Edge-Aware Inter-Domain Routing for Realizing Next-Generation Mobility Services*

Shreyasee Mukherjee, Shravan Sriram, Dipankar Raychaudhuri
WINLAB, Rutgers University, North Brunswick, NJ 08902, USA
Email: {shreya, sshravan, ray}@winlab.rutgers.edu

Abstract*—This work describes a clean-slate inter-domain routing protocol designed to meet the needs of the future mobile Internet. In particular, we describe the edge-aware inter-domain routing (EIR) protocol which provides new abstractions of aggregated-nodes (aNodes) and virtual-links (vLinks) for expressing network topologies and edge network properties necessary to address next-generation mobility related routing scenarios which are inadequately supported by the border gateway protocol (BGP) in use today. Specific use-cases addressed by EIR include emerging mobility service scenarios such as multi-homing across WiFi and cellular, multipath routing over several access networks, and anycast access from mobile devices to replicated cloud services. Simulation results for protocol overhead are presented for a global-scale Caida topology, leading to an identification of parameters necessary to obtain a good balance between overhead and routing table convergence time. A Click-based proof-of-concept implementation of EIR on the ORBIT testbed is described and used to validate performance and functionality for selected mobility use-cases, including mobile data services with open WiFi access points and mobile platforms such as buses operating in an urban area.

I. Introduction

The inter-domain routing architecture of the Internet is currently based on the border gateway protocol (BGP) standards [1]. BGP, which was introduced about 25 years ago, represented a major advance in networking because it provided fully distributed, non-hierarchical routing mechanisms between autonomous systems (ASes) at a global scale. More importantly, BGP provides a flexible framework for policy-based routing taking preferences and business objectives of individual ASes [2] into account. However the Internet is currently going through a fundamental change driven by the rapid rise of mobile end-points such as smartphones and embedded Internet-of-Things (IoT) devices [3]. There is an increased dynamism caused by end-point, network and service mobility which can take various forms, ranging from conventional end host mobility and edge network mobility to multi-homing and multi-network access associated with emerging hetnet and 5G cellular scenarios as well as dynamic cloud service migration across edge networks. This emerging "mobile Internet" requires different capabilities such as efficient anycasting to cloud-services, awareness of multiple disjoint network paths for multi-homing and alternative and more granular metrics for path selection, than currently supported by BGP.

*Research supported under NSF Future Internet Architecture - Next Phase (FIA-NP) grant CNS-1345295

Emerging Internet requirements have motivated several clean-slate Internet design projects such as Named Data Network (NDN) [4], XIA [5] and MobilityFirst [6]. Previously published works on these architectures have addressed mobility requirements at the intra-domain level [7], [8], but support for end-to-end mobility services across multiple networks remains an important open problem. In this paper, we first motivate the need for clean-slate approaches to inter-domain, and then describe the key features of a specific new design called EIR (edge-aware inter-domain routing) intended to meet emerging requirements. The proposed protocol provides new abstractions for expressing network topology and edge network properties necessary to support a full range of mobility services such as multi-homing over WiFi and cellular, multipath routing over multiple access networks, disconnection tolerant routing and anycast access to cloud services.

It is noted here that such inter-domain mobility service enhancements are also expected to be useful for emerging "5G" scenarios [9] which are associated with heterogeneous access technologies, multiple interfaces, ad-hoc connectivity, etc. Emerging software defined network technologies make it possible to introduce enhanced 5G mobility services in a single domain, but there is still a need for inter-domain solutions which continue to work as mobile users migrate from one operator's network to another or use multi-homing across multiple radio access networks.

The edge-aware inter-domain routing protocol is being proposed as a part of the MobilityFirst Future Internet Architecture (FIA) project [6] aimed at a clean-slate redesign of the IP protocol architecture. Clean-slate research projects like MobilityFirst do recognize the fact that the Internet cannot be changed overnight particularly when dealing with core protocols such as inter-domain routing. However, with the advent of software-based network functionality, it is now increasingly practical to introduce new Internet protocol concepts on a trial basis. In particular, the recently proposed "SDX (software-defined exchange)" concept makes it possible for networks to voluntarily participate in enhanced or new protocol frameworks for inter-domain routing [10]. For example, EIR can be implemented by a small number of ASes as an SDX-hosted function that supplements BGP with the goal of efficiently supporting a specific service such as multi-homing over WiFi and cellular networks. We consider some of these use-cases in further detail in the following section and discuss their implications on the routing layer.

II. Mobility services and their requirements

In this section we describe a few emerging mobility services such as multipath, network mobility and service anycast, and summarize their requirements from the routing plane.

A. Multipath support

A typical mobile device can see multiple cellular or WiFi networks at the same time. Although the majority of current business models restrict a user to a single carrier at a time, with the increasing popularity of "hetnets", mobile devices may soon be able to connect simultaneously to a dynamically changing set of networks [11]. Current solutions for multi-homing [12], [13] rely on transport layer enhancements that require an end-point to inform the sender about its multiple interfaces, and the sender stack to adapt to the packet rate of each interface. However there is no mechanism by which users can specify under what conditions, and in what manner the multiple interfaces are to be used (most economical, highest throughput, all interfaces, etc.). In addition, any end-host driven solution without feedback from the network performs poorly with high latency cellular or lossy wireless links [14]. Therefore, the key requirements for in-network multipath support are: (i) visibility of multiple disjoint paths through the network, and, (ii) mechanism to infer dynamically changing *qualities* of each path in order to utilize them efficiently.

B. Dynamic network formation with disconnection

Another emerging mobility scenario is that of dynamic network formation along with network mobility. For example, a fleet of cars on a highway can form an ad-hoc vehicular network and peer along the edge to different ASes as the network moves. Temporary disconnections and large variation in link quality are also common in such cases. Looking forward to 5G, both the backhaul and the fronthaul will be wireless and may soon become mobile, such as Google's network of aerial balloons as LTE basestations [15]. Managing a global scale of highly mobile radio units is challenging and BGP's partial point solutions [16] cannot scale to a network of hundreds of mobile nodes or respond to changing link quality/capacity at the edge. Inter-domain routing therefore, should have better support for, (i) temporary disconnections, and, (ii) allow finer granularity of path qualities for fast changing links.

C. Service anycast

Emerging cloud-based service applications for on-demand computing or storage often require anycast routing for finding the "closest available resource" based on specialized metrics such as latency or bandwidth. Selection of inter-domain paths based on more than just the BGP reachability metric becomes necessary in such cases and is difficult to achieve without setting up of costly overlays [17]. We believe that the inter-domain routing protocol should (i) provide metrics other than the traditional shortest AS hop count, and, (ii) allow means of flexible path selection based on these metrics.

As we can see different emerging mobility services require a different set of capabilities from the routing plane which is summarized in Table I.

TABLE I: Networking requirements for next-gen applications

Applications	Network Stack and Architectural Requirements	
	Meta-level Service	Network Service
Vehicular	Uni/multi/broadcast	Delay tolerance, high variability
End-user	Uni/multicast	Mobility, multipath
Sensor-data	Multi/reverse-multicast	In-network storage, computation
Cloud	Anycast	Multipath, alt. path metrics
Disaster mgmt	Multi/broadcast	Delay tolerance, storage

III. Edge-aware inter-domain routing

Based on the above considerations, in this section we present the design rationale and the key building blocks of our proposed edge-aware inter-domain routing (EIR).

A. Design Concepts

1) In-network mapping of names to addresses: The concept of separating names from addresses has been used in several proposals (MobilityFirst [6], LISP [18], HIP [19]). The infrastructure for mapping names to addresses can either be external to the network and accessed only by end nodes, or alternatively be implemented in-network and be accessible by both end-hosts and routers. We make use of the in-network approach, to ensure delivery of packets in the case of fast end-host mobility. All objects (devices, routers, access points, etc.) are assigned unique names and a global name resolution service (GNRS) maintains mappings between a name and its routable address(es). Several past works have shown the feasibility of Internet-scale, distributed, in-network mapping infrastructure with extremely small query-response time [20], [21].

2) Increased visibility of alternative paths: Between any two networks, multiple routes are usually available and these routes can entail vastly different properties [22]. In BGP, a network might learn about alternate routes but can only select and propagate one "best" route to other networks, which leads to a myopic view of the network graph. Even though solutions extending BGP [23], [24] do exist, they are mostly used for failure-recovery or *best-path* applications. In order to support the increasingly important use-cases of multipath and multi-network operations, EIR entails network-wide visibility of multiple possible paths between each pair of networks.

3) Propagating alternative properties in inter-domain routing: BGP does not propagate link-quality metrics in its routing updates making it difficult to differentiate paths based on metrics other than hop count. For e.g., in an early in-flight WiFi implementation, Boeing associated each flight with an IP address block which was announced into BGP from different locations as the plane moved [16]. Other networks receiving such announcements had no idea that the last hop had a ground-to-plane high variability wireless link instead of the usual high-capacity peering-point wired link. In EIR, coarse-grained information about aggregate links is propagated through the routing protocol to enable networks to make forwarding decisions based on alternate network properties.

B. Key Building Blocks

1) Aggregated nodes (aNodes) and virtual links (vLinks): Each AS has the option of dividing its routers and other network elements (such as access points and base-stations)

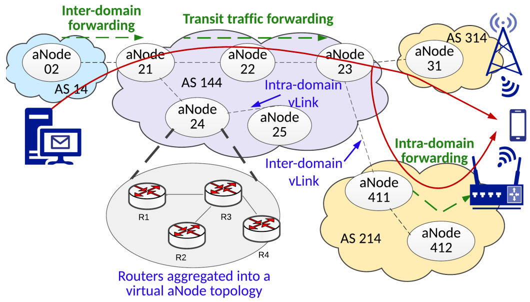

Fig. 1: aNode-vLink topology abstraction for an AS

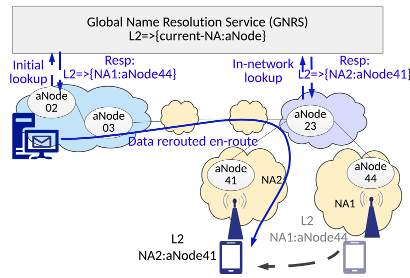

Fig. 2: Late-binding of data to counter destination mobility

into one or more groups (called aggregated nodes or aNodes) as shown in Fig. 1. Entities belonging to the same aNode typically share some common operational or physical attributes. Possible compositions of aNodes include: the entire AS (similar to BGP); group of routers in a geographical area; all routers that support flow-based routing (for e.g. through OpenFlow); and wireless routers on bus/train/plane networks. Connectivity between aNodes is expressed through virtual links (vLinks) with aggregate link properties such as mean latency, average bandwidth, variability, etc.

The aNode-vLink abstraction allows a network to partially expose its internal connectivity structure while limiting it to a level of detail that fits its needs. Networks that do not wish to expose internal structure describe themselves as a single aNode. A network state packet (nSP) is used to inform other ASes of the network's aggregate topology graph along with the vLink properties. Similar to the previously proposed Pathlet routing [25], this abstraction is quite versatile and allows different ASes to exert different levels of control over a fragment of the end-to-end path. By choosing different levels of aggregation, a domain can control traffic patterns that traverses into and inside its network and can also offer its clients flexible route selection as a value-added service.

2) Telescopic flooding of network state: Internal to an AS, routers use intra-domain routing to build an internal network-graph [7]. Border routers construct nSPs by combining the internal network topology with the management enforced aggregation-level and export policies to build the virtual aNode-vLink topology. The nSPs are then announced to neighboring ASes and propagated throughout the network. However, in order to limit the control overhead, the border routers relay nSPs that originated from other ASes in a *telescopic manner*. This means that, on receiving an nSP, a border router holds it for a certain amount of time which is proportional to the distance i.e. AS hop count from the source of the nSP, before forwarding it out. As a result, an AS will get more frequent (hence up-to-date) routing updates from ASes that are closer to it. The term "telescopic" comes from the analogy of distant nodes seeing each other through the reverse-end of a telescope, i.e. they are visible but less clearly so. Telescopic flooding ensures that every network has a global view of the inter-domain topology, but with acceptable control overheads. Different telescopic functions can be defined by changing the relation between the hold-delay (time for which a border router holds a received nSP before relaying it to other neighbors) and the hop-count. We have extensively analyzed the effect of different hold functions on nSP distribution overhead and route convergence time. Please refer to our technical report [26] for details.

As a side-effect of telescopic route updates, network states observed from far away could be stale. However, our intuition is that, up-to-date path quality information for the entire path is possibly unnecessary at the source. For e.g. if a source located in China wishes to send data to a mobile network in a bus in USA, the path quality of the last few hops near the source, is probably stale. However, from the source point of view, the route does not change until the packet reaches a border router in the US, at which point, the path information is much fresher to route through appropriate alternative paths. This is a key difference between source routing concepts of Pathlet and EIR both of which use similar aggregate topology information, but different dissemination techniques.

3) Late-binding for mobility support: As a side-effect of mobility, a destination may move during transit of a data packet and thus result in delivery failure. To address this, EIR incorporates in-network name-to-address binding during the transit of a packet. This serves as a fail-safe mechanism that allows routers to actively react to link variations and mobility of end nodes. In particular, EIR makes use of a fast in-network name-to-address resolution through the GNRS [20] in order to retrieve the current network location of the destination. As shown in Fig. 2, network-address mapping of in-transit data can be looked up at an intermediate location within the network to properly route to a new location, without failure in delivery. Different late binding algorithms have been studied in detail in our work [27] and one of them has been evaluated on a real-world dataset with our prototype in Sec. IV.

In addition to the above mentioned key features, EIR also supports a variety of routing algorithms and has flexible policy specification semantics that go beyond standard BGP policies. Please refer to our technical report for further details [26]. Fig. 3 brings all the discussed features together to show how end-host multihoming can benefit from EIR. In this e.g., a device with name *E2*, is connected to two different networks through WiFi and LTE at the same time and wishes to receive data across both the interfaces. The GNRS stores the up-to-date mapping of *E2*'s name to addresses along with its intent of receiving data across both. In-network GNRS lookup binds *E2* name to *NA1* and *NA2* as well as expresses *E2*'s delivery intent through the use of a service identifier (SID) in the packet

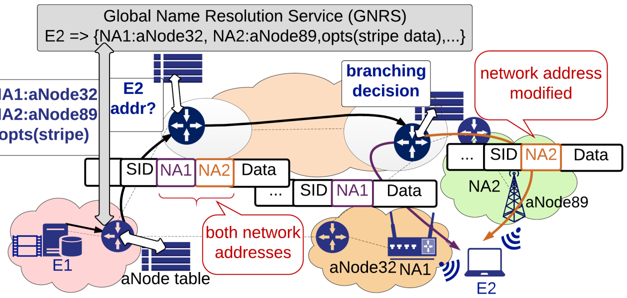

Fig. 3: Multihoming with data delivery through WiFi & LTE

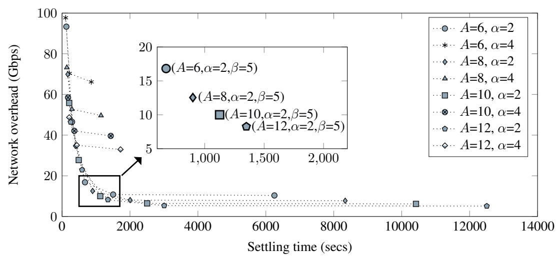

(a) Overhead vs. settling time for different parameters of the constant-exponential-constant telescopic function

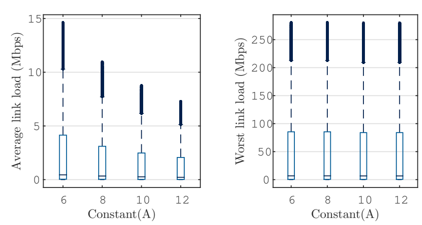

(b) Average and worst case load on links for values that provide a good tradeoff

Fig. 4: Internet scale simulations

as shown. Every border router looks at *NA1* and *NA2* and takes an independent decision based on their aNode forwarding table whether or not to bifurcate the data stream and therefore, can adapt well to network changes and fluctuations in link quality.

IV. Evaluation

In this section we evaluate the EIR protocol in terms of scalability and mobility service performance through an Internet-scale simulation and a Click-based prototype evaluation.

A. Simulation for Internet scale overhead and scalability

One of the main challenges in propagating network state packets (nSPs) throughout the Internet is scalability. By using different telescopic functions, this overhead can be reduced, but at the cost of slower route convergence. Our detailed NS3 simulation in [26] using a 200 node Jellyfish topology (which closely resembles the Internet topology) and link latencies from the Dimes database [28] has shown a constant-exponential-constant telescopic function to have *reasonable* route convergence time. This telescopic function of hold-delay, y vs. hop count, x is represented as:

$$y = \begin{cases} A, \text{if } x < \alpha \\ A\ e^{(x-\alpha)}, \text{if } \alpha \le x < \beta \\ A\ e^{(\beta-\alpha)}, \text{if } x \ge \beta \end{cases} \tag{1}$$

where, A, α, β are constants. In order to choose *reasonable* values of these constants (values that result in low overhead as well as low convergence time), we simulate a complete Internet topology dataset available at Caida [29] in our custom simulator. This dataset is composed of 47,445 ASes and 200,812 links using which, we simulate the generation and propagation of 1000 byte nSPs across the network. Fig. 4(a) shows the global routing overhead vs. settling time for different values of A, α and β. As shown in the highlighted box in the figure, there is a small subset of parameter values ($\alpha = 2, \beta = 5$ and $A \in \{6, 8, 10, 12\}$) that result in low overhead as well as low settling time, which, therefore, can be used for setting the telescopic function at each AS in a real-world setting. Notice, that even to achieve very low settling times, the worst case network overhead is about 100 Gbps. This is a negligible fraction of the total Internet traffic of ∼182 Tbps as of 2014 [3]. Fig. 4(b) further plots the average and worst case link load for these subset of parameter values. As seen from the plot, average link load decreases with increasing the hold delay, since routers propagate nSPs less frequently. It however does not exceed 15 Mbps while the worst case link load does not exceed 300 Mbps. The latter also remains almost constant for different parameter values due to the bursty nature of the instantaneous load on a link.

B. Mobility experiments with prototype

To measure the performance and implementation feasibility of EIR, we have built a prototype router based on the Click [30] and evaluated it on the ORBIT testbed [31]. In order to evaluate mobility support in EIR, we used a realistic inter-domain topology and a probabilistic mobility transition matrix which is briefly described below.

1) Topology generation and probabilistic mobility: We start with a 2012 Caida dataset with point-of-presence (PoP) topology and parse the dataset based on cities, specifically focusing on San Francisco, which has 28,052 PoP nodes belonging to 354 ASes. We consider a cooperative scheme where a multitude of ASes agree to share connectivity among their customers, i.e. a user can decide to switch from one network provider to another when moving, provided the latter provides a better coverage in the region. In order to keep the number of nodes in the experiment tractable, we choose 15 random ASes which participate in this cooperative scheme. Since AS tier information was not available, a random choice ensures that we get a good mix of ASes from different tiers. Given the reduced topology of 1,327 PoPs, a corresponding aNode topology is developed based on geographical proximity, i.e., PoPs belonging to the same AS and located close to each other are clustered to the same aNode. This results in a final inter-domain EIR topology of 53 aNodes.

TABLE II: Probabilistic transition for user mobility

Basic parameters:	
Z	avg number of network transitions/sec
K	total number of network transitions
T	granularity of transition (sec)
r	avg distance to neighbors (meters)
s	avg speed of mobility (m/sec)
$w = s/r$	average transition rate/sec
α	probability of transition to a network
Transition probability from node N_j:	
$\alpha(wT)/N_j$	to each of N_j's neighbors
$(1-\alpha)(ZT)/K$	to each of K non-neighbors

In order to realistically model user mobility across domains, we generate a probability matrix for network transition taking into account: (i) local mobility within a certain radius (denoted as r), with biased transitions between aNodes belonging to the same AS (users tend to remain connected to the same network provider as they move, unless no connectivity by the current provider is available at the new location); (ii) equal probability of transition to the rest of the aNodes if local AS aNodes are not available; and, (iii) biased transitions (determined by α) to a random, k number of "macro mobility" points based on the average number of networks visited by a user per day [32]. Table II explains the transition probability computations.

2) Mobility support through late binding: Based on the San Francisco topology and a transition matrix generated for a typical mobile user, we analyze the advantage of late binding for user mobility support. The evaluated parameter is path stretch incurred with and without late binding and defined as the ratio of number of hops traversed by a packet to the number of hops across the shortest path. Note that without late binding, failure in delivery would result in rebinding through a GNRS re-lookup at the previous point of attachment. On the other hand, late-binding would re-bind the network address at an intermediate router, as explained in Sec. III-B3. The late-binding algorithm for this evaluation chooses the aNode with the highest degree along the path as the late-binding point. The intuition behind this logic is that a highly connected node would have shorter path stretch to the next point of association for the user. Fig. 5 highlights the improvement in path stretch when packets are late binded along the way. Notice that the solid blue and the dotted red curves are fairly close since only the packets in transit are rerouted and suffer a path stretch, whereas newer packets are automatically sent to the new destination, from the source, following a GNRS lookup. We are also looking at alternative late binding algorithms, so as to minimize path-stretch and improve latency of data delivery across a broad range of mobility scenarios [27].

3) Network mobility support: Next, based on the same topology, we consider a scenario where a source sends data to a network with mobile aNodes. To realistically model network mobility, we use actual bus traces from San Francisco Municipal Transit system [33] and measure data delivery failure rate for different telescopic hold delay rates. Since MobilityFirst architecture ensures packet delivery through storage and rebinding, for the purpose of this experiment, we calculate failure anytime a packet needs to be stored and rebinded.

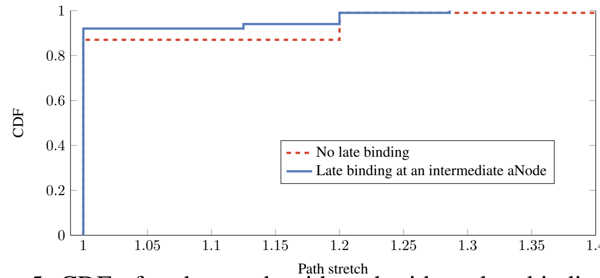

Fig. 5: CDF of path stretch with and without late binding for end-user mobility

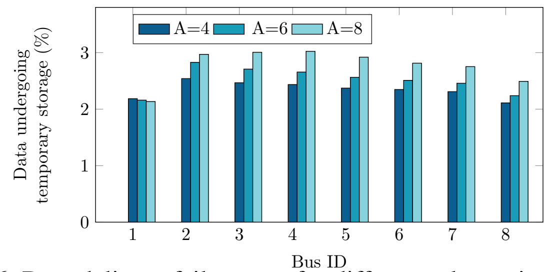

Fig. 6: Data delivery failure rate for different telescopic update intervals for network mobility

Fig. 6 shows the percentage of packets undergoing rebinding at mobile buses on 9 randomly picked routes, for different hold delays of the telescopic function in Eq. 1 from Sec. IV-A. Note that increasing the hold delay parameter (A) leads to slower update of routing tables and therefore leads to more packets being routed across a stale path. Similar to our previous experiments, the values of α and β were kept constant at 2 and 5 respectively as they provided *reasonable* overhead and settling time, based on our Internet-scale simulation. We also looked at the number of AS transitions for each trace which determines the failure rate and observed that 2 hops AS transition tend to dominate these mobility events. Of the 9 randomly picked traces, trace 1 resulted in a scenario that had primarily 1-hop transitions and hence the data delivery rate is almost similar for different telescopic hold time. Whereas in the other traces, there are a few transitions to ASes that are multiple AS-hops away. Consequently, the failure rate increases with A as the up-to-date reachability information is not known for a longer period of time due at the source.

V. Related Work

There has been a considerable amount of work done in improving inter-domain routing as (1) extensions to BGP, and through (2) clean-slate routing proposals.

Proposals such as path splicing [34] and route-deflections [35] are loose source routing based schemes, where the end-hosts are assumed to be intelligent enough to to explicitly choose a path alternative to the default BGP-computed route. [35] provides a limited choice of paths, whereas [34] provides path diversity without addressing scalability. MIRO [23] moves the decision of path choice from the end-host to the AS which could request alternate paths if *not satisfied* with the default BGP route. This handles scalability effectively, but reduces path diversity. [24]

proposes similar failover path set-up techniques in order to reduce disconnectivities on link failures.

There has also been a growing interest in the Internet community to look for alternatives of BGP that could be incrementally deployed. As mentioned before, the aggregation techniques in EIR is similar to Pathlet [25]. However, our path-selection approach is quite different from Pathlet, which performs loose source routing. Instead EIR allows routes to be updated while a packet is in transit through telescopic route updates and late-binding of names to addresses. HLP [36] proposes a hybrid link-state and path-vector approach that effectively improves scalability of the protocol but restrictive for policies beyond simple business relationships. NIRA [37] offers more choice to end-users using a hierarchical provider-rooted address scheme. However, similar to HLP, the basic protocol provides limited support for policies other than business relationships. In comparison, EIR's aggregations scheme is quite flexible for realization of a wide range of policies [26].

VI. Conclusion

In this paper, we have proposed the edge-aware inter-domain (EIR) routing protocol as a potential routing solution for the future mobile Internet. The proposed architecture has been shown to provide improved support and flexibility for routing to wireless devices, network-assisted multipath routing, routing to multiple interfaces (multi-homing) and service anycast. Our results show that even with increased expressiveness of network structure and node/link properties, the protocol can be designed to have reasonably small overhead via telescopic dissemination of the nSPs. Further, prototype evaluations of the protocol using Click software routers on the ORBIT testbed were conducted to show proof-of-concept level feasibility. Experimental results for selected use-cases show good service level performance can be achieved in highly mobile scenarios. For further work, we plan to deploy EIR on the GENI large scale testbed to evaluate service capabilities and performance in more realistic global network scenarios.

References

[1] Yakov Rekhter, Tony Li, and Susan Hares. A border gateway protocol 4 (bgp-4). Technical report, 2005.

[2] Matthew Caesar and Jennifer Rexford. Bgp routing policies in isp networks. *Network, IEEE*, 2005.

[3] Cisco Visual Networking Index. Global mobile data traffic forecast update, 2014-2019. *White Paper, February*, 2015.

[4] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, Patrick Crowley, Christos Papadopoulos, Lan Wang, Beichuan Zhang, et al. Named data networking. *ACM SIGCOMM CCR*, 2014.

[5] Dongsu Han, Ashok Anand, Fahad R Dogar, Boyan Li, Hyeontaek Lim, Michel Machado, Arvind Mukundan, Wenfei Wu, Aditya Akella, David G Andersen, et al. Xia: Efficient support for evolvable internetworking. In *NSDI*, 2012.

[6] Dipankar Raychaudhuri, Kiran Nagaraja, and Arun Venkataramani. Mobilityfirst: a robust and trustworthy mobility-centric architecture for the future internet. *ACM SIGMOBILE Mobile CCR*, 2012.

[7] Samuel C. Nelson, Gautam Bhanage, and Dipankar Raychaudhuri. GSTAR: Generalized storage-aware routing for MobilityFirst in the future mobile Internet. In *Proceedings of MobiArch'11*, 2011.

[8] AKM Hoque, Syed Obaid Amin, Adam Alyyan, Beichuan Zhang, Lixia Zhang, and Lan Wang. Nlsr: Named-data link state routing protocol. In *Proc. of Workshop on ICN*. ACM, 2013.

[9] Intel. 5G: A Network Transformation Imperative. *White Paper, December*, 2015.

[10] Arpit Gupta, Laurent Vanbever, Muhammad Shahbaz, Sean P Donovan, Brandon Schlinker, Nick Feamster, Jennifer Rexford, Scott Shenker, Russ Clark, and Ethan Katz-Bassett. Sdx: A software defined internet exchange. In *Proceedings of ACM SIGCOMM*. ACM, 2014.

[11] Project Fi. https://fi.google.com.

[12] Janardhan R Iyengar, Paul D Amer, and Randall Stewart. Concurrent multipath transfer using sctp multihoming over independent end-to-end paths. *IEEE/ACM Transactions on networking*, 2006.

[13] Alan Ford, Costin Raiciu, Mark Handley, Sebastien Barre, and Janardhan Iyengar. Architectural guidelines for multipath TCP development, RFC 6182. 2011.

[14] Parishad Karimi and Dipankar Raychaudhuri. Achieving high-performance cellular data Services with multi-network access. In *Proceedings of IEEE Globecom 2016 (to appear)*. IEEE, 2016.

[15] Project Loon. http://www.google.com/loon/.

[16] B Abarbanel. Implementing global network mobility using bgp. In *NANOG Presentation, http://www. nanog. org/meetings/nanog31/abstracts. php*, 2004.

[17] Hitesh Ballani and Paul Francis. Towards a global ip anycast service. In *ACM SIGCOMM CCR*. ACM, 2005.

[18] D Meyer and D Lewis. The locator/id separation protocol (lisp). *IETF RFC 6830*, 2013.

[19] Robert Moskowitz, Pekka Nikander, Petri Jokela, and Thomas Henderson. "host identity protocol". *IETF RFC 5201*, 2008.

[20] T. Vu et al. DMap: A Shared Hosting Scheme for Dynamic Identifier to Locator Mappings in the Global Internet. In *Proceedings of ICDCS '12*, 2012.

[21] Changhoon Kim, Matthew Caesar, and Jennifer Rexford. Floodless in seattle: a scalable ethernet architecture for large enterprises. In *ACM SIGCOMM CCR*. ACM, 2008.

[22] Yi Wang, Ioannis Avramopoulos, and Jennifer Rexford. Design for configurability: rethinking interdomain routing policies from the ground up. *Selected Areas in Communications, IEEE Journal on*, 27(3):336–348, 2009.

[23] Wen Xu and Jennifer Rexford. Miro: Multi-path interdomain routing. In *PhD thesis, Princeton University*, 2009.

[24] Feng Wang and Lixin Gao. Path diversity aware interdomain routing. In *INFOCOM 2009, IEEE*.

[25] P Godfrey, Igor Ganichev, Scott Shenker, and Ion Stoica. Pathlet routing. In *ACM SIGCOMM CCR*. ACM, 2009.

[26] Shreyasee Mukherjee, Shravan Sriram, Tam Vu, and Dipankar Raychaudhuri. EIR: Edge-aware Interdomain Routing Protocol for the Future Mobile Internet. Technical Report WINLAB-TR-443, WINLAB, June 2015.

[27] Suja Srinivasan and Dipankar Raychaudhuri. GNRS assisted inter-domain routing in MobilityFirst. In *MS thesis, Rutgers University*, 2016.

[28] Yuval Shavitt and Eran Shir. Dimes: Let the internet measure itself. *ACM SIGCOMM CCR*, 2005.

[29] The CAIDA UCSD Internet Topology Data Kit. http://www.caida.org/data/internet-topology-data-kit.

[30] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M Frans Kaashoek. The click modular router. *ACM TOCS*, 2000.

[31] Dipankar Raychaudhuri, Ivan Seskar, Max Ott, Sachin Ganu, Kishore Ramachandran, Haris Kremo, Robert Siracusa, Hang Liu, and Manpreet Singh. Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols. In *WCNC*. IEEE, 2005.

[32] Zhaoyu Gao, Arun Venkataramani, James F Kurose, and Simon Heimlicher. Towards a quantitative comparison of location-independent network architectures. In *SIGCOMM*. ACM, 2014.

[33] SFMTA Municipal Transport Agency. https://www.sfmta.com/.

[34] Murtaza Motiwala, Megan Elmore, Nick Feamster, and Santosh Vempala. Path splicing. In *ACM SIGCOMM CCR*. ACM, 2008.

[35] Xiaowei Yang and David Wetherall. Source selectable path diversity via routing deflections. In *ACM SIGCOMM CCR*. ACM, 2006.

[36] Lakshminarayanan Subramanian, Matthew Caesar, Cheng Tien Ee, Mark Handley, Morley Mao, Scott Shenker, and Ion Stoica. *HLP: a next generation inter-domain routing protocol*.

[37] Xiaowei Yang, David Clark, and Arthur W. Berger. Nira: A new inter-domain routing architecture. *IEEE/ACM Transactions on Networking*, 2007.