares the new path information using its inter-domain routing policy. If the route specified in the **path** is allowed and is more preferable than the current route to MD based on inter-domain routing policy (i.e., **path** \prec route to MD), then it updates its routing table by inserting the new route to the top.³ It then appends its own MANET ID in front of **path** and rebroadcasts the information to its inter-domain

³This is based on the assumption that the preference of a route is determined by the order in the routing table.

and intra-domain neighbours.

Algorithm 3 is a separate thread that takes care of the exchange of beacons among the gateways in the same domain. Periodically, a gateway sends out a beacon to all intra-domain gateways notifying its presence. When it does not receive a beacon from one or more of the gateways in its intra-domain, it updates $G^{\text{intra}}(i)$. Similarly, when it receives a beacon from a gateway g that is not currently in the list of intra-domain gateways, it updates its entry. When these changes are detected, the gateway initiates a route announcement process to update its neighbours.

Algorithm 3 Beaconing among Intra-domain Gateways

```

while (true) do
  if (timer > beacon interval) then
    send beacons to every gateway in  $G^{\text{intra}}(i)$ 
  end if
  for all (gateway  $g$  in  $G^{\text{intra}}(i)$ ) do
    if (no beacons from  $g$  within time limit) then
      // network has partitioned
       $G^{\text{intra}}(i) \leftarrow G^{\text{intra}}(i) \setminus \{g\}$ 
      raise change flag
    end if
  end for
  if (received a beacon from  $g$  not in  $G^{\text{intra}}(i)$ ) then
    // network merge event OR new gateway
     $G^{\text{intra}}(i) \leftarrow G^{\text{intra}}(i) \cup \{g\}$ 
    raise change flag
  end if
  if (change flag is up) then
    // generate a new route announcement
    call Algorithm 1
    send the route announcement to  $G^{\text{inter}}(i)$ 
    reset change flag
  end if
  increment timer and sleep
end while

```




Figure 3: We deploy gateways in the setting of Figure 1.

Received by gateways a_2 and a_3		Received by gateways b_2 and b_3	
Destinations	Inter-domain Route	Destinations	Inter-domain Route
$\{a_0, a_1, a_2, a_3\}$	Internal	$\{a_0, a_1, a_2, a_3\}$	$b_2 \rightarrow [A:a_2:a_3]$
$\{b_1, b_2, b_3\}$	$a_2 \rightarrow [B:b_2:b_3]$	$\{b_1, b_2, b_3\}$	Internal
$\{c_1, c_2\}$	$a_2 \rightarrow [B:b_2:b_3] \rightarrow [C:c_2]$	$\{c_1, c_2\}$	$b_3 \rightarrow [C:c_2]$

Received by gateway c_2	
Destinations	Inter-domain Route
$\{a_0, a_1, a_2, a_3\}$	$c_2 \rightarrow [B:b_2:b_3] \rightarrow [A:a_2:a_3]$
$\{b_1, b_2, b_3\}$	$c_2 \rightarrow [B:b_2:b_3]$
$\{c_1, c_2\}$	Internal

Table 1: The received route announcements by the gateways in Figure 3, before the topology change.

Received by gateway a_2		Received by gateway a_3	
Destinations	Inter-domain Route	Destinations	Inter-domain Route
$\{a_1, a_2\}$	Internal	$\{a_1, a_2\}$	$a_3 \rightarrow [B:b_2:b_3] \rightarrow [A:a_2]$
$\{a_0, a_3\}$	$a_2 \rightarrow [B:b_2:b_3] \rightarrow [A:a_3]$	$\{a_0, a_3\}$	Internal
$\{b_1, b_2, b_3\}$	$a_2 \rightarrow [B:b_2:b_3]$	$\{b_1, b_2, b_3\}$	$a_3 \rightarrow [B:b_2:b_3]$
$\{c_1, c_2\}$	No Route	$\{c_1, c_2\}$	$a_3 \rightarrow [C:c_2]$

Received by gateways b_2, b_3		Received by gateway c_2	
Destinations	Inter-domain Route	Destinations	Inter-domain Route
$\{a_1, a_2\}$	$b_2 \rightarrow [A:a_2]$	$\{a_1, a_2\}$	$c_2 \rightarrow [A:a_3] \rightarrow [B:b_2:b_3] \rightarrow [A:a_2]$
$\{a_0, a_3\}$	$b_3 \rightarrow [A:a_3]$	$\{a_0, a_3\}$	$c_2 \rightarrow [A:a_3]$
$\{b_1, b_2, b_3\}$	Internal	$\{b_1, b_2, b_3\}$	$c_2 \rightarrow [A:a_3] \rightarrow [B:b_2:b_3]$
$\{c_1, c_2\}$	No Route	$\{c_1, c_2\}$	Internal

Table 2: The received route announcements by the gateways in Figure 3, after the topology change.

5.3 Illustration of the Operations

We now illustrate the control plane operations of IDRM. Consider Figure 3; nodes a_2, a_3, b_2, b_3, c_2 are gateway nodes, whereas a_0, a_1, b_1, c_1 are non-gateway nodes. Suppose the inter-domain routing policies of domains A, B, C as follows:

- A is willing to provide transit service for B , but not for C ,
- B, C are willing to provide transit service for others.

In the operation, the internal gateways first detect the connectivity among each other. The node IDs of connected intra-domain gateways establish three MANET IDs as $[A:a_2:a_3], [B:b_2:b_3], [C:c_2]$. By the path vector protocol, the received route announcements by the gateways are given in Table 1.

In Figure 3, after the topology changed, gateways a_2 and a_3 find that they cannot detect each other. This partition event will triggers both a_2 and a_3 to (1) create new MANET IDs as $[A:a_2]$ and $[A:a_3]$, (2) collect the reachable intra-domain nodes, and (3) update the route announcements. The updated route announcements will be propagated to other inter-domain gateways though the path vector protocol. The final route announcements received by the gateways are given in Table 2. Note that B cannot reach

C through $[A:a_3]$ as A is unwilling to provide transit service for domain C . Hence, $[A:a_2]$ also cannot reach C .

We now look at the forwarding process in the data plane in Figure 3 by considering the cases of OLSR and DSR as the intra-domain routing protocols. After the topology change, suppose a_1 wants to send packets to a_0 . Suppose A is using OLSR. Based on the link state information forwarded by MPRs (not shown in the figure), a_1 know that both the destination a_0 and default gateway a_3 are disconnected from the MANET. Hence a_1 sends packets to another gateway, a_2 , that can be reached. With the updated route announcements from b_2 , a_2 knows that a_0 is located in another MANET with ID $[A:a_3]$ through MANET $[B:b_2:b_3]$. Hence, it will forward the packets to gateway b_2 .

Now suppose A is using DSR. Then a_1 will start route discovery, by employing route request to destination a_0 . When a_2 receives the route request, a_2 recognizes that there is an inter-domain route to a_0 with the updated route announcements from b_2 . Then a_2 acts as a proxy and replies the route request with the full path information from a_1 to a_2 . Using this information a_1 can send packets to a_2 , and a_2 forwards the packets to b_2 based on the inter-domain routing table.

The rest of the packet forwarding process is similar in MANET $[B:b_2:b_3]$ from b_2 to b_3 and then subsequently in MANET $[A:a_3]$ from a_3 to a_0 . If the domain is using OLSR, then the source will forward the packets to the destination based on the internal routing table. If the source is using DSR, then the source will first look at its route cache to determine a route to the destination, and optionally employ route discovery if a valid route is not found.

6 Evaluation

In this section, we present some preliminary evaluation of the effectiveness of the inter-domain routing in various MANET operation scenarios. We note that the main criteria to gauge the effectiveness of an inter-domain routing protocol is the end-to-end reachability between nodes. We study the impact of number of gateways, mobility patterns of the nodes, different number of domains on the reachability metric. We also provide an asymptotic overhead analysis of the proposed inter-domain routing protocol. The analytical result provides us an insight as to the relation between various operational parameters and the overhead incurred by the protocol.

6.1 Simulation Setting

In our simulations, nodes are randomly deployed in a 1000×1000 m² area. The transmission radius is 250 m, and the random walk mobility model was used. Unless otherwise specified, the average speed of the nodes is 5 m/s. For each set of simulation, we randomly select some portion of nodes as gateways, which can directly communicate with the gateways in other domain. Each node belongs to only one domain, and the nodes belonging to different domains cannot communicate to each other unless they are gateways, or there exists a path through gateways. For each set of simulation, we average our results from 5 simulations, each last for 5000 seconds simulation time.

6.2 Effectiveness of Inter-domain Routing

In this section, we study the effectiveness of inter-domain routing in MANET using the end-to-end reachability as the key evaluation metric. We define the reachability as the percentage of source-destination pairs that are connected either by intra-domain paths or inter-domain paths. Note that the source and destination may belong to different domains.




Figure 4: End-to-end reachability vs. % of gateways (with different number of nodes and node speed)

(1) **Density of Gateway:** Figure 4 presents the impact of having different number of gateways in the network (represented as the percentage of gateways among all nodes) to the end-to-end reachability. We present the case of two domains of equal size in terms of number of nodes. We plot the reachability results with different network size (25, 50, 100 total nodes) and different node speeds (average speed of 5 m/s and 20 m/s).

From this result, we make the following observations. First, as the percentage of gateways increases, the reachability increases rather quickly. This trend is especially visible in dense networks (i.e., 50 and 100 node case). In particular, for a network with 100 nodes, the reachability becomes almost 100% with only 10% gateway nodes. We note that only the 100 node network is fully connected when there is no inter-domain routing support (because it has 50% reachability when there is no gateway). On the contrary, when there are 25 nodes, they are not fully connected even when every node is a gateway. In general, we find that inter-domain routing improves the reachability of networks regardless of the network density.

Second, we observe that the speed of nodes does not affect the overall reachability much when we compare the case with node speed of 5 m/s and 20 m/s. This is because even when the nodes move fast, in a dense network, a regular node can find a path to a gateway, and a gateway can find a neighboring gateway to establish an inter-domain connection with reasonable probability. We note that this reachability result is an upper bound because in our simulation inter-domain route updates were instantaneous. In general,

increased node speed does not fundamentally affects the operation of inter-domain routing as long as the domain level topology does not change dramatically.




Figure 5: End-to-end reachability vs. % of gateways (with 100 nodes with different number of domains).

(2) **Number of Domains:** Figure 5 presents the reachability result for different number of domains when the total number of nodes were fixed to 100 nodes. We simulate the case with 2, 3, 4, 5 domains in the coalition. From the figure, we observe that the reachability generally decreases as the number of domains increases. This result comes from the fact as we increase the number of domains, the node density for each domain decreases such that some nodes may be isolated from the connected component of its own domain. On the other hand, when we fix the number of nodes in each domain, the number of domains does not affect the reachability result (not shown in this paper due to space constraints).

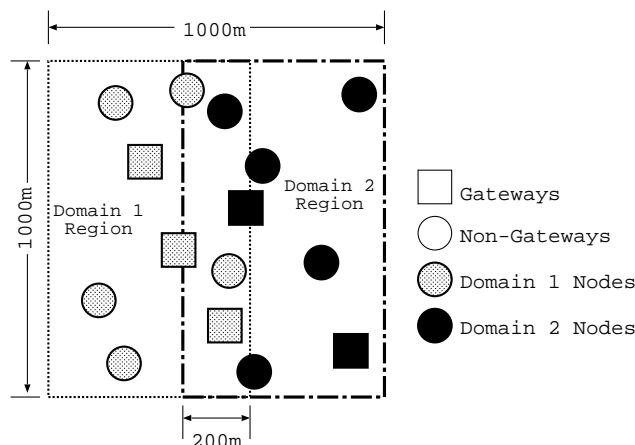


Figure 6: Nodes from two different domains are restricted to different sides of the area with overlap.

(3) **Group Mobility Pattern:** In the previous examples, we have studied the case when the nodes from different domains are deployed in the same area. However, in reality,

nodes may exhibit a group mobility or constrained mobility. In this section, we consider a very simple group mobility pattern as shown in Figure 6 where nodes from different domain are restricted to different sides of the field with some overlap.

Figure 7 presents the reachability with and without group mobility patterns. From the figure, we observe that the group mobility does not affect the reachability significantly in 50, 100 node cases. On a closer examination, we find that group mobility provides a better intra-domain connectivity for both 25 and 50 node cases (i.e., when there is no gateway). Overall, in the group mobility case, when the gateway density is low, the reachability suffers a bit because the probability of making inter-domain connection is relatively small. We expect that if the gateways were deployed near the border the reachability would increase in the group mobility case. We note that the full evaluation of inter-domain routing performance with more realistic mobility patterns will be an interesting future research topic.

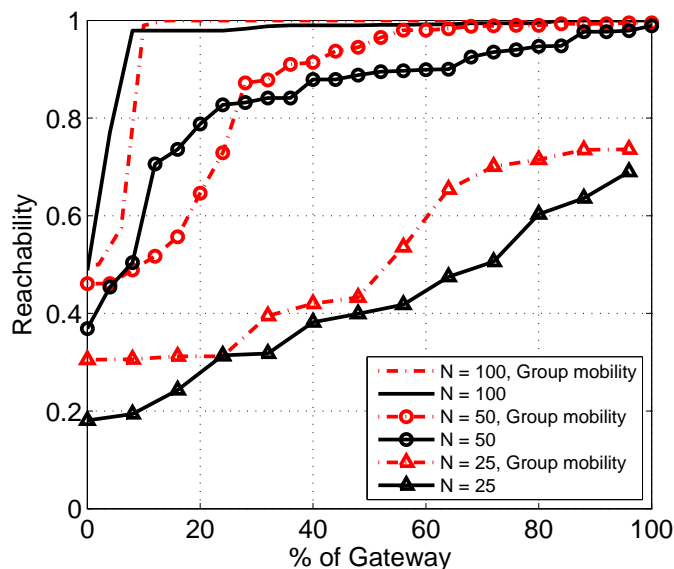


Figure 7: End-to-end reachability vs. % of gateways (with and without group mobility pattern).

6.3 Overhead Analysis

Next, we present a feasibility study by estimating the overhead of IDR. In the analysis, we assume that there is no congestion, so that no control packets are dropped or retransmitted, and the inter-domain routing policies are simple, so that gateways will not change existing forwarding paths except when the paths are disconnected. Our analysis aims to convey a basic estimated picture of the overhead, without involving the precise analysis at the detailed steps of protocols. It follows a similar analysis for proactive and reactive routing protocols in [4].

First, consider a single domain with the parameters as defined in Table 3. Assume that the mobility process of nodes is stationary and be confined to a bounded area. For a pair of nodes, if one node moves out of the transmission radius of other, then the link between them breaks. So, the average lifespan of a link is $\Theta\left(\frac{r}{v}\right)$, and the average number

Symbol	Defintion
N	Total number of nodes in a domain
G	The set of all gateways in a domain
r	Transmission radius
\bar{v}	Average speed of a node
\bar{E}	Average number of links in a domain

Table 3: Definitions of parameters

of link breakages per second due to mobility is $\Theta\left(\frac{\bar{v}}{r} \cdot \bar{E}\right)$.

Since the mobility process of nodes is stationary where there is no net links are created or broken over time, the average numbers of link creations per second due to mobility in the domain is also $\Theta\left(\frac{\bar{v}}{r} \cdot \bar{E}\right)$. Hence, the average number of link state changes (creations or breakages) per second is $\Theta\left(\frac{\bar{v}}{r} \cdot \bar{E}\right)$. The control overhead of intra-domain routing protocols is determined by the number of link state changes.

Next, we estimate the overhead for proactive intra-domain routing protocols, reactive intra-domain routing protocols, and inter-domain routing protocol, respectively.

(1) **Proactive Intra-domain Routing Protocols:** Each node periodically broadcasts hello packets to its neighbours. Based on the received hello packets, each node announces a new link-state/distance-vector packet that will be propagated throughout the MANET. Let λ^{hel} be the number of hello packets broadcast by each node per second. The total number of hello packets per second is $\lambda^{\text{hel}} \cdot N$.

Since the average number of link state changes per second is $\Theta\left(\frac{\bar{v}}{r} \cdot \bar{E}\right)$, the total number of link-state/distance-vector packets per second broadcast in the MANET is

$$O\left(\frac{\bar{v}}{r} \cdot \bar{E}^2\right).$$

This is an upper bound because optimized broadcast-based protocols (e.g., OLSR) normally requires less than \bar{E} transmissions for each link-state/distance-vector packet to propagate throughout the network. Thus, the estimated number of control packets per second for intra-domain proactive routing protocols is

$$\lambda^{\text{hel}} \cdot N + O\left(\frac{\bar{v}}{r} \cdot \bar{E}^2\right). \quad (1)$$

This is also the control overhead per second (at domain level) carried out by IDRMM to detect network partition and merging.

(2) **Reactive Intra-domain Routing Protocols:** IDRMM requires beaconing among gateways to detect network partition or merging. The number of gateway pairs that will beacon each other is upper bounded by $O(|G|^2)$. Let λ^{bea} be the beaconing rate between a pair of gateways. Then total number of beacons per second initiated by the gateways is

$$O(\lambda^{\text{bea}} \cdot |G|^2)$$

Let \bar{L} be the average number of hops between a pair of nodes in the MANET. The number of link state changes per second for a path between a pair of gateways is

$$\Theta\left(\frac{\bar{v}}{r} \cdot \bar{L}\right)$$

Since each link state change will incur maintenance overhead in reactive routing protocols, it is reasonable to assume that the number of control packets is proportional to the number of link state changes and the beaconing traffic. Hence, the estimated number of control packets per second required by IDRМ to detect network partition and merging is

$$O\left(\lambda^{\text{bea}} \cdot |G|^2 \cdot \frac{\bar{V}}{r} \cdot \bar{L}\right) \quad (2)$$

(3) **Inter-domain Routing Protocol:** Suppose there are m^{pro} domains running proactive routing protocols and m^{rea} domains running reactive routing protocols. Also assume each domain has the same parameters as in Table 3. Note that the path vector protocol in IDRМ behaves like a proactive routing protocols, but with different parameters. Let λ^{inter} be the number of inter-domain hello packets broadcast by each gateway per second in the path vector protocol. The total number of hello packets generated in the multi-domain MANET per second is $\lambda^{\text{inter}} \cdot (m^{\text{pro}} + m^{\text{rea}}) \cdot |\bar{G}|$, where $|\bar{G}|$ denotes average number of gateways in each domain.

If a pair of intra-domain gateways stay in the same MANET, there may be multiple paths connecting them. Let $\frac{1}{\mu}$ be the average lifespan of the connectivity between a pair of intra-domain gateways. That is, μ is the connectivity breakage rate of connected pairs of intra-domain gateways due to mobility. By stationarity of mobility process, μ is also the rate of change for the connectivity status of intra-domain gateways. Since IDRМ will carry out new membership management and announcement when the connectivity status between a pair of intra-domain gateways is changed, the estimated number of connectivity status changes is

$$O\left(\mu \cdot (m^{\text{pro}} + m^{\text{rea}}) \cdot |\bar{G}|^2\right)$$

Hence, the total number of control packets per second for path vector protocol is

$$\lambda^{\text{inter}} \cdot (m^{\text{pro}} + m^{\text{rea}}) \cdot |\bar{G}| + O\left(\mu \cdot (m^{\text{pro}} + m^{\text{rea}}) \cdot |G|^2 \cdot \bar{E}^{\text{inter}}\right) \quad (3)$$

where \bar{E}^{inter} is the average number of pairs of connected inter-domain gateways in the $m^{\text{pro}} + m^{\text{rea}}$ domains.

When a network is given, $m^{\text{pro}}, m^{\text{rea}}, \bar{G}, \bar{E}$, and λ^{inter} are fixed. However, it is not straightforward to estimate how μ will look like. Thus, we obtain this value from simulation. In Figure 8, we observe μ decreases as the number of nodes increases because as a MANET becomes denser, the connectivity between a pair of gateways becomes more stable, whereas node speed adversely affects the stability of links almost linearly.

Note that Eq.(3) only provides an asymptotic result for the *total* control overhead incurred by IDRМ without any optimization. Since the overhead will be distributed among all the gateways and various optimization can be applied (e.g., suppression of hello, adaptive adjustment of probing interval), the overhead incurred at each gateway for inter-domain routing operation will be quite moderate.

We compare this estimation to the normal routing overhead (not incurred by inter-domain routing). The overhead for proactive domains is Eq.(1), and the same for reactive domains is Eq.(2). Note that N and \bar{E} are typically orders of magnitude greater than the other parameters. Thus, the overhead from reactive domains and inter-domain operations are substantially small compared to proactive domains.




Figure 8: Average lifespan of the connectivity between a pair of intra-domain gateways.

To summarize, in a coalition MANET consisting of proactive and reactive domains, the overall control overhead will be dominated by that of the proactive domains, and the added overhead by the inter-domain routing protocol is expected to be moderate.

(4) **Comments on the Membership Digest:** To understand the overhead incurred by membership announcements, we compare two methods: (a) a plain membership list without any compression, and (b) a Bloom filter-based summary. We calculated the expected size of the plain membership digest in comparison with a more scalable membership digest using a Bloom filter for various network sizes from 100 nodes to 1000 nodes. For this calculation, we assumed 4 bytes for each node ID (e.g., IPv4 address) in the plain digest case. For the Bloom filter case, we considered three hash functions with a bit vector size 16000 to have a false positive rate smaller than 0.25%. In this case, for 500 nodes, both of them require about 2KB for membership announcement. For 1000 nodes, the plain digest required about 4KB and the Bloom filter required about 2KB resulting in about 50% space reduction. However, the benefit of the Bloom filter-based approach comes at the expense of extra processing for hash computation, and the non-trivial false positive rates. While it is possible to handle false positives (for instance, by sending explicit membership queries to the domains that claim to have the node when a false positive is suspected), we conclude that the extra overhead outweighs the space benefit in most MANET scenarios, and Bloom filter-based approach is valid only in very large networks (e.g., VANETs in a highly populated urban area).

7 Discussion

This paper has presented one viable approach to achieve inter-operation among multiple MANETs, and there are several interesting directions that we can further investigate.

(1) **Different Path Vector Protocol Options:** The proposed inter-domain protocol is based on a semi-proactive path vector protocol that only partially requires internal

network knowledge. One may also consider a fully reactive path vector protocol that does not require internal gateway detection or internal network knowledge. Compared with proactive routing protocols, in general, reactive protocols can handle dynamic network topology better, but also creates larger invocation delays when a route must be discovered. Proactive path vector protocol like BGP is more suitable for relative stable network topology (e.g., mesh networks) and time-sensitive data (e.g., real-time surveillance), whereas reactive path vector protocol is suitable for rapidly varying network topology (e.g., vehicular networks) and delay-tolerant data (e.g., non-emergency sensing). In this paper, we have proposed a semi-proactive path vector protocol so that we can strike a balance between the proactive approach and the reactive approach. However, the actual routing choice will be dictated by application scenarios and operation environments. It will be also interesting to explore an adaptive approach based on network dynamics. Table 4 summarizes different options for inter-domain operations.

Path Vector Protocol	Internal Gateway Detection	Internal Network Knowledge
Proactive	Required	Required
Semi-proactive	Required	Partially Required
Reactive	Not Required	Not Required

Table 4: Variants of path vector protocols.

(2) **Security Considerations:** Securing inter-domain MANET routing entails major future research. Firstly, the protocol itself should be secured in terms of authentication between peers. While similar techniques to those used for securing BGP may be employed [6], the key distribution costs may make them inappropriate. Secondly, while policy routing controls ingress, egress and transit of traffic, other resource management may be needed. For example, while topologically two MANETs might be disjoint, there are still radio resource implications if they topographically intersect and share spectrum. Thirdly, malicious nodes may not only inject false routing updates, but a subverted node may intercept unsecured routing control traffic and data traffic, and infer both routing policy and user’s data patterns. This leads to the most interesting potential avenue for future research since this can also be used by well behaved nodes to observe both routing exchanges and the effect of alleged routing policies by nearby neighbour nodes. For example, consistency checks can be made between observed routing messages and others reported to the witness nodes. Inconsistent updates from another node could result in witness nodes proposing the eviction of the bad node. Fourthly, we assumed that all gateways know the IDs of all nodes and other gateways within its own domain initially. But this may not always be true when the nodes are deployed in ad hoc manners. Hence, it requires a secure bootstrapping mechanism to discover and distribute these IDs.

8 Conclusion

Inter-domain routing offers a means for heterogeneous MA-NETs to interoperate. This paper has identified the challenges of inter-domain routing in MANETs, and proposed IDRM as a viable solution. A key contribution here is that despite dynamic network topology and diverse intra-domain ad hoc routing protocols, opaque interoperation among

heterogeneous multi-domain MANETs can be supported through IDRМ, a policy-based routing protocol that retains the merits of the widely used BGP. IDRМ not only improves the interoperability of MANETs, but also enhances the usefulness of today's MANETs for supporting future MANETs.

Acknowledgments

The authors gratefully acknowledge funding from the ITA Programme, and our partners in Honeywell and Roke Manor Research for valuable discussion on this topic.

References

- [1] M. Abolhasan, T. Wysocki, and E. Dutkiewicz. A review of routing protocols for mobile ad hoc networks. *Ad Hoc Networks*, 2:1–22, 2004.
- [2] A. Broder and M. Mitzenmacher. Network applications of bloom filters: A survey. *Internet Mathematics*, 1(4):485–509.
- [3] Y. Chen, A. Liestman, J. Liu. Clustering algorithms for ad hoc wireless networks. In *Proc. Ad Hoc and Sensor Networks '04*.
- [4] T. Clausen, P. Jacquet, and L. Viennot. Analyzing control traffic overhead versus mobility and data traffic activity in mobile ad-hoc network protocols. *ACM Wireless Networks journal (Winet)*, 10(4), July 2004.
- [5] J. Crowcroft et. al.. Plutarch: an argument for network pluralism. *ACM Computer Communication Review*, 33(4):258–266, 2003.
- [6] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. In *Proc. ACM SIGCOMM*, 2004.
- [7] D. McPherson et al. RFC 3345: Border Gateway Protocol persistent route oscillation condition.
- [8] V. Ramasubramanian, Z. J. Haas, and E. G. Sirer. SHARP: A hybrid adaptive routing protocol for mobile ad hoc networks. In *Proc. ACM MOBIHOC*, June 2003.
- [9] Y. Rekhter and T. Li. RFC 1771: a Border Gateway Protocol 4 (BGP-4), March 1995.
- [10] S. Schmid, L. Eggert, M. Brunner, and J. Quittek. TurfNet: An architecture for dynamically composable networks. In *Proc. of WAC 2004*, October 2004.
- [11] W. Ma, M. Chuah. Comparisons of inter-domain routing schemes for heterogeneous ad hoc networks. In *Proc. of WOWMOM '05*.