Value Flows: Inter-Domain Routing over Contract Links

Hasan T. Karaoglu, Murat Yuksel
University of Nevada - Reno, Reno, NV 89557.
karaoglu@cse.unr.edu,yuksem@cse.unr.edu

***Abstract*—Due to the Internet's commercialized ISP structure, inter-domain routing is perceived as more of a commercial service (and market decision) instead of an optimization problem subject to purely technical criteria. In this perspective, routing could be defined as a service which matches perishable resources to requested level of service demands. This diversity in demanded services and proliferation of diverse set of existing and emerging applications cannot be responded by a single mindset of priorities, limited set of metrics or protocol characteristics. Instead, we propose a contract-based routing service architecture which allows manageable dynamism or routing protocol characteristics. By considering routing as a contracted service, network traffic can be treated as "value flows" rather than bare bits. Such value-based treatment of routing will significantly help improving inter-domain routing economics and dynamics, where policy can be expressed as contracting terms. In this paper, we present a Link-State Contract Routing (LSCR) protocol for long-term services and explore how inter-domain routing dynamics could be managed by adjusting "contract term" or routing service lifetime.**

I. Introduction

A. Motivation and Problem Definition

Today, the Internet is a highly commercial medium where user choice is a fact and matching diversity in offered and demanded services is a requirement to survive. However, the current Internet architecture needs market mechanisms that can both (i) equip customers with simple knobs and interfaces so that they can express their willingness to pay for enhanced services and (ii) enable providers to inform customers about the cost of their choices.

Recent research findings show that routers are becoming more busy due to increased and diversified demand on routing. Just like Internet faced threatening strain due to routing pathologies earlier [1], today it is threatened by the strain on the routing function [2]. The recent increase in forwarding dynamics, multi-homing and visibility of prefixes, and growing BGP routing tables are considered to be major challenges that cannot be solved via patch solutions and require architectural revisions [3].

In some cases, requirements and interests on routing services may not be necessarily in harmony with each other, such as seeking maximum available bandwidth and stability of routes at the same time or quick convergence and sensitivity on dynamic conditions on route. It would be a big burden on routing process by itself to achieve such diverse e2e requirements by keeping the point-to-anywhere inter-ISP settlements.

This work is supported in part by National Science Foundation awards 0721600 and 0721609.

Moreover, the current Internet architecture needs a major overhaul so as to recognize the significance of value exchanged through traffic flows and business interactions over service-level agreements. Asymmetric traffic ratios, disputes on revenue-sharing structure and ambiguity in direction of value flows imbalance the relative positions of ISPs in terms of economics [4]. These pitfalls lead to complicated bilateral agreements shaped by power play of negotiating parties which threaten the stability of the Internet.

B. Solution Proposal

Our proposal is to define inter-domain routing as a service described by a sequence of "contracts", each of which explicitly states the service and performance as a set of contracting terms. In this contracting framework, economic components such as cost of providing enhanced services and the direction of value exchanged can be determined via point-to-point negotiations. In our approach, network traffic can be treated as a "value flow" rather than bare bits. So, our service definition allows: (i) aggregation of value flows at network core while preserving packet switching capabilities at network edge, (ii) isolate traffic flows with conflicting interest or performance requirements via different contracts, and (iii) define the value of service both in terms of compensation and direction of flow.

Our proposed inter-domain routing architecture employs "contracts" as its building blocks, and allows flexible, finer-grained and dynamic contracting over multiple providers. With such capabilities, the Internet itself will be viewed as a *"contract-switched"* network beyond its current status as a "packet-switched" network [5]. A contract-switched architecture will enable flexible and economically efficient management of risks and value flows in an Internet characterized by many tussle points [6].

C. Contributions and Major Findings

In this paper, we present an inter-domain routing framework, Contract Routing (CR), as an overlay on top of the existing Internet routing protocols. We build upon our prior work [5] by experimentally developing and evaluating various CR protocols. Particular contributions of this paper include:

- *Edge-to-Edge (g2g) Contract Links:* We define an g2g service abstraction called "contract link" which is the key building block for enabling technologies and services that are otherwise impractical to implement in the current Internet architecture.
- *Contract Routing (CR):* We outline CR protocols that enable automatic establishment of SLAs among service

providers. By means of CR, the providers can automatically compose e2e services by concatenating the contract links, and the users can automatically search for existence of an e2e contract path satisfying their requirements.

- *Link-State Contract Routing (LSCR):* As an overlay on top of OSPF and BGP, we develop a proactive CR protocol for long-term and large contract paths.
- *Contract Term Analysis:* We analyze CR behavior with varying service lifetime (or contract term). Our analysis shows that contract term is highly effective on determining routing behavior, performance and dynamics.
- *Value Flows:* We introduce the concept of "value flows". Since CR utilizes contract links with explicit economic components, traffic flows in CR become "value flows" recognizing the value exchanged via Internet traffic.
- *Tradeoffs in CR Dynamics:* We investigate several key tradeoffs in CR design, such as stability vs. adaptivity, contract term vs. QoS, and path availability vs. e2e QoS.

D. Organization

In Section I-E, we cover the literature on inter-domain routing architectures.In Section II, we define contract links as the building blocks of the contract routing framework and discuss key design considerations of our work such as service granularity, economics and routing via multiple protocols. In Section III, we present implementation of Link-State Contract Routing (LSCR) as an overlay on top of BGP and OSPF. In Section IV, we make an evaluation of LSCR behavior in a realistic network model with various contracting granularity and terms. We summarize our work in Section V.

E. Related Work

It has been long discussed by researchers that the current Internet architecture should be renovated, so that it can respond to the challenges and necessities of the future. One of the core issues being in these discussions is to improve routing performance by leveraging inter-ISP negotiations and more dynamism in inter-domain economics as proposed by Negotiation-Based Routing [7] and Path Trading [8] schemes. However, leveraging negotiations in routing process requires open interfaces and mechanisms to allow required exchange of information [6]. While clean-slate approaches like NIRA [9] and HLP [10] propose exploitation of hierarchy in the Internet topology, MIRO [11] has an approach to the same problem proposing multi-path routing as an alternative to single-path routing. So far, the common side of the routing proposals is to utilize the current Internet topology characteristics such as existing inter- and intra-domain path diversity [12], of which BGP currently can not explore well [13]. Pathlet Routing [14] offers advertising g2g services, called "pathlets", which allow composition of e2e paths by source in a rich policy setting which allows emulation of wide range of routing protocols.

Creation of a market of Internet connectivity services based on g2g services is proposed by several researchers [5], [15], and has significant potential for attaining bundles for high quality e2e services [16], [14], [17]. We propose routing over g2g service abstractions called "contracts", which explicitly embed economics into the routing protocols. In our routing framework, traffic flows are treated as "value flows" which allows better management of "routing economics".

II. Contract Routing (CR) Framework

In order to integrate economic flexibilities into the inter-domain routing and inter-ISP interactions, we formulate the problem of finding economically and technically efficient e2e paths as a "contract routing" problem. In our contract routing framework, each ISP is abstracted as a set of g2g virtual links, which we call "contract links". These unidirectional virtual links are connecting edge routers called Contract Routers residing at the edge of ISP domain boundaries. So, a single domain is represented by a virtual contract topology consisting of contract routers and contract links connecting these routers at ingress and egress points. This is a significantly different approach than the existing routing architecture which abstracts each ISP as a single dot in the inter-domain topology map, which was mainy motivated by the need to hide an ISP's internal network information. In regards to the confidentiality of an ISP network, our contract routing framework does not expose more information than what can be inferred from the existing BGP announcements [18].

Once this abstraction of virtual contract topology is made, then g2g services can be defined on top of these g2g contract links with various technical and economic promises. [17] Given availability of such g2g contract links, composing e2e services becomes a calculation of "shortest" contract paths, where "shortest" might mean different measures (e.g., lowest cost, highest quality) for each customer.

Table I comparatively shows the key differences between our contract routing framework and the existing inter-domain routing.

A. Key Design Considerations

Routing via Multiple Protocols: Before explaining novelty of our approach, we ask a simple question: *What makes a routing protocol different than others?* Definition of route information, dissemination and exchange mechanism of these routes, path calculation techniques and resulting forwarding dynamics can be listed as answers. Varying characteristics shaped by these components are actually the reflection of various requirements imposed by medium, traffic patterns and application-specific performance metrics. We believe that future Internet architecture should not be limited to a pattern, characteristic or a mindset imposed by a single routing protocol. A single routing protocol cannot answer requirements of networks as extreme as quantum networks, DTNs or mobile networks.

Contract routing carries out routing tasks via multiple routing protocols. Common to all these protocols, they operate on contracts as building blocks. Since contract definitions are flexible, they can be used to compose e2e services which handle wide variety of routing characteristics demanded by diverse set of applications or medium access technologies.

Contracting Granularity: Desired routing characteristics are various and may look conflicting. However, we believe that in a service setting which requires reservation of limited resources, these characteristics can be generalized according to their temporal behavior and their required level of dynamism.

TABLE I
COMPARISON OF CONTRACT ROUTING VS. CURRENT INTER-DOMAIN ROUTING

CURRENT INTER-DOMAIN ROUTING	CONTRACT ROUTING
An ISP is abstracted as a single node	An ISP is abstracted as a set of contract links
Point-to-anywhere inter-ISP settlements	Possible point-to-point inter-ISP sett lements
SLA establishments take long time	SLA establishments are automated
Min-hop routing (with the exception of policies)	"Shortest" might be longer but cheaper
Economics is implicitly embedded in routing	Economics is explicitly embedded into routing decisions
Only destination-based and single-path routing is possible	Source-based, multi-path routing is possible
Users cannot choose among different e2e paths	Users can choose among multiple alternative paths
Users can only choose the access ISP	Users can choose intermediate ISPs
Same price for all ingress-to-egress flows	Different prices are possible for ingress-to-egress flows

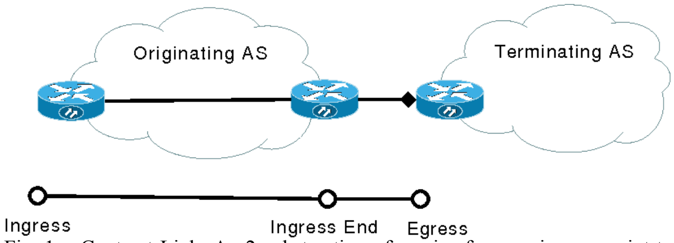

Fig. 1. Contract Link: A g2g abstraction of service from an ingress point to an egress point

For example, applications requiring an on-demand, on-the-fly composition of dynamic services can be served by a path-vector protocol which operates on short-term, micro time-scale contracts (e.g. span of tens of minutes to several hours). Whereas, applications which demand stable e2e routes for a long time span can be served by a less dynamic but more stable link-state protocol operating at macro-time scales (e.g. several days, months or years). An example of e2e service composition for link-state CR protocol is given in Figure 2.

B. Contract Links: Edge-to-Edge (g2g) Service as a Point-to-Point Contract

In CR, service providers advertise their g2g services as contracts, a.k.a. contract links (see Figure 1). These contract links can be defined as an embedding of three major components: (i) *performance*, (ii) *financial*, and (iii) *time*. The performance component of a contract link can include QoS metrics such as delay, loss and guaranteed bandwidth. The financial component will include various fields to aid negotiating parties of the contract link in making financial decisions related to value and risk tradeoffs involved in engaging in the contract. The basic financial component fields can be various prices, e.g., spot, forward, and usage-based. It is possible to design interesting financial component fields identifying financial security and viability of the contract, e.g., whether or not the contract is insured or has money-back guarantees. The time component can include operational time-stamps and be useful for both technical decisions by network protocols and economic decisions by the contracting entities. Example time component fields are the contract duration and the time left for the insured term when the money-back guarantee will expire. Notice that all three components operate over an aggregation of several packets instead of a single packet.

III. LINK-STATE CONTRACT ROUTING

Link-State Contract Routing (LSCR) is an optimized link-state protocol which operates at *macro* time-scales, e.g., several hours, days or longer. Although LSCR requires global coordination, it is significantly different from other well-known protocols in the same characteristic family (e.g. OSPF

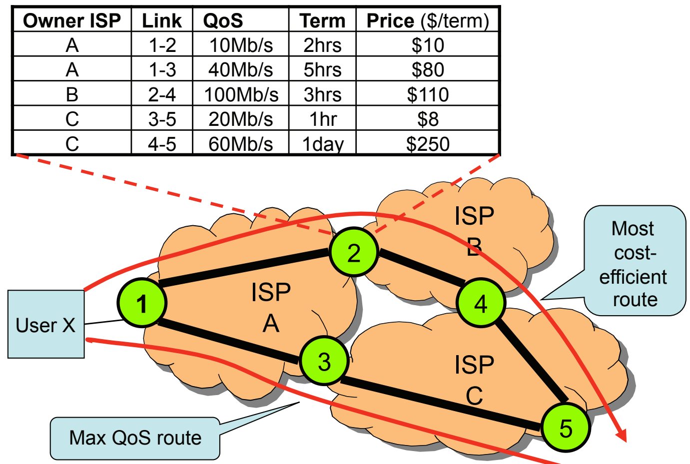

Fig. 2. Scenario for LSCR: The route 1-2-4-5 is the most cost-efficient contract path (i.e. (10Mb/s*2hrs + 100Mb/s*3hrs + 60Mb/s*24)/($10 + $110 + $250) = 27.2Mb/s*hr/$), while the 1-3-5 route is better in terms of QoS

[19]). In LSCR, there are no regular link-state updates, Hello messages which would be prohibitive considering the nature of inter-domain routing and scale of the Internet.

LSCR treats each contract link being advertised by an ISP as a "contract-link-state", and requires these contract-link-states to be updated under two cases: (i) when an ISP wants to advertise a new contract link and (ii) when an ISP wants to update/withdraw an already-advertised contract link. The frequency of these two cases can be too high if the contract durations are too short and accumulated changes over the contract link's g2g path become significant. This is the design rationale for why LSCR operates at macro time-scales. The frequency of the latter case depends on how robust the network is in the ISP originating the contract link as well as how rigorous the promised contract terms are. The ISPs will have to carefully assess their network dynamics and promise performance guarantees based on their risk-benefit tradeoff.

Like in other link-state routing protocols, LSCR does require a broadcast of an update on a contract-link-state. This means all the contract routers will have to be informed about the update on a contract-link-state. This can be a cumbersome process if done too frequently. Further, depending on the agreements among LSCR-capable ISPs, this update process may potentially cost money to the ISP who asks other ISPs to update their contract routing tables. Thus, service providers, who want to advertise their g2g services via LSCR, should make a long-term assessment and planning of their infrastructure and resources, so that their advertised contracts will not need to be updated for long duration of time. As previous research has shown, routes to popular Internet destinations are usually stable and their lifetimes can span in the scale of weeks and longer [20]. We want to exploit this route stability via long-term enhanced e2e paths composed by LSCR. Matching long-term traffic demand to these stable routes is the right way of managing LSCR's scalability.

A. Composition of Contract Links

We implemented LSCR on realistic scenarios where it interacts with both OSPF as intra-domain and BGP as inter-domain routing protocols. So, contract routers (which sit at domain borders) running LSCR protocol keep track of current intra-domain topology and available g2g resource capacity by monitoring information bases of BGP and OSPF-TE protocols. Central coordinators aggregate these topology information fed by contract routers. The central coordinators can instruct contract routers according to business policies. Considering business policies and risk evaluations, contract routers will decide to advertise, update or withdraw the contract links.

As seen in Figure 1, a contract link is a unidirectional virtual link which consists of two integral components: 1) intra-domain g2g part, 2) inter-domain connection. The first part represents intra-domain virtual link which connects entry and exit points of ISP A, while the second part represents the inter-domain link which connects exit point of ISP A to the entry point to its neighbor domain ISP B. So, entry point of a contract link represents an interface address on originating (ingress) contract router and exit point represents an interface address on terminating (egress) contract router.

So, for an ISP with N ISP neighbors, the problem of composing contract links boils down to the following steps:

- Calculating total bandwidth capacity between its ingress and egress points
- Determining what portion of this total capacity is to be advertised by an individual contract link announcement
- Attaching a price to contract link which will be determined as a result of offered bandwidth, current g2g path utilization and pre-assessed risk of providing this service (e.g., robustness of links and routers along the g2g path, burstiness of traffic flows usually takes this path)

For the first step, an ISP can employ OSPF-TE capabilities to assess total g2g capacity of its domain. In our implementation, we developed an intra-domain monitoring mechanism which probes core routers of the topology to access shortest-path trees calculated by OSPF protocol on these routers. Matching these path information with link capacity information, we calculated bandwidth capacities for all combinations of $N \times (N-1)$ unidirectional g2g paths. We opt for max-min share of resources on common links which are used by more than one g2g path.

B. Composition of Contract Paths

1) Reservation Signaling and Forwarding: Having global map of contract topology consisting of advertised contract links, each contract router can compute e2e contract paths crossing multiple provider boundaries. As an LSCR-capable service provider, an ISP has to match its user demand with available e2e contract paths composed by concatenating contract links advertised by other LSCR-capable service providers. So, contract path calculations are based on different performance metrics shaped by user demand patterns. Once such a contract path is calculated, LSCR initiates a signaling process similar to RSVP.

When reservations are granted by providers, using tunneling technologies, these forwarding path establishments can be handled without overloading intra-domain routing. ISPs only need to update forwarding state at inter-domain level and then stitch these inter-domain tunnels with already-established intra-domain tunnels.

2) Path Calculation: LSCR neither enforce nor depend on any specific path calculation algorithm. However, to show how a representative path calculation method can be applied on contract links, we choose shortest-widest path algorithm. At this point, we need to differentiate between two types of contract links: 1) *Transit* 2) *Sink*. Transit contract links are the ones which connect two contract routers as was described in the previous section. Sink contract links, however, start and end at the same contract router, and are simply the announcements that advertise which particular prefixes are reachable through which contract routers. A typical sink contract link would have a zero price. Sink contract links also provide flexibility for ISPs to manage their incoming traffic.

In this design, a proposed path calculation method is to first build a map between contract routers which consists of transit contract links, and then populate this map with sink contract links to calculate a forwarding table which supports prefix-based lookups. Such a forwarding table is desired to support multiple entries which possibly represents routes to the same destination calculated by different path calculation methods favoring differing performance metrics (e.g. shortest-widest, cheapest-shortest) and algorithms (e.g. minimal inference).

IV. Evaluation

We performed simulation experiments to analyze contract routing behavior and dynamics as contracting granularity varies. We essentially vary the "contract term" (or the duration of contract) and observe key routing performance measures such as convergence speed and path quality. Varying contract term determines: (i) the duration of time during the network resources are dedicated to contracted links, (ii) the durability of e2e routes established on these contracted links, and (iii) the frequency of changes on virtual contract topology. According to these parameters, messaging load generated by the LSCR protocol, quality and stability of routes calculated, and overall system connectivity are affected. Although LSCR operates at macro time-scales (e.g. several hours, days and longer), here we have included cases at duration of tens-of-minutes to compare the convergence characteristics of LSCR with BGP.

A. Experimental Setup

Our first target is to achieve e2e connectivity through contract link abstractions (e.g. path calculations, address resolutions and forwarding). Furthermore, we want to realistically capture how LSCR interacts with the underlying OSPF-BGP protocols. For this reason we avoid any simplification of the protocol stack and the network topology. However, considering a typical intra-domain topology (e.g. Sprint: 315 nodes), we had to reduce size of our inter-domain topology to 15 ASes to run our experiments. We use BRITE [21] to construct our inter-domain topology. We matched each AS in the inter-domain topology to one of the 6 Rocketfuel [22] intra-domain topologies. In these embedded Rocketfuel topologies, we then classified edge and core routers based on connectivity degrees and distance to center of the domain. After that, we applied a link capacity assignment method which assigns more capacity to those links closer to the center of the domain. We selected 34 edge routers across the 15 ASes and 1000 traffic flows among those edge routers for experimentation. To model user traffic demand, we defined a traffic matrix for the these flows according to the gravity method [23] by considering the population of the cities where the edge routers are located. We repeated the simulation experiments 15 times for each parameter setting. During a simulation run, we took a snapshot of the system every 4 minutes (slightly longer than BGP's hold-down timer).

B. Contract Term vs. QoS

First, we investigated the routing behavior under high stress conditions where we chose high network utilization values so that network can not fully respond to high traffic demand due to limited capacity. We measured the routing performance as the ratio of satisfied traffic demand over total traffic demand, which we call as the "success ratio".

We ran our simulations with various contract terms from extreme 10 minute duration to 24 hours. As seen in Figure 3, for all contract terms, LSCR was able to find better quality e2e paths in comparison to BGP. This can be attributed to single-path constraint of BGP whereas LSCR can exploit alternative paths if available. As the contract term gets longer, LSCR performs higher success ratios as a result of increased stability of established e2e paths. In this comparison, we did not allow BGP to recalculate its paths due to BGP session resets when congested spots emerge in the topology. In real BGP operation, sessions among BGP routers would reset when TCP suffers from high load, which would make the BGP performance even worse than what appears in Figure 3.

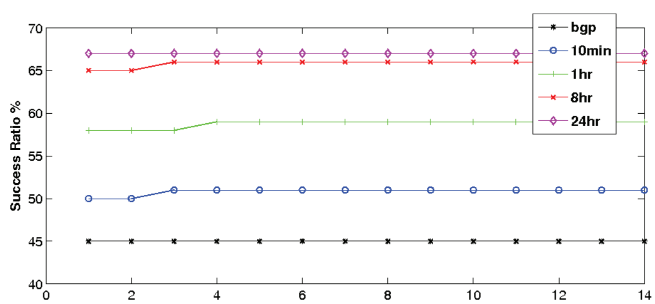

Fig. 3. QoS vs. Reachability Dilemma under High Load: QoS Performance (Supplied Bandwidth over Total Traffic Demand)

C. Contract Term vs. Reachability

We also investigated differences in reachability of destinations with varying contract terms. Considering the number of all destinations, we calculated X, the percentage of unreachable prefixes at each snapshot. Then, we calculated the probability of the event that X percent of the prefixes become unreachable. In Figure 4, we plotted the cumulative distribution function of that probability. As an example, for 10 minutes contract term, 50 percent of the time, observed percentages of unreachable prefixes were 30 percent or less.

As shown in Figure 4, as the contract term gets longer, more IP prefixes become unreachable. Since the network resources are reserved for longer durations by particular contracts, it becomes less possible for others to gain access to these network resources. Figure 4 together with Figure 3 presents the trade-off between e2e path quality and reachability. Increased contract terms cause e2e path quality to increase while reachability to decrease.

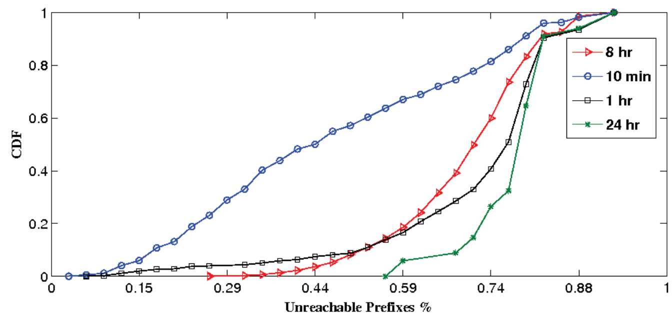

Fig. 4. QoS vs. Reachability Dilemma (High Load): Unreachable Prefixes

D. Messaging Overhead

One of the important factors shaping forwarding dynamics is the number of link-state updates and overall messaging load. Our first observation is that as contract term gets longer, the number of updates decreases since the rate of change in overall system capacity gets slower. Figure 5 shows the number of LSCR update messages normalized to the number of BGP update messages. We observed that, in high load cases, the number of updates could be five times as many as the number of BGP update messages for extreme 10 minutes contract term. For other cases, values are well below that number.

E. Path Stretch

Another key indicator of routing performance is the additional cost of e2e paths in comparison to the shortest path, which is also known as "path stretch". In LSCR, since the cost of a path is not determined solely by the number of hops (e.g. number of ASes), it becomes important to observe how long the contract paths can be in comparison to the BGP-calculated shortest paths. In our case, we investigated the case where LSCR takes the shortest e2e path which provides

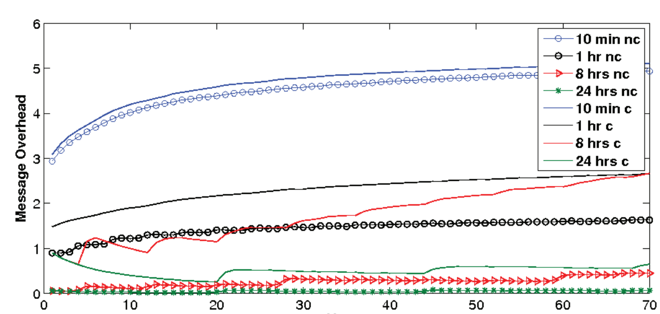

Fig. 5. CR Messaging Overhead: c=high load, nc=moderate load

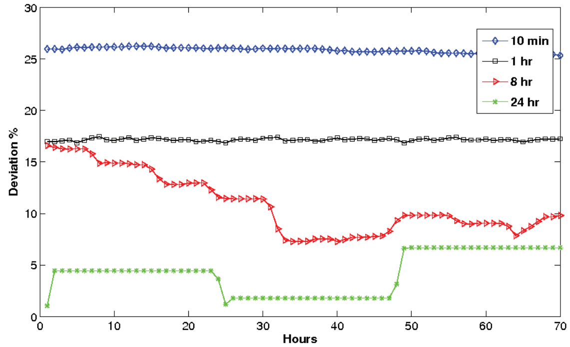

Fig. 6. Path Stretch under Moderate Load

highest quality-of-service in terms of guaranteed bandwidth in return of minimum price. Figure 6 shows that, stretch of LSCR paths in comparison to BGP-calculated paths varies between 25% (for 10 minutes contract term) and 7% (for 1 day contract term). However, this stretch could be limited further using advanced heuristics for contract path calculation.

V. Summary and Further Issues

A. Incremental Deployment

We imagine two-phased deployment transition to the proposed Contract Routing (CR) framework. At its initial phase, we consider a CR-based network core as a secondary market complementing the packet-switched Internet. In this market, existing ISPs operate as CR-enabled providers which offer their leftover bandwidth as g2g contracted services. Since CR supports definition of services on top of the current popular Internet protocols (e.g OSPF and BGP), cost of that minimal transition can be considered as self-sustaining deployment process with the expectation of additional revenue generated by enhanced contracted services. Such secondary market enhanced services will give rise to realization of the potential e2e value existing in the Internet. Later, we expect emergence of a pure CR-enabled infrastructure owners and CR service resellers which only operate through contracts.

B. CR and Network Neutrality

Provisioning of differentiated levels of backbone service quality has been under a heavy debate recently. At one side, keeping the service levels at the backbone "neutral" maintains the innovation at the end systems and enables fast growth of e2e applications. At the other side, blocking the differentiated services at the backbone obsoletes the backbone service provisioning as a profitable business. The Internet performance is not driven only by the end systems or the backbone. We believe that innovation should continue at all parts of the network in order to keep a steady increase in its the overall performance. In this context, our CR framework furthers the competition at the backbone and thus help increasing the innovation and investment at the backbone. Rather than blocking the existing e2e applications, the CR framework enables new set of applications with guaranteed e2e performance. Since the CR framework can be a secondary market, it will not hurt the performance of existing best-effort applications. A clear analogy is the co-existence of express mail and regular mail in the postal services. The CR framework will enable high quality e2e services (like the express mail of postal service) which are not possible due to the rigidity in the existing inter-ISP routing economics.

C. Summary

In this work, we considered a framework where routing over "contract" abstraction is realized. Contracts are defined as an embedding of various components such as technical performance, economics and temporal flexibilites. These flexibilities are essentially incorporated into the inter-domain routing protocol, which we called "contract routing". We built the Contract Routing framework which can house multiple protocols using contracts as routing units. More specifically, we showed that it is possible to implement LSCR as an overlay on top of popular routing protocols BGP and OSPF. Our future work will include experimentation of path-vector style contract path compositions.

References

[1] C. Labovitz, G. Malan, and F. Jahanian, "Internet routing instability," *IEEE Trans. on Network.*, vol. 6, pp. 515–528, 1998.
[2] L. G. Roberts, "A radical new router," *IEEE Spectrum*, vol. 46, pp. 34–39, 2009.
[3] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "Bgp routing dynamics revisited," *ACM CCR*, vol. 37, no. 2, pp. 5–16, 2007.
[4] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A.Berger, and W. Lehr, "Complexity of internet interconnections: Technology, incentives and implications for policy," in *Proc. of 35th TPRC*, 2007.
[5] M. Yuksel, A. Gupta, and S. Kalyanaraman, "Contract-switching paradigm for internet value flows and risk management," in *Proc. of IEEE GI Symposium*, 2008.
[6] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: Defining tomorrows internet," *IEEE/ACM Trans. on Network.*, vol. 13, no. 3, 2005.
[7] R. Mahajan, D. Wetherall, and T. Anderson, "Negotiation-based routing between neighboring isps," in *Proc. of USENIX NSDI*, 2005.
[8] Y. Shavitt and Y. Singer, "Trading potatoes in distributed multi-tier routing systems," in *Proc. of the NetEcon*, 2008, pp. 67–72.
[9] X. Yang, D. Clark, and A. Berger, "Nira: A new inter-domain routing architecture," *IEEE Trans. on Network.*, vol. 15, pp. 775–788, 2007.
[10] 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. of SIGCOMM*, 2005, pp. 13–24.
[11] W. Xu and J. Rexford, "Miro: multi-path interdomain routing," *ACM CCR*, vol. 36, no. 4, pp. 171–182, 2006.
[12] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "Characterizing and measuring path diversity of internet topologies," in *Proc. of the SIGMETRICS*, 2003, pp. 304–305.
[13] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The end-to-end effects of internet path selection," in *Proc. of SIGCOMM*, 1999, pp. 289–299.
[14] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *Proc. of SIGCOMM*, 2009, pp. 111–122.
[15] N. Feamster, R. Johari, and V. Vazirani, "Agora: A market for internet connectivity," in *PRESTO*, 2007.
[16] Y. Zhu, R. Zhang-Shen, S. Rangarajan, and J. Rexford, "Cabernet: connectivity architecture for better network services," in *Proc. of CoNEXT*, 2008, pp. 1–6.
[17] W. Liu, H. T. Karaoglu, A. Gupta, M. Yuksel, and K. Kar, "Edge-to-edge bailout forward contracts for single-domain internet services," in *Proc. of IEEE IWQoS*, 2008.
[18] D. G. Andersen, N. Feamster, S. Bauer, and H. Balakrishnan, "Topology inference from bgp routing dynamics," in *Proc. of the ACM IMC*, 2002.
[19] J. Moy, "OSPF version 2," *IETF RFC 2328*, April 1998.
[20] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "Bgp routing stability of popular destinations," in *Proc. of ACM IMW*, 2002, pp. 197–202.
[21] A. Medina, A. Lakhina, I. Matta, and J. Byers, "Brite: An approach to universal topology generation," in *Proc. of the MASCOTS*, 2001, p. 346.
[22] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *Proc. of SIGCOMM*, 2002.
[23] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: Existing techniques and new directions," in *Proc. of SIGCOMM*, 2002.