PAPER Special Section on Trust, Security and Privacy in Computing and Communication Systems

Using Regional Routing to Improve the Scalability and Security of Inter-Domain Multipath Routing*

Bin DAI[†a)], Feng WANG[††], *Nonmembers*, Baokang ZHAO[†], *Student Member*, *and* Jinshu SU[†], *Nonmember*

SUMMARY Multipath routing has been extended to Border Gateway Protocol (BGP), the current de facto inter-domain routing protocol, to address the reliability and performance issues of the current Internet. However, inter-domain multipath routing introduces a significant challenge for scalability due to the large scale of the inter-domain routing system. At the same time it also introduces new challenges in terms of security and security related overhead. In this paper, we propose a regional multipath approach, Regional Multipath Inter-domain Routing (RMI), where multiple paths are only allowed to be propagated within a well-defined range. With multipath routing in a region, we enable inter-domain routing with rich path diversity and improved security, and no longer have to sacrifice scalability. We show how to propagate multiple paths based on the region by theoretical analysis and by extensive simulations. Our simulations show that the number of messages generated using this approach and the convergence delay are much less than those of BGP and BGP with full multipath advertisement.

key words: *BGP, inter-domain multipath routing, network performance, prefix hijacking*

1. Introduction

Routing as core of the Internet is of great importance to successfully transmit packets from one end to another. A scalable, reliable and secure inter-domain routing method is critical when deploying new and large scale mission critical applications, such as Voice-over-IP (VoIP) applications, multiplayer games, and video conferencing. BGP (Border Gateway Protocol), the de facto inter-domain routing protocol, has been known to be slow to react and recover from network changes. Furthermore, BGP is known to be particularly vulnerable to a variety of mis-configuration and attacks.

Many efforts have been focused on improving the reliability and security of BGP. The most straightforward solution to address the reliability issue is to ensure rich path diversity. Multipath routing has been extended to BGP to improve the reliability and efficiency of the current Internet [4], [8], [15], [18], [20], [22], [24]–[26], [31]. However, since the inter-domain routing system is one of the largest distributed systems today, designing a scalable solution to provide multipath routing is challenging. There is concern about the overhead in multipath discovery and maintenance, which exponentially increases as the number of end users and their networks exponentially grow. Thus, if not carefully designed, the multipath solutions could exacerbate the scalability challenges.

Moreover, using multipath routing algorithms can introduce new challenges in terms of security and security related overhead. Although multipath routing can use end-to-end authentication and integrity to verify the paths [13], [16], [23], the cost of verifying multiple paths is higher than the cost to verify one path. Many non-cryptographic solutions attempt to reduce the computational overhead [5], [12], [14], [17], [19], [28], [29]. However, storing historical routing data and maintaining AS level topology in inter-domain multipath routing require a large storage overhead, which is impractical. Hence, a natural question arises: *Is it possible for inter-domain multipath routing to achieve reliability and security while remaining scalable?*

In this paper, we propose a novel inter-domain routing protocol based on *regional routing*, which is referred to as Regional Multipath Inter-domain Routing (RMI) protocol. The primary design goal behind RMI is to balance reliability gains with scalability and security in multipath routing. Here, being scalable means that the multipath overhead at each node must grow very slowly with the increase of the network size. Previous multipath based inter-domain routing methods may generate many unnecessary paths that could be too far away from the end users for them to use. To achieve the goal of linear scalability, we exploit the structure of an AS's *neighborhood region*, which is defined as a collection of an AS's providers and customers, and is used to capture provider-level and customer-level topology information. Essentially, in RMI, multiple paths are only allowed to be propagated within an AS' neighborhood region. ASes outside the region only have summary routing information about the regional paths. By analyzing the routing updates archived in RouteViews [2], we find that the size of neighborhood regions is usually very small compared with the size of the Internet, which makes it possible for RMI to advertise multiple paths without impacting the scalability. Our simulation results show that RMI is able to reduce the convergence delay, and decrease the amount of routing updates advertised throughout the network.

To achieve the goal of improved security, we exploit neighborhood region to impose the consistency between routing information and its associated neighborhood re-

Manuscript received March 15, 2011.

[†]The author is with the School of Computer, National University of Defense Technology, Changsha, 410073 China.

[††]The author is with the Faculty of School of Engineering and Computational Sciences, Liberty University, Lynchburg, VA 24502, USA.

[*]This work was supported by the National Basic Research Program of China (973 Programs) (No.2009CB320503), and NSF of China (No.61070199 and No.61003301).

a) E-mail: bin.danieldai@gmail.com

DOI: 10.1587/transinf.E95.D.94

gions. RMI provides an efficient and light-weight approach to prevent attacks and mis-configuration. In this paper, we consider two common types of Inter-domain routing attacks, namely invalid path attacks and invalid origin attacks. We show that RMI can successfully detect invalid origin attacks, and confine the possible invalid path attacks to a very small region. Note that in BGP a wide range of attackers can launch invalid route attacks. Thus, if combined with RMI, other proposed approaches, such as regular public/private key, can be used to provide path validation without imposing too much computational overhead.

In summary, RMI is a dirty slate solution to scale inter-domain multipath routing. That is, RMI is based on path vector routing. Because RMI inherits many features of BGP, it can support most of BGP's routing policies, for example, import and export policies. In addition, in RMI, providers forward multipath only to and from their customers, so that they have the incentive to implement regional multipath for their customers.

The rest of the paper is organized as follows. In Sect. 2, we introduce the basic concept of Regional Path Vector. In Sect. 3, we introduce multipath routing based on Regional Path Vector. We discuss the security improvement by utilizing region based multipath in Sect. 4. In Sect. 5, we evaluate the performance of RMI. In Sect. 6, we survey some previous work on solving the scalability of inter-domain routing. We conclude in Sect. 7.

2. Regional Path Vector

In this section, we focus on illustrating single Regional Path Vector (RPV) routing, in which each AS advertises only one single path. We first introduce the concept of neighborhood region, which forms the key design principle of RPV. After that, we introduce regional path vector routing, which consists of intra-region routing and inter-region routing. Finally, we provide analysis of RPV. In the next section, we introduce RMI, which advertises multiple regional paths. Note that in the rest of paper, we use terms *route* and *path* interchangeably.

2.1 Neighborhood Regions

In a *hierarchical inter-domain system*, the neighbors of an AS can be classified as providers, customers or peers according to their commercial agreements. We consider two types of AS relationship: *provider-to-customer* and *peer-to-peer*. Typically, provider-to-customer relationships among ASes are hierarchical. The hierarchical structure arises because an AS typically selects a provider with a network of larger size and scope than its own. A *direct* neighbor is defined as an AS with whom direct communication can be established, and is thus one hop away. For example, in Fig. 1, AS D is the provider of AS B, which in turn is the provider of AS A. We call AS B as a *direct* provider of AS A, and AS D as an *indirect* provider of AS A. Similarly, AS A is a *direct customer* of AS B, and an *indirect* customer of AS D.

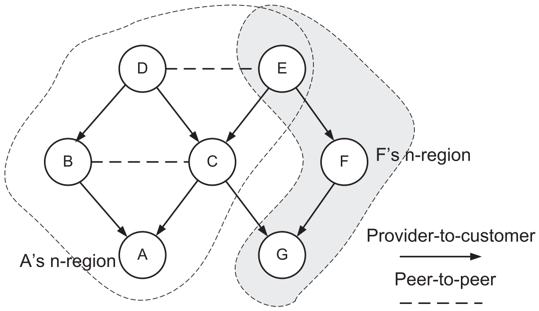

Fig. 1 An example of neighborhood region.

In RPV, an AS maintains routes to a destination within a local neighborhood, which we refer to as the AS's *neighborhood region*, or "n-region". More precisely, an AS's neighborhood region is defined as a collection of ASes which are direct or indirect providers or customers of the AS. An AS's n-region is composed of *provider region* and *customer region*. Provider region is considered to be a group of direct and indirect providers, while customer region is a collection of direct and indirect customers. In this paper, we use *P_Region*(*A*) to denote AS A's provider region.

Figure 1 illustrates the neighborhood regions at AS A and AS F. In this example, AS A's provider region contains 4 providers: B, C, D, and E. Since AS A is a stub AS, it does not have any customer region. AS F has one provider and one customer so that its provider region and customer region have only one AS.

2.2 Neighborhood Region Discovery

The construction of n-region requires an AS to first know who its direct providers and customers are. The providers and customers can be inferred from the commercial contractual agreements. After that, the customer region can be easily derived from the routes sent by its direct customers. Suppose that there is an AS path (A B C) learned from one customer. In this case, A, B, and C are all customers because of no-valley routing policy. No-valley policy means that a customer can only forward its own or its customers' prefixes to providers. On the contrary, the provider region cannot be inferred from the routes sent by direct providers. The reason is that the providers may advertise a route that traverses peer-to-peer or customer-to-provider links. Therefore, the ASes in the route may not necessarily be its providers. For example, suppose that an AS learns a path "A B C" from its provider A. From this path, the AS only knows that AS A is its provider because AS B and AS C could be the customers of AS A.

In RPV, identification of an AS's provider region is implemented through a separate Provider Discovery Protocol (PDP). Here, we describe the basic idea of the protocol. In the next section, we will present the details. The process starts from tier-1 ASes, the ASes that don't have any provider. Each tier-1 AS periodically broadcasts a beacon prefix to its customers. On receiving an announcement about the beacon prefix from the providers, each AS constructs its own provider region, appends its AS number to the received route and sends it to its customers. For example, in Fig. 1, AS D and AS E are tier-1 ASes, and each of them broadcasts a beacon prefix to its customers. AS A receives two messages generated by AS D, and one message from AS E. Based on those beacon messages, AS A knows that its provider region contains four ASes: B, C, D and E.

Note that a n-region is defined regarding each AS. Each stub AS only has one provider region, while a transit AS has both regions. Tier-1 ASes only have customer regions. Each AS discovers and maintains its own n-region. To let other ASes determine the propagation range, each AS needs to distribute the information. In RMI, only the information about provider region, rather than the whole n-region, is needed to disseminate to other ASes. That means, each AS does not need to distribute the information about its customer region. This is because of two reasons. First, as we shall see later, the information about provider region is sufficient for other ASes to determine the propagation range for the routes to the destination. Second, from our measurement based on today's Internet, which will be shown in Sect. 5, we know that the size of provider regions is much smaller than that of customer regions. Therefore, the number of routing announcements does not increase significantly. Furthermore, the provider list can be implemented by using a Bloom Filter to reduce the size of messages, which will be discussed in Sect. 3.

2.3 Dissemination of Provider Region

Here, we consider how to distribute the information about provider region. As mentioned above, RPV is based on path vector. Just like BGP, each AS advertises its prefixes to all its neighbors, including its providers, customers, and peers. Before an AS announces a prefix, the AS first generates a list containing its providers, which is based on its provider region, and sends the list with the prefix announcement. We call the list as *provider list*. Here we emphasize that the provider list should include a *complete* list of providers. At normal situation, each AS can obtain all its providers from PDP. However, in some cases the list only contains a subset of providers. For example, in BGP, an AS may connect to multiple providers via static configuration or using a private AS number, which results in the Multiple Origin Autonomous System (MOAS) [30]. From other ASes' perspective, it appears as if the routes have multiple origin ASes, which are its providers' AS numbers. The same cases can occur in RPV. To correctly distribute provider region, RPV requires each of the providers to advertise the customer's provider list, not its own provider list.

For example, in Fig. 2 AS 1 uses a private AS number, and has two providers. This results in the AS's prefix as if is AS 2 and AS 3's prefix. In RMI, when AS 2 and AS 3 announce AS 1's prefix, they need to advertise AS 1's provider list, rather than their own lists. In this example, the provider list for AS 1 contains AS 2, AS 3, and AS 4. On the contrary, AS 2 and AS 3's provider lists contain AS 4 only.

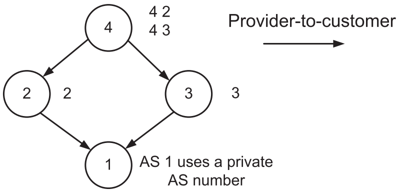

Fig. 2 An example of multiple origins. AS 1 uses a private AS number.

2.4 Intra-Region Routing

In this section, we discuss how to setup the routes to reach the destination inside a n-region. Inside an AS's n-region, a route destined to the AS is referred to as an *intra-region route*. An intra-region route is similar to a route in BGP, which contains a *regional path* to reach the destination. In this paper, we denote an intra-region route as (u_i u_{i-1} ... u_0), where u_0 is the origin AS. An intra-region route has the following property:

Property 1. *For a given intra-region route (u_i u_{i-1} ... u_0) to a destination originated by AS u_0, AS u_i is a provider or a customer of AS u_0.*

For example, in Fig. 1, A has a regional path (A C E) to reach a destination originated at E, and E has another regional path (E F G) to reach G. AS A is a customer of the origin AS E, and E is a provider of AS G.

When an origin AS advertises its prefixes, according to the type of its neighbors the AS may construct different types of routes. When advertising a prefix to a provider or a customer, the AS should generate an intra-region route. In particular, the origin AS generates a prefix announcement, inserts its AS number into the regional path, and sends it to its providers and customers. Each intra-region route advertisement concerns a particular prefix and includes a provider list of the origin AS. On the contrary, if it has peer-to-peer AS relationship with a neighbor, the origin AS has to construct another type of prefix announcement–an inter-region route, which will be introduced in the next section.

After an AS announces a prefix, other ASes construct the intra-region routes to it by successively propagating the advertisement between pairs of ASes. Just like BGP, before accepting an intra-region route, the receiver checks for the presence of its own AS number in the path to avoid routing loops. And then, the AS must decide whether or not to use this route. When an AS has several intra-region routes to the destination, the AS needs to select the best intra-region route. Note that in this section we focus on the single path routing. Since RPV is based on path vector, we can use the same BGP route selection and apply the same routing policies to it. In this paper, we assume that every AS applies *typical routing policies* [10]. That is, an AS announces its customer routes to all neighbors but its peer or provider routes to its customers only. Besides, every AS prefers its customer routes over its peer routes and then over its provider routes.

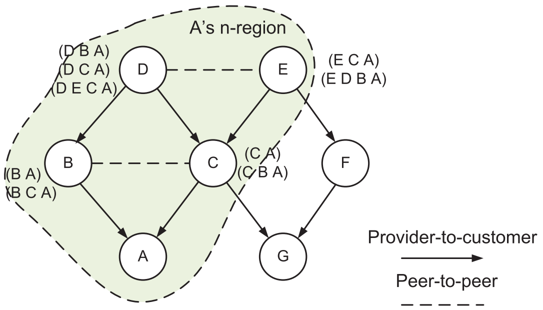

Fig. 3 An example of intra-region routes. The AS paths around a node represent the available paths in the node's routing table, which are ordered in the descending order of local preference.

After finishing the best intra-region path selection, the AS continues to determine whether or not to propagate the route to neighboring ASes (after adding its own AS number to the AS path), which is determined by its export policies and the origin AS's n-region. If the route advertisement is allowed by its export routing policies, the AS advertises the route according to the following dissemination rules for intra-region routes.

Rule 1 (Uphill)**.** *Each AS directly sends intra-region routes to its providers.*

The Uphill rule implies that the AS can directly forward an intra-region route to its uphill providers without examining the provider list coming with the prefix. In this case, the intra-region routes must come from the AS's customers. Otherwise, due to no-valley policy, the intra-region routes are not allowed to be sent to providers. We use the example shown in Fig. 3 to illustrate the Uphill rule. The example has the same network topology as the previous example (in Fig. 1). AS A advertises a prefix to AS B and AS C. According to the Uphill rule, AS C and AS B send the route to their provider AS D directly.

Rule 2 (Origin AS's Provider)**.** *Each AS sends intra-region routes to its neighbors that are the providers of the origin AS.*

An AS can advertise an intra-region route to a neighbor who is not its provider. The condition is that the neighbor must be the origin AS's provider. For example, in Fig. 3, AS D wants to forward an intra-region route learned from AS B to AS C. Since AS C is not AS D's provider, AS D continues to determine if AS C is in AS A's provider list. Searching the provider list, AS D finds that AS C is a provider of AS A. Consequently, AS D sends the intra-region route to AS C.

Rule 3 (Origin AS's Customer)**.** *Each AS sends intra-region paths to its customers that are the customers of the origin AS.*

This rule implies that an AS cannot just simply advertise an intra-region path to its customers. Instead, the AS needs to ensure that the customers are the origin AS's customers as well. We use an example to demonstrate this rule.

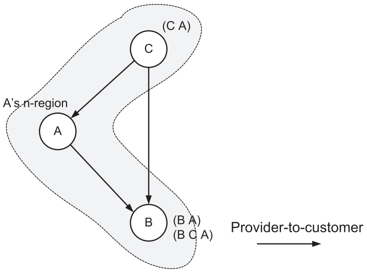

Fig. 4 An example of Origin AS's Customer Rule (Rule 3). AS C determines the origin AS's customer region by using AS B's provider list.

In Fig. 4, AS A originates a prefix, and advertises it to AS B and AS C. Before AS C sends an intra-region path to AS B, it first checks if AS B is in the origin AS's provider list according to Rule 2. AS C finds that AS B is not AS A's provider. We know that AS B is AS A's customer so that it is in AS A's n-region. However, only the origin AS has the information about its customer region, which is not visible to other ASes. To solve the problem, AS C uses AS B's provider list, instead of AS A's provider list, to investigate if AS A is AS B's provider. More specifically, AS A first finds an intra-region path originated by AS B. And then, AS B investigates the provider list associated with the path to see if AS A is in the list. In this example, AS B's provider list contains A and C so that AS C can determine that AS B is a customer of the origin AS. This example shows that each AS can determine if a neighbor is in the origin AS's n-region just based on the provider lists.

Rule 4 (Tier-1 ASes' Prefix to Customers)**.** *Each AS sends intra-region paths to its customers if there is no available provider list associated with the paths.*

Tier-1 ASes do not have any provider so that they do not have provider region. In this case, other ASes simply forward the intra-region path destined to a tier-1 AS to its customers. For example, in Fig. 3, suppose that AS D originates a prefix and advertises it to AS B and AS C. AS B and AS C just forward the intra-region path to their customer, AS A. Note that AS B and AS C cannot forward the path to each other because of no-valley routing policy.

2.5 Inter-Region Routing

In this section, we discuss how to announce the routing paths to an AS that is located beyond the origin AS's n-region. As we mentioned above, if an intra-region route is not allowed to send by the intra-region route dissemination rules, the AS needs to construct an *inter-region route* to summarize the intra-region route, and advertise the new route to the neighbor.

Different from an intra-region route, an inter-region path is composed of a *Source Regional Path* (SRP) and a *Summary Path Metric* (SPM). When an AS needs to advertise a route outside the origin AS's n-region, the AS constructs an inter-region route to replace the intra-region route.

The AS summarizes the intra-region route by using a summary path metric, such as the shortest path distance or propagation delay. Similar like an intra-region path, a SRP is a sequence of ASes along which the path receiver can reach the bounder of the origin AS's n-region. When an AS originates an inter-region route, the AS inserts its AS number into the SRP. In this paper, we denote an inter-region route by "$[(P) : n]$", where P is the SRP and n is the summary path metric.

As we described before, an AS uses an intra-region route to announce its own prefixes to its providers and customers. When the AS's neighbor is a peer, the AS uses an inter-region route to advertise the prefixes to the neighbor. In this case, the origin AS inserts its AS number into SRP, and assigns 0 to SPM, which indicates the prefix originated by the AS. After that, the AS sends the route to the neighbor. The neighbor must be either the AS's *peer* or *customer*, but cannot be a provider. Based on this, we find that an inter-region route has the following property:

Property 2. *For a given inter-region route* $[(u_i \ldots u_{k+1} u_k){:}n]$ *to a destination originated by AS* u_0*, AS* u_k *is a provider of AS* u_0*, and AS* u_{k+1} *is a customer or a peer of AS* u_k

Note that when an AS originates an inter-region route, provider list, which is coming with the corresponding intra-region route *should be sent with an inter-region path*. There are two reasons behind this. First, even though ASes are not willing to reveal their provider-customer relationships, in practice it is possible to infer most provider-customer relationships from routing updates. Second, in Sect. 4, we will show that the provider list can be used by other ASes to validate routing information.

After an AS constructs an inter-region route and sends the route, we next determine how other ASes forward the route. Upon receiving an inter-region route, each AS first appends its AS number into the SRP, and then follows the following rule:

Rule 5 (Downhill Only)**.** *Each AS sends inter-region routes to its customers only.*

An AS forwards an inter-region route only to its customers without examining the origin AS's n-region or its n-region. Since an inter-region route is only allowed to be advertised to customers, the route is composed of a sequence of providers. In addition, each AS still uses the same BGP route selection to select the best inter-region route.

For example, in Fig. 5, AS C advertises an inter-region route $[(C) : 1]$ to AS G, and AS E sends route $[(E) : 2]$ to AS F. Note that in this example and following examples, we use the shortest path distance to summarize an AS's intra-region path. After AS F receives the inter-region route from AS E, it adds its AS number to SRP, and forwards the route to AS G.

2.6 Correctness of RPV

First, we prove that the intra-region route dissemination rules, which are described in previous section, can ensure the propagation of intra-region routes within the origin AS's n-region.

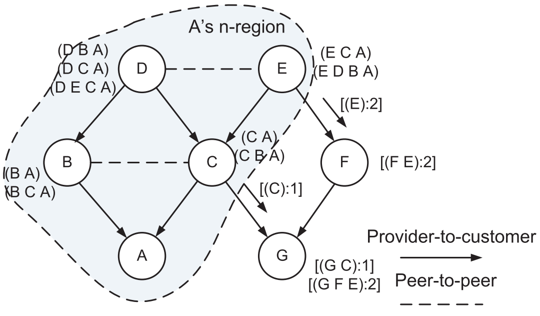

Fig. 5 Example of inter-region routes.

Theorem 1. *If every AS follows the intra-region route dissemination rules and typical routing policies, intra-region routes are propagated within the origin AS's n-region.*

Proof. We prove this by contradiction. Suppose that AS u_i has an intra-region route ($u_i\ u_{i-1}\ \ldots\ u_1\ u_0$). We assume that $u_{i-1}, \ldots, u_1$, are in the origin AS u_0's n-region, but u_i is outside the region, or AS $u_{i-1}, \ldots, u_1 \in P\_Region(u_0)$ but $u_i \notin P\_Region(u_0)$. Since AS u_0 advertises intra-region paths only to its providers and customers, we consider the two cases:

Case 1: u_1 is u_0's provider. After u_0 advertises its prefix to its provider u_1, because of no-valley property of Internet paths, the advertisement can traverse one or more customer-to-provider links followed by zero or one peer-to-peer link or one or more provider-to-customer links to u_i. According to Uphill rule, the ASes along a series of customer-provider links belong to u_0's region. Thus, the intra-region path must have a peer-to-peer link or one or more provider-to-customer links to reach u_i. If u_i has a peer-to-peer link to u_{i-1}, according to Origin AS's Provider rule, u_{i-1} ensures $u_i \in P\_Region(u_0)$. If u_i has a provider-to-customer link to u_{i-1}, according to Origin AS's Customer rule, u_{i-1} ensures $u_i \in P\_Region(u_0)$. The two cases contradict the assumption $u_i \notin P\_Region(u_0)$.

Case 2: u_1 is u_0's customer. After u_0 advertises its prefix to u_1, the advertisement can traverse one or more provider-to-customer links to u_i. According to Origin AS's Customer rule, u_{i-1} ensures $u_i \in P\_Region(u_0)$, which contradicts the assumption.

The two cases show that there is not such intra-region route outside the origin AS's n-region. □

Second, we prove that following the regional route dissemination rules, including intra-region route rules and inter-region route rules, RPV is loop-free.

Theorem 2. *If every AS follows the intra-region route and inter-region route dissemination rules, and every AS applies typical routing policies, then RPV is free of loops at every instant.*

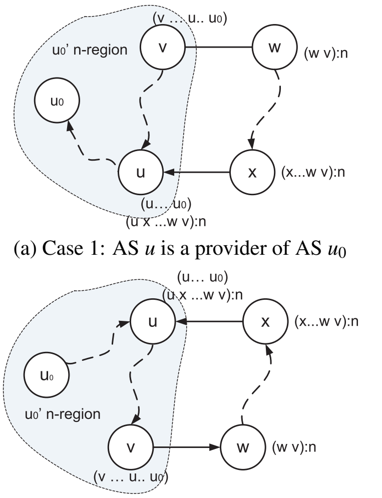

(a) Case 1: AS u is a provider of AS u_0

(b) Case 2: AS u is a customer of AS u_0

Fig. 6 Proof of Theorem 2. A straight line represents either provider-to-customer or peer-to-peer AS relationship. A straight arrow line denotes provider-to-customer relationship. A curved arrow dashed line represents a sequence of ASes that each pair of them has provider-to-customer relationship. The distance between AS v and the origin AS is n hops.

Proof. Since path vectors are used to propagate routing information within an origin AS's n-region, there is no loop within the region. Thus, we need to prove the case that it is impossible for every AS in any intra-region route to receive the corresponding inter-region path that summarizes the intra-region route. We show this by contradiction. We assume that there is an intra-region path that is advertised by AS u and sent to AS v. As shown in Fig. 6, AS v constructs the corresponding inter-region path and sends it to AS w. Finally, The path is advertised to AS u by AS x. Because AS u has an intra-region route, we consider the two cases: 1) AS u is a provider of AS u_0, and 2) AS u is a customer of AS u_0.

Case 1: AS u is a provider of AS u_0. As shown in Fig. 6 (a), AS u finally receives an inter-region path $[(ux \dots wv) : n]$, which hides the path from v to the origin AS. Thus the inter-region path contains a loop. According to Rule 5 (Downhill Only), AS x must be AS u's provider. Since AS u is a provider of AS u_0, AS x must also be an indirect provider of AS u_0. Subsequently, AS w must be AS u_0's provider. Both AS w and AS x must be in AS u_0's n-region. However, AS v constructs and sends an inter-region path to w, which contradicts Rule 1 (Origin AS's Provider).

Case 2: AS u is a customer of AS u_0. As shown in Fig. 6 (b), AS u receives an inter-region route $[(ux \dots wv) : n]$ from AS x, which contains a loop. We now look at the construction of the inter-region route. First, AS u has an intra-region route and sends to AS v. Since the route is from AS u's provider, according to no-valley routing policy, AS u can only send the route to its customers. Thus, AS u must be AS v's provider. Similarly, AS v cannot advertise the intra-region path from AS u to other providers or peers. As a result, AS w must be AS u's customer. And then, AS v constructs the inter-region route, and sends to it customer, AS w. According to Rule 5 (Downhill Only), AS x must be AS v's customer either. When AS x advertises the inter-region path to AS u, AS u must be AS x's customer. As a result, AS u is a customer of AS v, which results in a contradiction.

The two cases show that there is not such inter-region path containing a loop. Therefore, RPV is loop-free. □

2.7 Diversity of Regional Path

The diversity of regional route means the types of region paths that each AS may have. Recall that there are two types of region routes in RPV: intra-region routes and inter-region routes. Understanding the path diversity can help each AS to validate routing information. Here, we focus on understanding the relationship between path diversity and provider region. In the next section, we will show how to utilize the provider region concept to achieve path validation.

To investigate the relationship, in the following discussion we consider a scenario where an origin AS, AS u_0, advertises its prefix to its neighbors. We investigate the region path diversity at AS u_i, which receives one or several paths to reach the origin AS.

First, we consider the case that AS u_i belongs to AS u_0's n-region. We have the following claims to describe the regional paths at AS u_i.

Claim 1 (All Intra-region Routes at Provider)**.** *If AS u_i is a provider of an origin AS u_0, AS u_i's routes to reach the destination originated by AS u_0 must be all intra-region routes.*

Proof. Since AS u_i is a provider of the origin AS, according to Uphill rule, AS u_i must have at least one intra-region route to reach AS u_0. Thus, we need to prove that all of AS u_i's routes are intra-region routes if it has more than one route. We show this by contradiction. We assume that AS u_i has an inter-region route $[(u_i\ u_{i-1}\ \dots\ u_{k+1}\ u_k) : n]$. Here, u_k originates an inter-region path, and the distance from u_k to the destination is n. According to Property 2, AS u_k is a provider of the origin AS. At the same time, the region path $[(u_i\ u_{i-1}\ \dots\ u_{k+1}\ u_k) : n]$ implies that u_k is a provider of u_{k+1}, which in turn is a provider of u_i. Since both AS u_i and AS u_k are providers of the origin AS, u_{k+1} must be a provider of the origin AS's n-region either. As a result, AS u_k needs to send an intra-region path instead of an inter-region path to u_{k+1}, which results in a contradiction. Thus, the inter-region route does not exist. □

Claim 2 (Hybrid Routes at Customer)**.** *If AS u_i is a customer of an origin AS u_0, AS u_i's routes to reach the destination originated by AS u_0 must be either all intra-region routes, or some of them but not all are inter-region routes.*

Proof. Since AS u_i is a customer of the origin AS, due to no-valley routing policy, AS u_i can obtain routes to reach the destination only from its providers. Suppose that AS u_i has a set of direct providers, $p_1, p_2, \dots, p_k$. If all those providers are located in the origin AS's n-region, according to Rule 3 (Origin AS's Customer), all the routes via those providers are intra-region paths. On the other hand, among the *k* providers, we assume that there are *m* providers are located outside the origin AS's n-region. According to Rule 3 (Downhill Only), the paths via the *m* providers to reach the destination are inter-region paths. □

Based on the above two claims, we can prove the theorem about the case that an AS has both inter-region routes and intra-region routes to reach the same destination.

Theorem 3. *If AS u_i has both inter-region routes and intra-region routes to reach the same destination originated by AS u_0, AS u_i must be a customer of AS u_0.*

We also present one theorem regarding the case that an AS has all inter-region routes to reach the same destination.

Theorem 4. *AS u_i's all routes to reach the same destination originated by an origin AS u_0 are inter-region routes if and only if AS $u_i \notin P\_Region(u_0)$ and AS $u_0 \notin P\_Region(u_i)$.*

Proof. (Sufficiency) if AS $u_i \notin P\_Region(u_0)$ and AS $u_0 \notin P\_Region(u_i)$, we need to prove that all AS u_i's path to the destination are inter-region paths. Since AS u_i is not in AS u_0's provider region and AS u_0 is not in AS u_i's provider region, AS u_i can only get paths to u_0 from AS u_0's providers due to no-valley policy. According to intra-region path advertisement rules, AS u_0's providers will send inter-region paths to AS u_i.

(Necessity) if AS u_i's all paths to reach the origin AS u_0 are inter-region paths, according to intra-region path advertisement rules, AS u_i is not in AS u_0's provider region. Next, we need to prove that AS u_0 is also not in AS u_i's provider region. We show this by contradiction. Assume that AS u_0 is in AS u_i's provider region, or AS u_i is in AS u_0's customer region. According to Claim 1 and 2, AS u_i must have intra-region routes from AS u_0, which contradicts the necessary condition. □

Finally, we have a theorem about the case that an AS must have at least two intra-region routes to the same destination. The theorem will be used to validate the consistency between regional routes. We will show this in detail in Sect. 4.

Theorem 5. *If AS u_i has an intra-region route ($u_i\ u_{i-1}\ \dots\ u_0$) to a destination originated at AS u_0, and has peer-to-peer AS relationship with AS u_{i-1}, AS u_i must be a provider of the origin AS and must have another intra-region route via its customer to the same destination.*

Proof. According to no-valley policy, AS u_{i-1} must be a provider of the origin AS u_0. Otherwise, it cannot advertise the path to its peer, AS u_i. Since AS u_i has an intra-region path from AS u_{i-1}, based on Property 1, we know that AS u_i must be either a provider or a customer of the origin AS. Let's first consider the case that AS u_i is a customer of the origin AS u_0. Since AS u_{i-1} is a provider of the origin AS, AS u_i should be a customer of AS u_{i-1} as well. However, the AS relationship between AS u_i and AS u_{i-1} is peer-to-peer relationship. It is very unlikely that an AS's customer would be a peer of the AS. Therefore, we do not consider this case. When AS u_i is a provider of the origin AS, it must traverse a set of provider-to-customer links to reach the destination. So, AS u_i must have another path via its customer rather than AS u_{i-1} to reach the destination. □

3. Regional Multipath Inter-domain Routing

In previous section, we introduce RPV, which is based on single path routing. In this section, we extend multipath routing to RPV. The goal of incorporating multipath routing in RPV routing is to increase path diversity. We first present the overview of RMI, and focus on the benefit of RMI on the scalability of multipath routing. And then, we present RMI implementation.

3.1 Overview of RMI

We use an example to show that RMI can achieve rich path diversity with an acceptable number of routing messages. As depicted in Fig. 7, suppose AS 1 is the origin and all ASes advertise all their known routes to their neighbors. As a result, AS 5 has three routes to reach AS 1, that are (5 3 4 1), (5 3 1) and (5 3 2 1). And path (5 3 1) is used as the primary path. In RMI, AS 5 summarizes these three routes into one path by using shortest path distance: [(5):2]. AS 9 has two routes: [(9 7 6 5):2] and [(9 8 6 5):2]. On the contrary, without RMI, AS 9 will have 6 routes, which are too far away from AS 9 for it to use.

If link 3 − 2 fails, AS 5 doesn't need announce any information to AS 6. Thus, both the message overhead and the convergence time could be reduced dramatically in the large network.

3.2 RMI Design

RMI consists of two parts. First part of RMI is an algorithm, Provider Discovery Protocol (PDP), for each AS to get its provider region. The second part of RMI is the region route distribution rules.

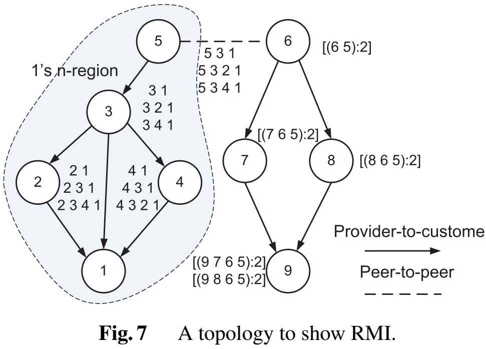

Fig. 7 A topology to show RMI.

3.2.1 PDP Implementation

The process of generating provider region starts from tier-1 ASes. Each tier-1 AS advertises a beacon prefix, called Topology Beacon, and sends the prefix in a message called Topology Information Packet (TIP) to its customers. The TIP contains three fields: 1) the source field indicating the tier-1 AS's beacon prefix, 2) a sequence number field to distinguish the TIPs received at different time, and 3) a list recording the providers that the TIP traverses. On receiving a TIP from the providers, each AS constructs its own provider region, appends its AS number to the received TIP and sends it to its customers. We show the procedure in Algorithm 1.

Algorithm 1: Procedure of PDP.

```
Set current AS = $M$
if $P\_Region(M)$ is empty then
    if $M.provider\_set$ is empty then
        Construct $P\_Region(M)$
        Construct $TIP$
        foreach AS $P$ in $M.customer\_set$ do
            Send $TIP$ to AS $P$
        end
    end
    if $M.provider\_set$ is non empty then
        if $M$ receives all $TIP$s from $M$'s providers then
            Construct $P\_Region(M)$
            Construct $TIP$
            foreach AS $P$ in $M.customer\_set$ do
                Send $TIP$ to AS $P$
            end
        end
    end
end
if $P\_Region(M)$ is non empty then
    if any link state in $P\_Region(M)$ is changed or $M$
    receives a new $TIP$ then
        Construct $P\_Region(M)$
        Construct $TIP$
        foreach AS $P$ in $M.customer\_set$ do
            Send $TIP$ to AS $P$
        end
    end
end
```

4. Security Improvement based on Provider Region

In Sect. 2, we investigate the relationship between the regional path diversity and provider region. We understand that intra-region routes and inter-region routes are restricted by the origin AS's provider region and the receiver's region. In addition, the provider list coming with a route implies the AS relationships between the origin AS and its neighbors. In this section, we investigate how to utilize the relationship to detect mis-configuration and routing attacks. We consider two common types of Inter-domain routing attacks: invalid path attacks and invalid origin attacks. In RMI, each AS checks the consistence between a receiving route and the provider list associated with the destination to validate routing information. Next, we present the guidelines for consistence check.

4.1 Consistence Check Guidelines

After an AS has received a set of regional routes, including intra-region or inter-region routes, the AS can validate those routes by examining if those regional routes are consistent with the provider regions associated with the destination. More specifically, when an AS u has received a set of regional routes, AS u uses the following guidelines, which are based on the analysis results in Sect. 2, to check the consistence.

1. If AS u has more than one route to the same destination, the provider lists associated with those routes must be consistent.
2. For each intra-region route, the provider list associated with the route must contain AS u, or the origin AS must be in AS u's provider region (based on Property 1).
3. For each inter-region route, the AS that generates the inter-region route must be in the origin AS's provider region (based on Property 2).
4. An inter-region route must come from AS u's providers or peers, but not its customers (based on Downhill Only rule).
5. If AS u has both inter-region routes and intra-region routes to the same destination, AS u's provider region must contain the origin AS (based on Theorem 3).
6. If AS u's paths to a destination are all inter-region routes, the origin AS should not be in AS u's provider region and AS u itself should not be in the provider region list associated with those paths (based on Theorem 4).
7. If AS u's intra-region route via a peer v, AS u must have another intra-region route via its customer. At the same time, if the intra-region route is via a provider p, the provider p must be in AS u's own provider region (based on Theorem 5).

If a received route violates any of the guidelines above, the route is considered to be invalid. In the rest of this section, we describe how to use those guidelines to detect mis-configuration and attacks.

4.2 Mis-configuration Detection

The above guidelines can be used to detect faulty configuration when an AS incorrectly advertises an intra-region route or inter-region route. One possible faulty configuration is due to the false positive of provider list. Since RMI aggregates provider region information into Bloom Filters, we need to consider the impact of the false positive caused by Bloom Filters. In RMI, the false positive means that an AS mistakenly considers one neighbor inside the origin AS's n-region, which actually is outside the region. In other words, the AS incorrectly sends an intra-region route instead of an inter-region route to the neighbor. On the other hand, the case that an AS incorrectly sends an inter-region route instead of an intra-region route due to using Bloom Filters is impossible to occur in RMI because Bloom Filters demonstrate false positive but not false negative.

For example, in previous example, as shown in Fig. 7, AS 5 should send an inter-region route to AS 6. Suppose that there is a false positive from the provider list matching, AS 5 incorrectly sends an intra-region route to AS 6. After AS 6 receives the route, it can use guideline 7 to detect the fault. According to guideline 7, due to the peer-to-peer relationship between AS 5 and AS 6, AS 6 checks its routing table to see if it has another intra-region route via its customer. If there is no such a route, the received intra-region route violates guideline 7 so that AS 6 can detect the fault. Note that this method can avoid the false negative result from the origin AS's provider list matching. Even though the example only shows that the consistence check can detect faulty configuration due to false negative of provider list, we believe that the consistence check can also be used to detect other types of mis-configuration.

4.3 Invalid Path Attack Detection

The consistence check guidelines not only can be used to detect faulty configuration but also can be used to detect invalid path attacks. A malicious AS can modify the path it receives from other ASes by inserting or deleting ASes from the path. And then, the AS advertises the incorrect route, which does not represent the true AS path to a destination prefix, to convince other ASes to route traffic for the prefix to itself. In this paper, we focus on the invalid path attacks launched by customers to attack their providers. The reason is that providers do not have incentive to send invalid paths to their customers since their customer traffic will always traverse their networks.

Based on the location of attackers and the types of invalid paths, we classify the invalid path attacks into four categories as described below. We use several examples to illustrate how to use the above guidelines to detect each type of invalid path attacks. To simplify the description, for an attacker, we define its *remote provider* as an AS which only has one or several inter-region routes to reach a destination, and a a *local provider* as an AS which only has one or several intra-region routes to reach the destination. In the following discussion, we show how a remote provider or a local provider can detect invalid path attacks. Here, we assume that the remote provider and the local provider must have at least one valid path. The reason is that highly interconnected Internet AS topology makes it possible to let each AS receive a valid route. Furthermore, we assume that all the paths, including the invalid path, must have the correct provider list. Otherwise, according to guideline 1, the conflict between the correct routes and false routes indicates a possible attack.

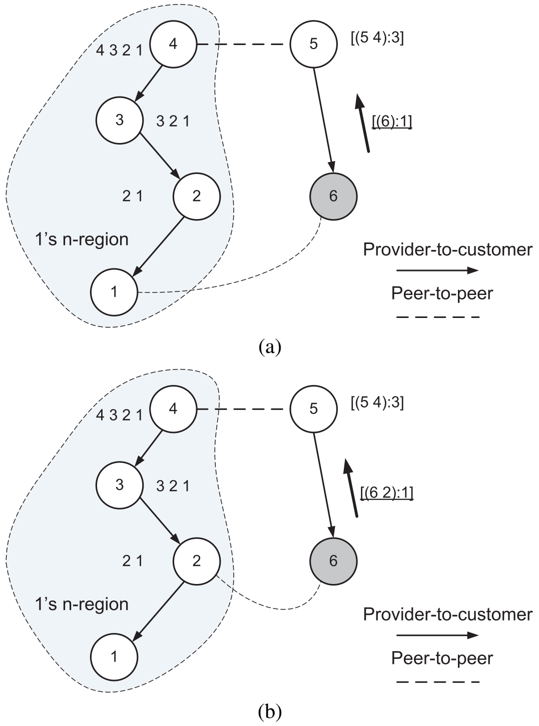

(a)

(b)

Fig. 8 Examples of invalid inter-region path to remote provider. In this example, AS 6 sends an invalid inter-region path to AS 5. The attacker is represented by a shadow node. The dashed curve is used to represent an invalid link, and the path with underline represents an invalid path.

Invalid Inter-region path to remote provider. We consider the case that the remote provider is in the attacker's provider region, and both of them are not in the origin AS's n-region. As shown in Fig. 8, the attacker can either pretend to originate the invalid inter-region path (Fig. 8 (a)), or forward the path originated by another AS (Fig. 8 (b)). In the two cases, the remote provider can detect the path attack by examining guideline 3 and 4. AS 5 can detect the invalid path because the invalid path [(6):1] in Fig. 8 (a) and [(6 2):1] in Fig. 8 (b) come from its customer, which violates guideline 3. In addition, AS 5 can use guideline 4 to detect the attack shown in Fig. 8 (a). That is, AS 5 can find that AS 6, who originates the inter-region path, actually is not in the origin AS's provider list.

Invalid Intra-region path to remote provider. Just like the previous case, the remote provider is in the attacker's provider region, and both of them are not in the origin AS's n-region. In this case, the attacker sends an invalid intra-region path instead of inter-region path to the remote provider. The remote provider can use guideline 5 to detect the attack. For example, as shown in Fig. 9, AS 6 sends an intra-region path (6 1) to AS 5. Upon receiving the path, AS 5 can use its provider list to verify if the origin AS is its provider or not. Since the origin AS is not a provider, which violates guideline 5, AS 5 can detect the attack.

Invalid inter-region path to local provider. Here, we consider the case that the attacker is in the origin AS's n-

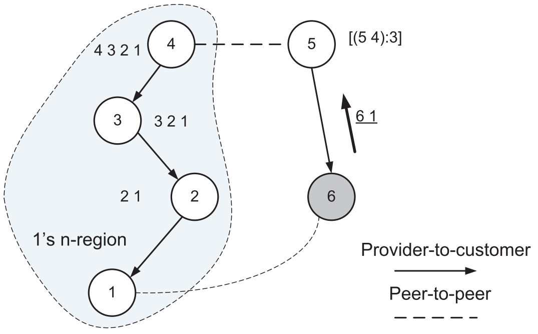

Fig. 9 An example of invalid intra-region path to remote provider. In this example, AS 6 sends an invalid intra-region path to AS 5.

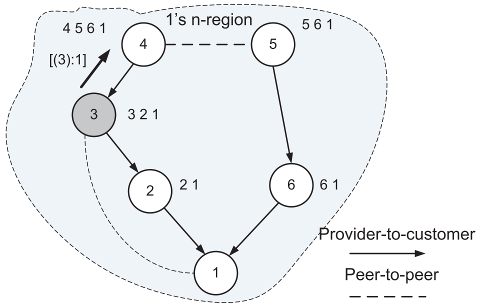

Fig. 10 An example of invalid inter-region path to local provider. In this example, AS 3 sends an invalid inter-region path to AS 4.

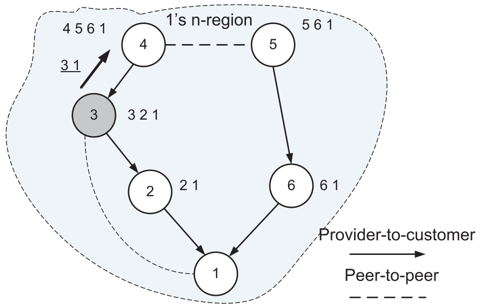

Fig. 11 An example of invalid intra-region path to local provider. In this example, AS 3 sends an invalid intra-region path to AS 4.

region. The attacker sends an invalid inter-region path to its provider. The provider can detect the attack according to guideline 4. That is, the inter-region path cannot come from its customer. For example, in Fig. 10, AS 3 sends an invalid inter-region path to AS 4. Since the route comes from its customer, which violates guideline 4, AS 4 can refuse to accept the invalid route.

Invalid intra-region path to local provider. In this case, the attacker is in the origin AS's n-region, and sends an invalid intra-region path to its provider. For example, in Fig. 11, AS 3 sends an invalid intra-region path (3 1) to AS 4. We can find that this type of invalid path attacks is similar to invalid path attacks in BGP. Since the provider lists only provide the membership of providers, it cannot be used to detect this type of attacks. The topology information, such as the links between ASes, is required to detect such attacks. Although provider lists cannot be used to detect this kind of attacks, RMI actually delimits and constrains the potential attackers–the attackers must be the origin AS's providers, and the type of invalid paths–the attackers must send intra-region paths. Otherwise, the invalid path will not pass the consistence check. Several possible methods can be used to detect this type of attacks. For example, each provider can use the information about its customer region to verify the path. The customer region could contain all possible paths from each customer to itself. Furthermore, with multipath routing, it is possible for each provider to obtain the complete path information for its customer region.

4.4 Invalid Origin Attack Detection

An invalid origin AS attack occurs when an AS attacker pretends to originate a prefix that it does not own. When the route to the bogus prefix propagates, some ASes will reroute to the hijacker instead of the legitimate host, making the prefix unreachable. As we described in Sect. 2, the routes to a destination may appear multiple origins due to static configuration or using private AS number. Thus, to detect such attacks, each AS should distinguish valid multiple origins and invalid origin attacks.

RMI utilizes the provider list to detect invalid origin attacks. Suppose that AS v has several paths to a destination that has different origin ASes. From those paths, AS v can find the associated provider lists of those multiple origins. Suppose that $P\_region(u_1), P\_region(u_2), \dots, P\_region(u_k)$ are the provider lists from ASes $u_1, u_2, \dots, u_k$, respectively. As we described in previous section, those ASes should advertise a complete provider list for the real origin AS. Thus, AS v can detect invalid origin attacks by verifying if those provider lists are the same (guideline 1), and if those origin ASes are located in the provider lists. If one AS is not in those lists, or the lists are not consistent, there is an invalid origin attack.

5. Evaluation

In this section, we focus on evaluating the scalability of RMI. Since RMI heavily depends on the propagation of provider region, a large provider list can result in a large update message, which can impact the scalability of RMI. To understand the size overhead, we measure the size distribution of provider region in the Internet. Through extensive experiments by simulations, we evaluate the performance of RMI. Our results show RMI's effectiveness in reducing the number of routing messages and convergence time.

5.1 Size Distribution of Provider Region

We use BGP updates from Oregon RouteView [2] to infer the provider region for each AS. In particular, we create Internet connectivity on AS level based on BGP updates and

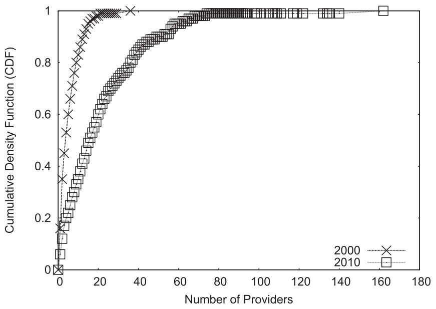

Fig. 12 The size distribution of provider regions in the Internet.

AS relationships [9]. We collect BGP updates at two different time, May 1, 2010 and May 1, 2000. The reason that we collect data from different time is to investigate the size of the Internet and the size of each AS's provider region as the Internet evolves over the past 10 years. Based on the AS connectivity, we derive the provider region for each AS.

Figure 12 shows the size distribution of each AS's provider region. We find that in 2000, about 60% of ASes have less than 6 providers in their provider regions, and 10% of ASes have more than 10 providers. The maximum number of providers in a provider region is 36. In 2010, only 25% of ASes have less than 6 providers and the majority of ASes (about 80%) have no more than 40 providers. The maximum number of providers in a provider region is 162. That means, as expected, the majority of ASes intend to connect more providers as the Internet evolves. However, the number of providers is still small compared with the size of customer region. We also measure the size distribution of customer region, which does not show in this paper. We find that the majority of ASes do not have customers (stub ASes), which is consistent with previous measurements. However, some ASes can have very large number of customers. For example, AS 2914 has 1,3854 direct and indirect customers in 2010. Our measurement also shows that as the Internet evolves the number of customers grows very fast.

As we described in previous section, RMI only advertises the information about provider region, and encodes the information into a Bloom Filter to control the size of routing updates. Because the majority of ASes have less than 40 providers, we can construct a Bloom Filter with a small number of bits, and ensure that the false positive rate requirement is met.

5.2 Simulation Results

We implement RMI based on the simBGP simulator, which is a message-level event driven simulator. We implement multipath advertisement, the Provider Discovery Protocol and route dissemination rules in SimBGP. The multipath mechanism we use is to advertise all available routes, called full multipath advertisement. Our simulation is based on the internet-like AS level topologies annotated with business relationships like peer-peer, customer-provider and provider-customer generated by [6].

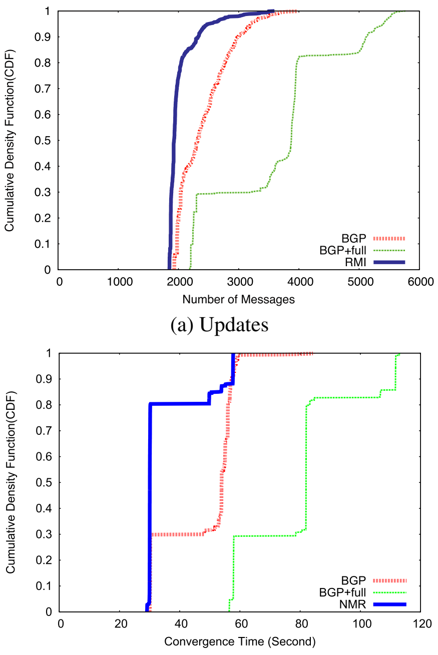

(a) Updates

(b) Convergence Time

Fig. 13 Experimental results for prefix announcement.

We evaluate the performance of RMI in terms of the message overhead and the convergence time during prefix announcement events and link failure events. The simulations are based on a topology with 1000 nodes. In the 1000-node topology, there are totally 818 stub ASes and 515 multi-homed stub ASes. In addition, we investigate the number of message generated by RMI as long as the size of network topology increases. When we present the simulation results, we compare the performance of RMI, BGP and BGP with full multipath advertisement, which is denoted as "BGP+full".

During Prefix Announcement Events. During a prefix announcement, we select one stub AS to announce a prefix. After the network converges to a stable state, we collect the results of the number of routing messages and measure the convergence time. Figure 13 (a) shows the CDFs of the overall number of routing messages each event. Note that we did not take account into the number of TIPs in the number of routing messages because TIP processing is a much lighter operation than routing message processing. In general, from the figure, we observe that RMI significantly outperforms BGP and BGP with multipath in terms of routing messages and convergence delay. More specifically, during the majority of prefix announcement events (about 70%), we observe that BGP+full method produces more than 3,000 updates. In the worst case, 5,777 updates are generated. On the contrary, RMI produces less than 2,000 updates during the majority of the events (around 74%). We also observe that BGP generates less than 2,000 updates only during 20% of the events. This implies that RMI even performs better than BGP despite the fact that RMI constructs more routes. The reason is that RMI limits not only the multipath prop-

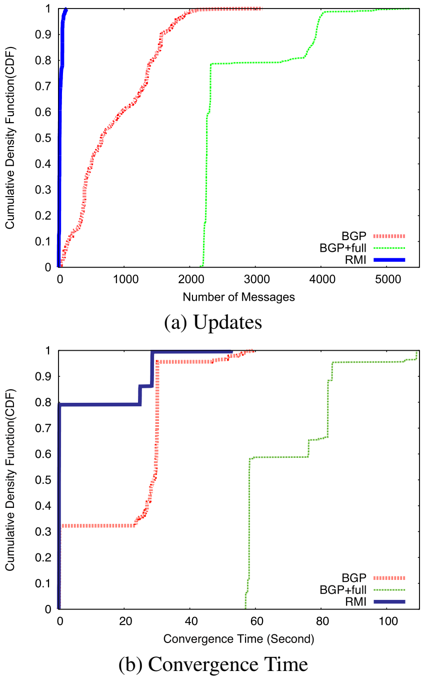

(a) Updates

(b) Convergence Time

Fig. 14 Experimental results for provider link failures of multi-homed ASes.

agation but also the path exploration during routing convergence. Furthermore, on average, RMI reduces the message overhead by a factor of 1.19 in comparison of BGP, and by a factor of 1.78 compared to BGP with full multipath advertisement.

We measure the convergence time by recording the duration from the time when a prefix is advertised to the one when the network becomes stable. Figure 13 (b) shows the CDFs of the convergence time in the event of prefix announcement. In general, we observed that RMI can converge much faster than BGP and BGP+full. Specifically, during the majority of prefix advertisement events (more than 80%), RMI has the convergence time less than 30 seconds. On the contrary, during the majority of prefix advertisement events (more than 70%), BGP and BGP with full multipath advertisement have more than 30 seconds and 57 seconds, respectively. On average, RMI reduces the convergence time by a factor 1.34 comparing to BGP, and a factor of 2.26 in comparison of BGP with full multipath advertisement.

During Link Failure Events. In each run, we choose a multi-homed AS to announce a prefix and wait for the network to become stable. And then, we disconnect one of the provider links to simulate a link failure event. Totally, we simulate 1198 link failure events. Figure 14 (a) shows the CDFs of the overall number of updates in each event. We observe the same result as that during prefix advertisement events. That is, RMI performs better than BGP and BGP+full in terms of message overhead and convergence delay. During all the events, RMI generates less than 130 updates. During 56% of the link failure events, RMI generates less than 50 updates. For BGP+full, due to full multipath announcement, the link failure will be propagated to the whole network. Thus, even in the best case, BGP+full generates 2,163 updates. On average, RMI generates only 28 updates comparing to 844 in BGP and 2,609 in BGP+full.

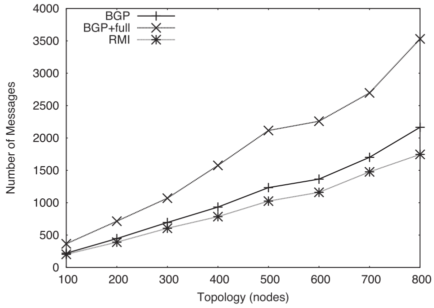

Fig. 15 The scaling of messages during prefix announcement event.

Figure 14 (b) shows the CDFs of convergence time during link failure events. We observe that during almost 79% of the events, RMI converges in less than a second. On average, RMI reduces the convergence time by a factor 3.5 compared with BGP, and a factor of 11.4 compared with BGP+full.

Moreover, we ran simulations on several different topologies, which range from 100 nodes to 800 nodes. The goal of this measurement is to provide an overview of the scalability of RMI with the increase of the network size. For each topology, we measure the number of routing updates during prefix announcement events. In Fig. 15, we show the average number of routing updates for each topology. The figure shows that the increase in the number of messages is linear with the increase of the network size. In summary, all of our simulation results show that RMI has a good scalability in terms of the number of routing messages and convergence delay.

6. Related Work

Several hybrid routing protocols are intended to improve the scalability of inter-domain routing. HLP [21], which is close to our approach, combines link-state routing protocol within a provider-customer hierarchy and path vector routing protocol across peering hierarchies. HLP addresses the reliability and scalability by relying on the link-state protocol. However, our measurement has shown that as the Internet evolves, the number of ASes in the provider-customer hierarchy grows so fast. Thus, distributing the link state information in a large scale communication network is impractical, which can impact the scalability. Furthermore, it has been shown that it is difficult to implement various routing policy based on link state routing protocol [27]. RMI differs from HLP in the sense that RMI is based on path vector routing. Even though both RMI and HLP divide the Internet into regions, the region in RMI is related to each AS. RMI requires each AS to keep and propagate its provider-customer relationships, i.e., its direct and indirect providers. On the contrary, the region in HLP includes all customers of tier-1 ASes, and each AS has to keep the whole region topology information. Thus, the region size in RMI is much smaller than that in HLP. Furthermore, because RMI is based on path vector, it can support most of BGP's routing policies. HAIR [7] presents a scalable routing architecture for future internet. It improves the routing scalability based on separation of locators and identifiers and a hybrid edge-based approach. However, RMI provides the scalability of inter-domain routing based on regional routing.

In addition, dividing inter-domain routing into intra-region routing and inter-region routing is not a new idea. DTIA [3] proposes to use regional routing to improve the scalability of inter-domain routing. Different from DTIA, RMI presents a new method to deal with the inter-region routing in which a route is composed of different portions. Also, RMI is prefix based routing while routing in DTIA is based on AS and region identifiers.

Many previous work focuses on the flexibility of multipath control that allows the end users to choose the paths [4], [8], [15], [20], [24]–[26], [31]. For example, MIRO [24] allows ASes to have more control over the flow of traffic in their networks, as well as enable quick reaction to path failures. MBGP [8] focuses on improving network bandwidth. However, MBGP is not an efficient solution for Internet-wide multipath routing because it uses message flooding to discovery the multiple paths. Path Splicing [18] takes advantage of alternate paths in BGP to discover multiple paths. However, Path Splicing might cause forwarding loops and violate routing policies. Pathlet Routing [11] enables a source to assemble an end-to-end route, and allows ASes to control the portion of routes that pass them. NIRA [25] allows end users to choose the sequence of Internet service providers a packet traverses, but it offers valley-free paths only. BANANAS [15] uses explicit AS-PATH forwarding technique to implement multipath routing. Some schemes [25], [31] utilize a link-state like routing to acquire the knowledge of the whole network for implementing source routing, which limits the scalability.

7. Conclusion

Despite the fact that extending multipath routing to inter-domain routing can improve the reliability of the Internet routing, designing a scalable and secure inter-domain multipath routing is challenging. This paper presents a scalable and secure inter-domain multipath routing protocol. We explore how regional path routing can provide multipath routing feature with acceptable message overhead and good convergence property. Based on the analysis and the simulations, RMI outperforms BGP and BGP with full multiple path advertisement. In addition, RMI can provide a certain level of route validation to improve the security of inter-domain routing. We believe that this paper presents the first in-depth study of multipath routing under constraint propagation. Our approach shows that controlling the visibility of multiple paths within a provider region allows us to kill two birds with one stone.

One of the main questions remaining to be studied is to use customer region to detect invalid path attacks occurring within the origin AS's region. We believe that using customer region will help us to detect such attacks. Another question is studying the computation overhead of regional path distribution and consistence check. Our future work is to implement a prototype of RMI based on XORP [1], and measure the computation overhead.

References

[1] The Extensible Open Router Platform (XORP). http://www.xorp.org/

[2] University of Oregon Route Views Project. http://www.routeviews.org/

[3] P. Amaral, F. Ganhão, C. Assunção, L. Bernardo, and P. Pinto, "Scalable multi-region routing at inter-domain level," Proceedings of the 28th IEEE conference on Global telecommunications, pp.641–648, IEEE Press, 2009.

[4] K. Argyraki and D.R. Cheriton, Loose source routing as a mechanism for traffic policies, Proc. Future Directions in Network Architecture, 2004.

[5] K.R.B. Butler, P. McDaniel, and W. Aiello, Optimizing bgp security by exploiting path stability, ACM Conference on Computer and Communications Security, pp.298–310, 2006.

[6] X.A. Dimitropoulos and G.F. Riley, Modeling autonomous-system relationships, PADS, pp.143–149, 2006.

[7] A. Feldmann, L. Cittadini, W. Mühlbauer, R. Bush, and O. Maennel, Hair: Hierarchical architecture for internet routing, Proceedings of the 2009 workshop on Re-architecting the internet, pp.43–48, 2009.

[8] H. Fujinoki, Multi-path BGP (MBGP): A solution for improving network bandwidth utilization and defense against link failures in inter-domain routing, Proceedings of ICON, 2009.

[9] L. Gao, On Inferring Autonomous System Relationships in the Internet, IEEE/ACM Trans. Networking, vol.9, no.6, Dec. 2001.

[10] L. Gao and J. Rexford, A Stable Internet Routing without Global Coordination, IEEE/ACM Trans. Netw., vol.9, no.6, pp.681–692, Dec. 2001.

[11] P.B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," SIGCOMM 09, pp.111–122, 2009.

[12] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working around bgp: An incremental approach to improving security and accuracy of interdomain routing," Proc. NDSS, 2003.

[13] Y.-C. Hu, A. Perrig, and M.A. Sirbu, "Spv: secure path vector routing for securing bgp," SIGCOMM, pp.179–192, 2004.

[14] J. Karlin, Pretty good bgp: Improving bgp by cautiously adopting routes, Proc. International Conference on Network Protocols, 2006.

[15] H.T. Kaur, S. Kalyanaraman, A. Weiss, S. Kanwar, and A. Gandhi, "BANANAS: an evolutionary framework for explicit and multipath routing in the internet," Proc. ACM SIGCOMM workshop on Future directions in network architecture, 2003.

[16] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, Secure border gateway protocol (s-bgp, IEEE J. Selected Areas in Communications, vol.18, no.103–116, 2000.

[17] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A prefix hijack alert system," Proc. USENIX Security Symposium, 2006.

[18] M. Motiwala, N. Feamster, and S. Vempala, "Path Splicing," Proceedings of ACM SIGCOMM, SEATTLE, WA, AUGUST 2008.

[19] J. Qiu, L. Gao, S. Ranjan, and A. Nucci, "Detecting bogus bgp route information: Going beyond prefix hijacking," SecureComm 2007, pp.381–390, 2007.

[20] B. Raghavan and A.C. Snoeren, "A system for authenticated policy-

compliant routing," Proc. ACM SIGCOMM, 2004.

[21] L. Subramanian, M. Caesar, C. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica, "HLP: A next generation inter-domain routing protocol," Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications, pp.13–24, ACM, 2005.

[22] F. Wang and L. Gao, "Path diversity aware interdomain routing," INFOCOM, 2009.

[23] R. White, "Securing BGP through secure origin BGP (soBGP)," BUSINESS COMMUNICATIONS REVIEW, vol.33, no.5, pp.47–53, 2003.

[24] W. Xu and J. Rexford, "Miro: multi-path interdomain routing," SIGCOMM '06, pp.171–182, 2006.

[25] X. Yang, "NIRA: a new Internet routing architecture," Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture, 2003.

[26] X. Yang and D. Wetherall, "Source selectable path diversity via routing deflections," ACM SIGCOMM Computer Communication Review, 2006.

[27] X. Zhang, A. Perrig, and H. Zhang, "Centaur: A hybrid approach for reliable policy-based routing," ICDCS '09: Proceedings of the 2009 29th IEEE International Conference on Distributed Computing Systems, pp.76–84, 2009.

[28] Z. Zhang, Y. Zhang, Y.C. Hu, Z.M. Mao, and R. Bush, "Ispy: detecting ip prefix hijacking on my own," SIGCOMM, pp.327–338, 2008.

[29] M. Zhao, S.W. Smith, and D.M. Nicol, "Aggregated path authentication for efficient bgp security," ACM Conferernce on Computer and Communication Security (CCS), pp.128–138, 2005.

[30] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S.F. Wu, and L. Zhang, "An analysis of bgp multiple origin as (moas) conflicts," IMW '01: Proc. 1st ACM SIGCOMM Workshop on Internet Measurement, pp.31–35, 2001.

[31] D. Zhu, M. Gritter, and D. Cheriton, "Feedback based routing," Proc. SIGCOMM Workshop on Hot Topics in Networking, 2002.

Bin Dai was born in 1982. He is a fourth year Ph.D. candidate with the School of Computer at National University of Defense Technology. He received his B.E. and M.S. degrees in the School of Computer at National University of Defense Technology. Between Feb, 2008 to Feb, 2009, he was working with prof. Scott Shenker at International Computer Science Institute. His current research focuses on interdomain multipath routing.

Feng Wang was born in 1972. He is an assistant professor with the School of Engineering and Computational Sciences at Liberty University. He received his Ph.D. degree in Electrical and Computer Engineering at the University of Massachusetts, Amherst. He received his B.E. degree from Zhejiang University in China, and M.S. degree from Yanshan University in China. His research interests include networked computer systems, Internet routing, and wireless networks.

Baokang Zhao was born in 1981. Currently, he is with the School of Computer Science, National University of Defense Technology. His current research interests include security and privacy in sensor networks, algorithms, protocols and system design in delay tolerant networks and Internet. Baokang Zhao served as reviewers for several journals, including Computer Communications (Elsevier), Security and Communication networks (Wiley), Journal of Computer Science and Technology (Springer), etc.

Jinshu Su was born in 1962. He is a professor with the School of Computer at National University of Defense Technology. He received his B.E., M.S. and Ph.D. degrees from the School of Computer at National University of Defense Technology. His research interests include Internet arcitecture, Internet routing, security, and wireless networks.