Bootstrapping Evolvability for Inter-Domain Routing

Raja R. Sambasivan*, David Tran-Lam†, Aditya Akella†, Peter Steenkiste*
*Carnegie Mellon University, †University of Wisconsin-Madison

ABSTRACT

It is extremely difficult to deploy new inter-domain routing protocols in today's Internet. As a result, the Internet's baseline protocol for connectivity, BGP, has remained largely unchanged, despite known significant flaws. The difficulty of deploying new protocols has also depressed opportunities for (currently commoditized) transit providers to provide value-added routing services. To help, we identify the key deployment models under which new protocols are introduced and the requirements each poses for enabling their usage goals. Based on these requirements, we argue for two modifications to BGP that will greatly improve support for new routing protocols.

Categories and Subject Descriptors

C.2.2 [**Network Protocols**]: Routing protocols

General Terms

Design

Keywords

BGP, evolvability, inter-domain routing

1. INTRODUCTION

BGP, the Internet's inter-domain routing protocol, is the critical glue that holds the Internet together. All services and content we hold dear are accessible because of the routing paths that it computes. But, this critical protocol is plagued with severe problems. For example, it does not provide domains (stubs or transit providers) sufficient influence to limit incoming traffic; its paths are slow to converge and prone to oscillations; it indiscriminately chooses a single best-effort path per router, robbing other domains of paths they may prefer more; and it is prone to numerous attacks, including prefix hijacking, traffic interception, and black-holing.

In response, researchers and operators have proposed a variety of critical fixes and improvements. Changes that only involve single domains (e.g., new forms of outbound route filtering and multi-protocol BGP to connect customer sites [1]) have been deployed quickly. However, broader changes that span multiple domains have proven more difficult to roll out (e.g., adding secure route announcements via S-BGP [11] or adding awareness of path costs to limit incoming traffic [15]). The research community has also explored even more disruptive protocols [19, 25, 27]. However, none have been deployed despite the clear benefits they offer.

We posit that the reason even critical fixes are difficult to deploy is because BGP cannot *bootstrap evolution*—i.e., help new protocols gain traction and seamlessly deprecate itself in favor of a replacement. Evolvability support is critical in order to rapidly upgrade a protocol—either across all or a subset of domains—whenever new use cases bring critical deficiencies to the fore. In the extreme, it can help the Internet transition from an old routing protocol to one that uses a fundamentally different paradigm (e.g., move from destination-based to path-based forwarding). Such evolution support could also facilitate the simultaneous co-existence of multiple disparate protocols, improving the richness of the Internet architecture as a whole.

In this paper, we ask: *given the benefit of hindsight, how would we redesign a BGP-like inter-domain routing protocol with support for bootstrapping evolvability?* In answering this question, our paper makes two key contributions.

First, we provide a systematic analysis of the space of deployment models for introducing new protocols. We identify three models: rolling out protocol fixes or new features; rolling out custom routing protocols, which are used for only select traffic; and, replacing routing protocols entirely.

For each model, we provide examples from prior research, allowing us to precisely enumerate the scope of architectural (control and data plane) enhancements entailed by the model and the requirements they impose for routing evolvability. Our requirements align with 4D's principles of providing clean abstractions for dissemination, discovery, and decision [6].

Second, we describe two modifications to BGP—integrated advertisements and pass-through modules—that we claim satisfy the requirements. They bootstrap protocol evolution by allowing multiple protocols' control information to be compactly carried in BGP-like advertisements. We provide concrete examples that show how these modifications can help a BGP-like inter-domain routing protocol seamlessly evolve into some recently proposed BGP enhancements/alternatives.

2. DEPLOYMENT MODELS

Based on a literature survey [3, 5, 11, 13, 15, 16, 19, 24, 25, 27], we have identified three commonly used deployment models for introducing new protocols. They differ in how they expect new protocols to be used. As such, it is the model, not individual protocols, that dictate requirements for routing

HotNets '15 November 16–17 2015, Philadelphia, PA USA

DOI: http://dx.doi.org/10.1145/2834050.2834101

evolvability. Multiple deployment models may be suitable for a single protocol. In such cases, operators choose a model based on protocol-specific goals (e.g., do they want the protocol to eventually replace the baseline or only be used for select traffic?). This section describes our models in detail. We first discuss the data-plane issues that can arise when deploying multiple protocols, which our models manage differently.

Assumptions: At the beginning of time, we assume that all domains, ASes for short, are using a baseline routing protocol for inter-connectivity that is BGP-like. It is a path vector protocol in which advertisements carry connectivity information upstream from traffic sinks to traffic sources. Data packets flow downstream from sources to sinks. Advertisements identify only one path to each sink. The discussion below and the mechanisms presented in Section 3 are agnostic to whether ASes use distributed control (i.e., routers choose paths) or centralized control (e.g., SDNs [8, 10]) and to whether ASes support different sets of protocols on different routers.

Terminology: *Islands* refer to a cluster of one or more contiguous ASes that support the same set of routing protocols. Neighbors of islands run a different set of protocols. *Baseline ASes / Islands* refer to those that run the baseline protocol (e.g., BGP). *Upgraded ASes / Islands* refer to those that support the new protocol being discussed. We refer to the set of baseline ASes separating two upgraded islands as *gulfs*.

2.1 Routing protocols & the data plane

The data plane or *network protocol* is responsible for enforcing routing protocols' path choices. When multiple routing protocols are deployed concurrently, consistency of routing decisions becomes an issue. If care is not taken to consistently enforce the same routing protocol's path choices for a destination address at every location (e.g., router or AS), the resulting end-to-end path may not be the result of any single protocol's choices. Also, paths chosen by one protocol at one location may prevent data packets from using better paths selected by a more preferred protocol at other locations. These issues can severely curtail new protocols' benefits. Whether or not a protocol needs its routing decisions to be consistently enforced informs the model to which it is best suited.

Enforcing consistency requires different mechanisms within islands and across islands. Our discussions assume an IP prefix as the destination address, but are equally valid for other types (e.g., content names [26]). To ensure consistency within islands, protocols must be careful not to install conflicting entries at different points in the path. This requires assigning different protocols different addresses that name the same physical destinations.

Additionally ensuring consistency for routing decisions across islands requires the relevant protocol's path choices to be enforced at locations that do not support it (i.e., within gulfs). Doing so requires data packets to be encapsulated and tunneled, thus hiding their within-island addresses from other protocols and islands.

We note that if routing protocols use different network protocols or use a network protocol that supports multiple address types (e.g., XIA [7]), consistency issues cannot arise.

2.2 Model A: Updating the baseline

This model assumes that new protocols do not need consistency for their routing decisions. It is safe for end-to-end routing paths to be an amalgamation of different protocols' individual path choices. It is most useful for deploying critical fixes or updates to the existing baseline protocol. Such updates disseminate extra control information to improve path selection or the protocol itself. Many proposed fixes to BGP are suited for this model, including Wiser [15], for fixing BGP's broken support for traffic management, S-BGP [11] for fixing BGP's susceptibility to route hijacking, and LISP [3] for supporting mobility.

Data-plane issues: Since consistency of routing decisions is not an issue, there are no data-plane issues. Protocols deployed using this model can be leveraged to support new network protocols, similar to IPv6 support using M-BGP [1]. Content-based routing [26], which forwards traffic based on content names, can be enabled in a similar way.

Example: Figure 1 shows a scenario in which ASes start to incrementally deploy Wiser [15] as an update to BGP. Wiser fixes BGP's broken support for inter-domain traffic management by modifying BGP's advertisements to include a *global path cost*, which is used to inform path selection. This field is unit-less and normalized across neighbors. The two ASes at the edge of the large Wiser island, E1 and E2, use BGP to advertise paths to their neighbors in the BGP gulf. Lines show paths advertised and arrows show the direction of the advertisement.

The figure illustrates two problems. First, the source, which supports Wiser, must use BGP to select paths because it cannot see global path costs. As such, it will choose the shortest path (due to BGP's decision criteria), which has the highest global path cost. Second, E1 and E2 are at a disadvantage because they must honor global path costs when selecting paths, but cannot express their own costs. They are at the mercy of upstream ASes' routing decisions. This may dis-incentivize them from

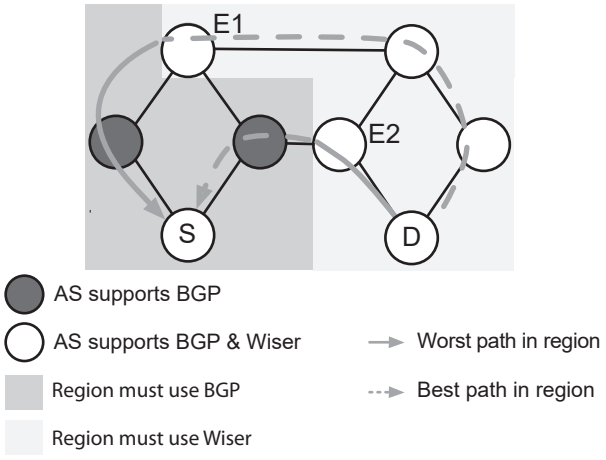

Figure 1: S cannot see path costs, so it will choose the highest-cost one.

supporting Wiser, especially if this requirement increases their payments to providers or peers.

Requirements: As the above example shows, today, non-contiguous islands or ASes that deploy updated baseline protocols cannot quickly leverage the improvements afforded by them. This is because updated baselines' extra control information cannot be disseminated across BGP gulfs. Thus, we end up with this requirement:

UB-R1 *Disseminate updated baseline's additional control information across gulfs.*

Also, the update must replace the baseline eventually:

UB-R2 *Allow the existing baseline to be eventually replaced.*

Constraints: This deployment model is only useful for a very restricted set of protocol improvements. For example (assuming a BGP baseline), it is limited to path-vector-based protocols. It assumes routing decisions need not be consistent and does not support off-path discovery of upgraded ASes.

2.3 Model B: Custom routing

This model assumes that protocols deployed using it require their routing decisions to be consistently enforced across the Internet. It also assumes that that new protocols will be used for only select traffic and that the baseline will be used for the rest. New protocols use out-of-band coordination to disseminate control information across upgraded islands (i.e., they use paths already established by the baseline).

In the literature, this deployment model is often used to introduce protocols that provide value-added services, which are sold for profit. Examples include selling alternate paths [13, 16, 24] and selling extra functionality on existing paths [16] (e.g., higher intra-domain or intra-island QoS). This deployment model can also be used to connect non-contiguous islands running a wide variety of protocols, including those that use different routing paradigms than the existing baseline (e.g., pathlet routing [5] or path-based routing [25,27]). For example, two non-contiguous islands could use a path-based protocol deployed using this model to explicitly coordinate the intra-island hops they will use for important traffic.

Data-plane issues: Islands will run multiple inter-domain routing protocols concurrently (e.g., the baseline and the new protocol). The new protocol's routing decisions must be enforced consistently, both within islands, and across gulfs. Assuming all routing protocols use the same network protocol and address types, separate address ranges must be assigned to custom protocols within islands. Packets must be encapsulated and tunneled across gulfs. Otherwise, the baseline protocol may divert packets from ever reaching an upgraded island.

Example: Figure 2 describes a scenario in which a transit AS (marked T) wishes to avoid the single poorly performing path advertised by BGP (the dashed path). An AS that supports MIRO [24] offers alternate paths for payment (the rightmost one). However, the transit AS cannot discover the MIRO-enabled AS because BGP does not allow discovery of

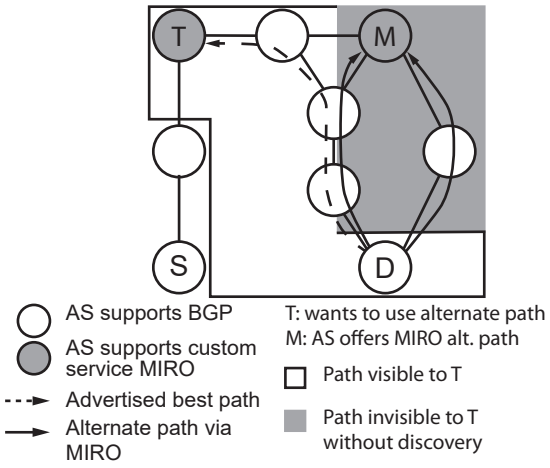

Figure 2: T cannot discover M's alternate path.

ASes' custom services or the extra coordination required to use them. This lack of discovery mechanism limits the MIRO AS's potential customers, perhaps only to its direct neighbors. It could use bespoke approaches for discovery (e.g., a web site), but these may go unnoticed.

Requirements: As the above example shows, ASes or islands supporting the new protocol must be able to both discover each other and how to coordinate out-of-band in order to exchange relevant control information, including protocol-specific information (e.g., alternate paths) and the type of encapsulation method that will be used to route packets across gulfs. Thus, we require:

CR-R3 *Facilitate discovery of custom services.*

Constraints: This model cannot be used for protocols that aim to replace the existing baseline. It will be difficult to deploy a large number of custom-routing protocols that use the same network protocol because each will have to be assigned increasingly smaller pools of addresses. In contrast, this model is attractive for routing protocols that use different network protocols (or different address types within an existing protocol), or for within-island protocol extensions.

2.4 Model C: Exclusive routing

This model involves deploying new routing protocols by completely *replacing* the baseline protocol with a new one in upgraded islands. So, the key difference between this and the previous model is that the new protocol is used for *all traffic* in these islands. Doing so is very aggressive model and likely to be only attractive if there are strong incentives or requirements that are impossible to meet with the baseline (e.g., high QoE for all traffic or specific economic relationships). As such, it is likely to be useful only within islands. However, multiple islands could use the same protocol and route traffic among each other using this model. In such cases, this model could be used to introduce radically different protocols that aim to eventually replace the existing baseline.

Protocols that could be introduced using this model include ones that use very different routing paradigms than

the baseline, such as HLP [19], which is a hybrid path-vector-based/link-state protocol, or path-based ones [5, 25, 27].

Data-plane issues: Within islands, consistency is not an issue because only a single routing protocol is used. The baseline protocol will not have alternate end-to-end paths to destinations controlled by upgraded islands, so consistency is also a non-issue when traversing gulfs. However, packets that traverse gulfs still need to be encapsulated so that they can be forwarded by both the new protocol and the baseline protocol (see requirements below).

Example: Figure 3 illustrates a scenario in which BGP is being replaced by SCION [27], a path-based protocol. Sources can be advertised multiple path options (exposed at the granularity of border routers). They pick which one they want to use by encoding the path in packet headers, which routers key on to forward traffic. In this case, the rightmost SCION region in the diagram exposes two paths to the destination.

The scenario illustrates two key problems: the SCION source in the diagram cannot discover other SCION islands or route traffic to them. Unlike the previous model, out-of-band coordination and tunneling cannot be used to address this problem as it will not scale to handle all traffic. Also, ASes within BGP gulfs cannot route to destinations within SCION islands. Both problems can be addressed by re-distributing SCION routes into BGP [14]. But, BGP can only advertise one path per router, so one of the SCION paths would be lost.

Requirements: Solving the above problems requires the ability to disseminate new protocols' control information in-band with the baseline protocol. Doing so sidesteps scalability issues and avoids redistribution issues that may result in loss of important information. Thus, we have:

ER-R4 *Enable in-band dissemination of new protocols' control information*

For protocols that aim to become the new baseline protocol, UB-R2 also applies.

With in-band dissemination, paths are jointly controlled by the baseline in gulfs and by the new protocol in upgraded islands. Packets routed along these paths be encapsulated so that they can be forwarded by both protocols. For example, in the example above, to route packets to the SCION destination, the SCION source must encapsulate data packets so that they contain both an IP header and their path choice within the rightmost SCION island.

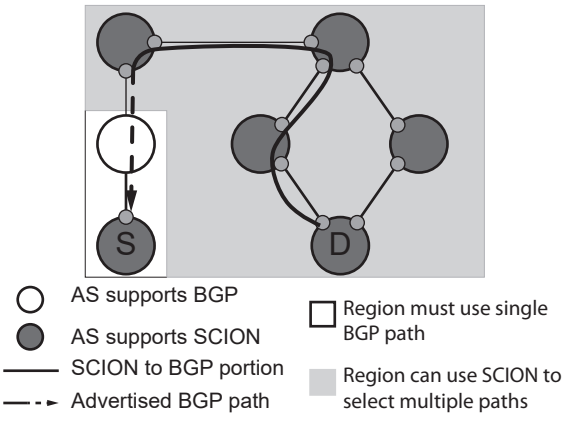

Figure 3: S cannot be advertised both paths to D.

Constraints: There are no limitation on the type of protocols that can be deployed using this model.

3. BOOTSTRAPPING EVOLVABILITY

In this section, we argue that two modifications to BGP—integrated advertisements and pass-through modules—satisfy the requirements derived in the previous section and would allow BGP to bootstrap evolution. We first describe our mechanisms, focusing on how they enable evolvability for updating the baseline, then discuss how they could be applied to enable evolvability for custom and exclusive routing.

3.1 Integrated advertisements

Integrated advertisements (IAs) transform BGP's advertisements into containers that can compactly carry multiple protocols in addition to the current baseline. This allows updated routers to use new protocols and legacy ones to fall back on the baseline. As more routers support the updated baseline, the current one can be eventually replaced *(UB-R2)*. Like BGP, each IA is associated with a destination (e.g., a prefix). However, different protocols encoded in an IA can name the same destination differently (e.g., using different address types). To combat potentially large message sizes, control information that is the same across protocols is shared initially and split when modified by upstream ASes.

Figure 4 shows the basic structure of an IA, which we believe is expressive (i.e., allows a wide range of protocols to be encoded using it) and maximizes potential for information sharing. It is composed of three elements. First, paths, which are encoded as nodes and edges. We allow for multiple paths to allow protocols that expose more than one path per router to express them (e.g., SCION [27]). Nodes determine path granularity. For example, they could be ASes, as in BGP, or border routers, as in SCION. Edges specify links between nodes.

The second element includes path descriptors, which describe properties of entire paths or parts of them (i.e., those of specific nodes or edges). Destination addresses (e.g., prefixes or content names) are included as path-level descriptors. Possible node-level descriptors might include intra-domain QoS objectives or S-BGP's route attestations [11]. Edge descriptors could include intra-domain congestion levels or BGP's MEDs. For accountability, we require node descriptors to include a field that states the AS that created the corresponding node.

The third element includes *AS descriptors*, which allows ASes to include important information about themselves (e.g., on-path or off-path services offered). They are also used for loop detection across all of the protocols included in an IA.

How IAs can be used to encode updated baseline protocols: Updated baselines' and current baselines' end-to-end paths can be an amalgamation of each others' routing decisions

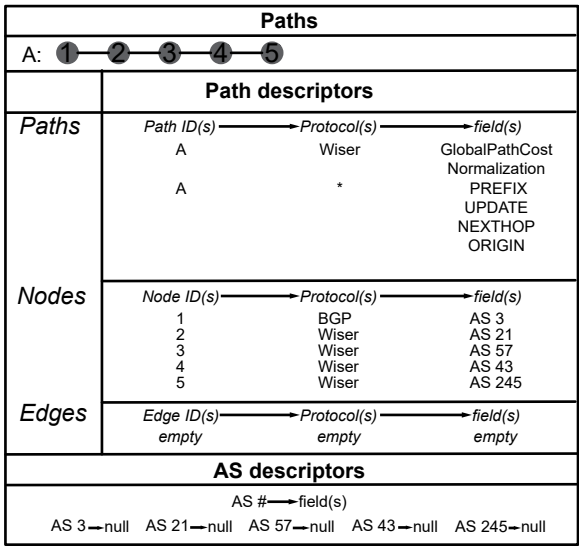

Figure 4: Example of an integrated advertisement. This advertisement is received by the source AS in the BGP [18] to Wiser [15] example (Section 2.2). Asterisks indicates information that is shared across protocols.

(necessarily meaning they will use the same node granularity). As such, routers can populate an IA with a single path for both the current baseline and any updated ones they support. Many fields can be shared. Protocols will use the same address type, so the destination address can be shared across protocols as a path-level descriptor.

Assuming both our modifications are implemented, Figure 4 shows the IA that would be received by the source AS in the example from Section 2.2. It is the advertisement for the lowest-cost Wiser path. It includes two Wiser-specific fields. The first is the global path cost. The second is a normalization factor, reflecting the total cost of all paths disseminated by the AS that created this advertisement. The latter is required to allow upstream Wiser-enabled routers to normalize their intra-AS costs with the global one before adding their contribution.

3.2 Pass-through modules

Pass-through modules on routers work in concert with IAs. They pass through control information for unsupported protocols with new IAs for paths chosen by supported ones. This allows control information for updated baselines to be disseminated across gulfs that support only the baseline (UB-R1).

Figure 5 shows a router that includes a pass-through module. It is similar to existing routers, except it runs multiple *decision modules* for each protocol it supports. Decision modules include protocol-specific path-selection algorithms (e.g., BGP's tie-breaking logic), import/export filters, and data structures (e.g., ADJ-RIBs). The pass-through module assumes responsibility for receiving IAs, interfacing with decision modules, installing forwarding entries corresponding to their path choices, and disseminating new IAs.

To work, pass-through modules include the following elements and interfaces. Import/export filters allow implementation of global policies (e.g., prefer paths learned through customers). The *control interface* is used to push control information relevant to supported protocols to their decision modules. The *path-selection interface* allows protocols' decision modules to return their path choices and any modified control information. The *data-plane interface* allows pass-through modules to install forwarding choices corresponding to chosen paths. For protocols that share the same network protocol and address types (e.g., updated baselines), pass-through modules will only use the most recent supported version.

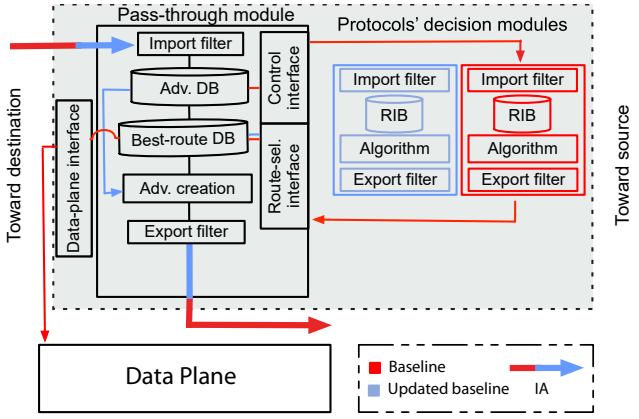

Figure 5: A router's pass-through module.

Pass-through modules store received advertisements in a database. When creating an IA for a chosen path, they index into this database to identify the message that advertised the path, and embellish the message as needed with new control information. They also add the protocol used to choose path(s) to the relevant AS descriptor, allowing upstream ASes to avoid paths chosen by undesired protocols.

How pass-throughs help with updating the baseline protocol: Passing through control information across gulfs allows non-contiguous islands to use an updated baseline when routing to each other. Figure 6 illustrates the result if the ASes in the scenario from Section 2.2 supported pass-throughs and IAs. The source AS is able to see Wiser's path cost (see Figure 4) and use it to select the lower cost, longer path. E1 and E2 are

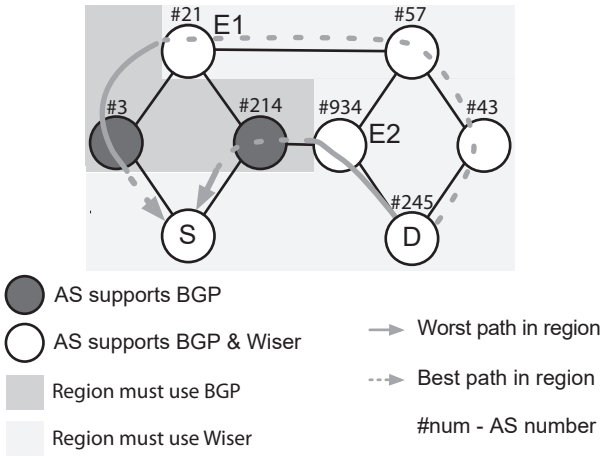

Figure 6: S sees path costs in IAs, so it chooses the lowest-cost one.

still at the mercy of ASes that run only BGP, but their situation incrementally improves as additional ASes deploy Wiser.

3.3 Custom & exclusive routing

IAs and pass-through modules enable discovery for custom routing and in-band dissemination for exclusive routing. For the former, IAs could carry descriptions of the custom services offered and the extra coordination needed to enable them within AS or node descriptors *(CR-R3)*. For the latter, routers could create IAs that encode control information for both the new protocol and the baseline protocol. It could also specify the data-plane encapsulation technique that must be used to forward traffic across gulfs *(ER-R4)*.

Our modifications allow the transit AS in the example from Section 2.3 to discover and use the MIRO AS's alternate path as follows. First, the MIRO AS uses IAs to advertise a path to a service portal it provides. The AS descriptor includes a description of the custom coordination required to use this service portal (e.g., a specific protocol). Second, the transit AS contacts the service portal to negotiate the alternate path to the destination and the data-plane encapsulation technique (e.g., an additional prefix) that will be used to cross gulfs and selectively route the transit's traffic. Third, the transit uses the necessary encapsulation technique to tunnel its traffic destined for the destination AS. As an optimization, the initial advertisement could include a list of the most popular alternate paths the MIRO-enabled AS provides (e.g., alternate paths to Google, or S-BGP [11] paths that avoid North Korea).

Figure 7 illustrates how IAs and pass-through modules address the example discussed in Section 2.4. The edge AS's border router in the SCION island creates an IA that includes control information for both BGP and SCION. The former includes a single path, an IP prefix, and an AS descriptor for the edge AS. The latter includes two SCION paths, AS descriptors for them, and an annotation in the edge AS's descriptor listing the encapsulation technique needed to bridge gulfs. In this case, it specifies that SCION packets should be encapsulated with an IP header that lists the prefix used in the IA as the destination. When receiving packets, the router at the edge of the SCION region de-encapsulates the IP header and forwards packets using the source's path choice, specified in the underlying SCION header.

3.4 Limitations

Our mechanisms are subject to the limitations and policies of the baseline protocol(s) used to bridge gulfs. For example, they allow Wiser islands to pick the lowest-cost path from the options given, but those options may include only high-cost paths because of ASes in BGP gulfs' poor path choices. Our mechanisms are also not sufficient to enable evolvability for protocols that are not path-vector-based (e.g., link-state).

4. OPEN QUESTIONS

As described in the paper, IAs are not lossily aggregated [18]. Doing so is important to reduce the total size of control messages sent by individual protocols (i.e., across all advertisements). But, because protocols will have differing aggregation policies for information that was previously shared, aggregation may result in larger IAs. To help, we are exploring mechanisms that allow protocols to cooperate during aggregation. We are also exploring how to accommodate protocols that differ in the rate they send advertisements [9] and whether our modifications will increase transient oscillations [21].

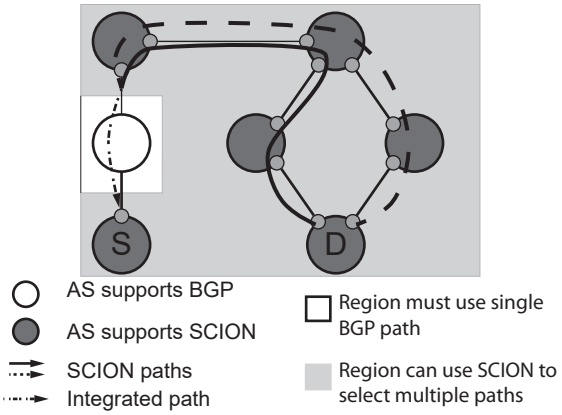

Figure 7: S sees both paths in the integrated advertisement.

5. RELATED WORK

Several previous research efforts focus on data-plane evolvability [7, 20, 22, 23, 26]. Our research complements these efforts by focusing on the control plane. Previous efforts have also identified requirements for network evolvability [2, 4, 17]. Those listed by Ratnasamy et al. [17] are compatible with our requirements, but we extend them to inter-domain routing.

Two features of BGP advertisements are similar in spirit to our mechanisms, but more limited in scope. Multi-protocol extensions to BGP allow advertisements to carry multiple network-protocol addresses (e.g., IPv4 and IPv6) [1]. Transitive community attributes are key-value pairs that *should* be always passed through. They can be used as building blocks for our IAs, but are not sufficient to enable evolvability on their own (e.g., they do not support information sharing).

Koponen et al. [12] propose using pathlet routing [5] to enable evolvability—i.e., as the new baseline—because of its ability to emulate many routing protocols. Our work can help such improved protocols gain traction on the Internet.

6. CONCLUSION

BGP cannot easily be evolved. This prevents new protocols from being widely deployed. Based on requirements identified by an analysis of key deployment models, we find that two modifications to BGP—IAs and pass-through modules–are promising starting points for making BGP evolvable.

Acknowledgements: We thank the reviewers, Mark Coatsworth, Aaron Gember-Jacobson, Michelle Mazurek, Ilari Shafer, and Brent Stephens for their feedback. This research was funded in part by NSF under award number CNS-1345305.

References

[1] T. Bates, R. Chandra, D. Katz, and Y. Rekhter. Multiprotocol Extensions for BGP-4. RFC 2283, feb 1998. http://www.rfc-editor.org/rfc/rfc2283.txt.

[2] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. *IEEE/ACM Transactions on Networking*, 13(3):462–475, June 2005.

[3] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis. The Locator/ID Separation Protocol (LISP). RFC 6830, January 2013. http://www.rfc-editor.org/rfc/rfc6830.txt.

[4] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox. Intelligent Design Enables Architectural Evolution. In *Proc. HotNets*, Nov. 2011.

[5] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica. Pathlet Routing. In *Proc. SIGCOMM*, Aug. 2009.

[6] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang. A Clean Slate 4D Approach to Network Control and Management. *SIGCOMM Computer Communication Review*, 35(5):41–54, Oct. 2005.

[7] D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. G. Andersen, J. W. Byers, S. Seshan, and P. Steenkiste. XIA: Efficient Support for Evolvable Internetworking. In *Proc. NSDI*, Apr. 2012.

[8] C.-Y. Hong, S. Kandula, R. Mahaan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving High Utilization with Software-Driven WAN. In *Proc. SIGCOMM*, Aug. 2013.

[9] G. Huston, M. Rossi, and G. Armitage. A Technique for Reducing BGP Update Announcements through Path Exploration Damping. *IEEE Journal on Selected Areas in Communications*, 28(8):1271–1286, Oct. 2010.

[10] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. stuart, and A. Vahdat. B4: Experience with a Globally-Deployed Software Defined WAN. In *Proc. SIGCOMM*, Aug. 2013.

[11] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, Sept. 2000.

[12] T. Koponen, S. Shenker, H. Balakrishnan, N. Feamster, I. Ganichev, A. Ghodsi, P. B. Godfrey, N. McKeown, G. Parulkar, B. Raghavan, J. Rexford, S. Arianfar, and D. Kuptsov. Architecting for Innovation. *SIGCOMM Computer Communication Review*, 41(3):24–36, July 2011.

[13] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs. R-BGP: Staying Connected in a Connected World. In *Proc. NSDI*, Apr. 2007.

[14] F. Le, G. G. Xie, and H. Zhang. Understanding Route Redistribution. In *Proc. ICNP*, Oct. 2007.

[15] R. Mahajan, D. Wetherall, and T. Anderson. Mutually Controlled Routing with Independent ISPs. In *Proc. NSDI*, Apr. 2007.

[16] S. Peter, U. Javed, Q. Zhang, D. Woos, T. Anderson, and A. Krishnamurthy. One Tunnel is (Often) Enough. In *Proc. SIGCOMM*, Aug. 2014.

[17] S. Ratnasamy, S. Shenker, and S. McCanne. Towards an Evolvable Internet Architecture. *SIGCOMM Computer Communication Review*, 35(4):313–324, Aug. 2005.

[18] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 4271, Jan 2006. http://www.rfc-editor.org/rfc/rfc1771.txt.

[19] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica. HLP: A Next Generation Inter-domain Routing Protocol. In *Proc. SIGCOMM*, Aug. 2005.

[20] D. L. Tennenhouse and D. J. Wetherall. Towards an Active Network Architecture. *SIGCOMM Computer Communication Review*, 26(2):5–17, Apr. 1996.

[21] K. Varadhan, R. Govindan, and D. Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report 1, Jan. 2000.

[22] A. Venkataramani, J. F. Kurose, D. Raychaudhuri, K. Nagaraja, M. Mao, and S. Banerjee. MobilityFirst: A Mobility-Centric and Trustworthy Internet Architecture. *SIGCOMM Computer Communication Review*, 44(3), July 2014.

[23] T. Wolf, J. Griffioen, K. L. Calvert, R. Dutta, G. N. Rouskas, I. Baldin, and A. Nagurney. ChoiceNet: Toward an Economy Plane for the Internet. *SIGCOMM Computer Communication Review*, 44(3), July 2014.

[24] W. Xu and J. Rexford. MIRO: Multi-path Interdomain ROuting. In *Proc. SIGCOMM*, Aug. 2006.

[25] X. Yang, D. Clark, and A. W. Berger. NIRA: A New Inter-Domain Routing Architecture. *IEEE/ACM Transactions on Networking*, 15(4):775–788, Aug. 2007.

[26] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang. Named Data Networking. *SIGCOMM Computer Communication Review*, 44(3), July 2014.

[27] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen. SCION: Scalability, Control, and Isolation on Next-Generation Networks. In *Proc. IEEE Symposium on Security and Privacy*, May 2011.