

Contract-Switching Paradigm for Internet Value Flows and Risk Management

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

Aparna Gupta
Rensselaer Polytechnic Institute
Troy, NY 12180
guptaa@rpi.edu

Shivkumar Kalyanaraman
IBM Research India
Bangalore, India
shivkumar-k@in.ibm.com

Abstract—The Internet’s simple design resulted in huge success in basic telecommunication services. However, in terms of providing end-to-end QoS services, the Internet’s architecture needs major shifts since it neither allows (i) users to indicate their *value* choices at sufficient granularity nor (ii) providers to manage *risks* involved in investment for new innovative QoS technologies and business relationships with other providers as well as users. Currently, users can only indicate their value choices at the access/link bandwidth level not at the routing level. Similarly, an enterprise that needs end-to-end capacity contracts between two arbitrary points on the Internet for a short period of time has no way of expressing its needs. To allow these much needed economic flexibilities, we introduce *contract-switching* as a new paradigm for the design of future Internet architecture. Just like packet-switching enabled flexible and efficient multiplexing of data, a contract-switched inter-network will enable flexible and economically efficient management of risks and value flows with many more tussle points.

Index Terms—Network Economics, Inter-domain Routing, Contract Routing

I. INTRODUCTION

The Internet’s simple best-effort packet-switched architecture lies at the core of its tremendous success and impact. Today, the Internet is firmly a commercial medium involving several competitive service providers and content providers. However, current Internet architecture neither allows (i) users to indicate their *value* choices at sufficient granularity nor (ii) providers to manage *risks* involved in investment for new innovative QoS technologies and business relationships with other providers as well as users. Currently, users can only indicate their value choices at the access/link bandwidth level not at the routing level. End-to-end QoS contracts are possible today via virtual private networks, but with static and long-term contracts. Further, an enterprise that needs end-to-end capacity contracts between two arbitrary points on the Internet for a short period of time has no way of expressing its needs.

We propose an Internet architecture that allows flexible, finer grained, 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. A contract-switched architecture will enable flexible and economically efficient management of risks and value flows in an Internet characterized by many tussle points [1] We view “contract-switching” as a generalization of the packet-switching paradigm of the current

Internet architecture. For example, *size of a packet* can be considered as a special case of the *capacity of a contract* to expire at a very short-term, e.g. transmission time of a packet. Similarly, *time-to-live* of packet-switching is roughly a special case of the *contract expiration* in contract-switching. Thus, contract-switching is a more general case of packet-switching with several additional flexibilities in terms of its economics and carefully reduced technical flexibilities due to scalability concerns particularly at the routing level.

This paper focuses on research issues behind creating a contract-switching network architecture for flexible value flows for future Internet, and for allowing sophisticated financial engineering tools to be employed in managing the risks involved in composition of end-to-end QoS contracts. We concentrate on the design of our contract-switching framework in the context of multi-domain QoS contracts. Our architecture allows such contracts to be dynamically composable across space (i.e., across ISPs) and time (i.e., over longer time-scales) in a fully decentralized manner. Once such elementary instruments are available and a method for determining their value is created (e.g., using secondary financial markets), ISPs can employ advanced pricing techniques for cost recovery and financial engineering techniques to manage risks in establishment of end-to-end contracts and performance guarantees for providers and users in specific market structures, e.g., oligopoly or monopoly. We build on top of our edge-based distributed dynamic capacity contracting (DDCC) framework [2], which was proposed for a single domain. As DDCC can operate over ISP peering points, we employ contracts involving these ISP peering points and illustrate ways of realizing a contract-switched Internet core.

In particular, we investigate elementary QoS contracts and service abstractions at micro (i.e., *tens-of-minutes*) or macro (i.e., *hours or days*) time-scales. For macro-level operation at high time-scales (i.e., *several hours or days*, potentially involving contracts among ISPs and end users), we envision a link-state like structure for computing end-to-end “contract routes.” Similarly, to achieve micro-level operation with more flexibilities at lower time-scales (i.e., *tens-of-minutes*, mainly involving contracts among ISPs), we envision a BGP-style path-vector contract routing. Though there are similarities to QoS routing, the composition of contracts can involve multiple attributes, involve derivative contracts, and are based




Fig. 1. Packet-switching introduced many more tussle points into the Internet architecture by breaking the *end-to-end circuits* of circuit-switching into *routeable datagrams*. Contract-switching introduces even more tussle points at the edge/peering points of domain boundaries by *overlay contracts*.

upon “contract-link-states” and “contract-path-vectors.” We illustrate some examples of such decentralized contract composition.

Several QoS mechanisms have been adopted within single ISP domains, while inter-domain QoS deployment has not become reality. Arguably the reasons for this include the highly fragmented nature of the ISP market and glut in core optical capacity due to overinvestment and technological progress of the late 1990s. BGP routing convergence and routing instability issues [3] also contribute to inter-domain performance uncertainties. Recent QoS research [4], [5] clearly identified a lack of inter-domain business models and financial settlement methods, and a need for flexible risk management mechanisms including insurance, money-back-guarantees. Specifically, attempts to incorporate economic instruments into inter-domain routing [6], [7] and to allow more economic inter-domain flexibilities to end-users [8] have surfaced. Our work directly focusses on these issues, and also relates to the Internet pricing research [2], [9], [10], [11].

In Section II, we first define the essence of contract-switching paradigm. Section III details architectural characteristics and challenges of contract-switching. Then, in Section IV, we provide the motivations for using financial engineering for risk management in the Internet. We summarize our work in Section VI.

II. CONTRACT-SWITCHING PARADIGM

The essence of “contract-switching” is to use *contracts* as the key building block for inter-domain networking. As shown in Figure 1, this increases the inter-domain architecture flexibilities by introducing more tussle points into the protocol design. Especially, this paradigm will allow the much needed revolutions in the Internet protocol design: (i) inclusion of economic tools in the network layer functions such as inter-domain routing while the current architecture only allows basic connectivity information exchange, and (ii) management of risks involved in QoS technology investments and participation into e2e QoS contract offerings by allowing ISPs to potentially apply financial engineering methods.

In addition to these design opportunities, the contract-switching paradigm introduces several research challenges. As the key building block, intra-domain service abstrac-

tions call for design of (i) single-domain edge-to-edge QoS contracts with performance guarantees and (ii) nonlinear pricing schemes geared towards cost recovery. Moving one level up, composition of end-to-end inter-domain contracts poses a major research problem which we formulate as a “contract routing” problem by using single-domain contracts as “contract links”. Issues to be addressed include routing scalability, contract monitoring and verification as the inter-domain context involves large-size effects and crossing trust boundaries. Several economic tools can be used to remedy pricing, risk sharing, and money-back problems of a contract-switched network provider (CSNP), which can operate as an overlay re-seller ISP (or an alliance of ISPs) that buys contract links and sells e2e QoS contracts. In addition to CSNPs, the contract-switching paradigm allows more distributed ways of composing e2e QoS contracts as we will detail later.

III. ARCHITECTURAL ISSUES

Inclusion of concepts such as values and risks into the design of network routing protocols poses various architectural research challenges. We propose to abstract the point-to-point QoS services provided by each ISP as a set of “overlay contracts” each being defined between peering points, i.e., ingress/egress points. By considering these overlay contracts as “contract links”, we envision a decentralized framework that composes an end-to-end (e2e) contract from contracts at each ISP hop. This is similar to path made up of links, except that we have “contracts” instead of links. Just like routing protocols are required for actually creating e2e paths from links, we need “contract routing” to find an inter-domain route that concatenates per-ISP contracts to compose an e2e path and an associated e2e contract bundle.

We define a contract as an embedding of three major flexibilities in addition to the contracting entities/ISPs: (i) *performance component*, (ii) *financial component*, and (iii) *time component*. The performance component of a contract can include QoS metrics such as delay or packet loss to be achieved. The financial component of contract will include various fields to aid entities in making financial decisions related to value and risk tradeoffs involved in engaging in the contract. The basic fields can be various prices, e.g., spot, forward, and usage-based. It is possible to design interesting financial components 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 term contract will expire and the time left for the insured term when the money-back guarantee will expire.

A. Dynamic Contracting over Peering Points

The first step towards contract-switching is the realization of necessary building blocks so that an ISP can compose an *advertisable* contract involving its *intra-domain* resources. A practical way of achieving this for an ISP is to compose




Fig. 2. Dynamic Capacity Contracting (DCC) framework. [2]




Fig. 3. An illustrative scenario for link-state contract routing.

contracts between peering points. Previous work showed that this kind of “edge-to-edge” dynamic contracting can be done in a distributed manner with low costs [2]. Figure 2 illustrates the big picture of such a distributed framework. Customers can only access network core by making contracts with the provider stations placed at the edge points. A key capability is that an ISP can advertise different prices for each edge-to-edge contract it offers. Therefore, we abstract the point-to-point QoS services provided by each ISP as a set of “overlay contract links” each being defined between peering points, i.e., ingress/egress points. ISPs would in practice choose to have only a portion of their network capacity available for such markets of flexible contracting. This distributed contracting architecture gives more flexibilities to users in choosing their e2e paths.

B. Contract Routing

Given contracts between peering points of ISPs, the “contract-routing” problem involves discovering and composing end-to-end QoS-enabled contracts from per-ISP contracts. We consider each potential contract to be advertised as a “contract link” which can be used for end-to-end routing. Over such contract links, it is possible to design link-state or BGP-style path-vector routing protocols that compose end-to-end “contract paths”. CSNPs and ISPs providing end-to-end services will need to be proactive in their *long-term* contracting and financial commitments and thus will need *link-state style* routing satisfying this need. Such proactive link-state routing components will be instrumental in attaining the flexibility of derivative contracts like forwards and options. On the other end, it is possible to achieve scalable *BGP-style path-vector* routing at *short time-scales* where frequent changes happen. Such short-term contract routing will be

instrumental for realizing dynamic pricing (e.g., congestion-based) and financial engineering (e.g., spot prices) capabilities.

1) *Macro-level Operations: Link-State Contract Routing:* One version of inter-domain contract routing is link-state style with long-term (i.e., *hours or days*) contract links. For each contract link, the ISP creates a “contract-link-state” including various fields. We suggest that the major field of a contract-link-state is the forward prices (or prices committed for a later deal) in the future as predicted by the ISP now (based upon anticipated future loads). Such contract-link states are flooded to other ISPs and CSNPs. Each ISP will be responsible for its flooded contract-link-state and therefore will have to be *proactively* measuring validity of its contract-link-state. This is very similar to the periodic HELLO exchanges among the routers in an OSPF domain. When remote ISPs obtain the flooded contract-link-states, they can offer point-to-point and end-to-end contracts that may cross multiple peering points. Though link-state routing was proposed in an inter-domain context [12], our “contract links” are between peering points *within* an ISP, and not between ISPs (see Figures 2 and 3).

To compute the end-to-end “contract paths”, the local agent of CSNPs or ISPs performs a QoS-routing like computation procedure to come up with source routes, and initiates a signaling protocol to reserve these contracts. Figure 3 shows a sample scenario where link-state contract routing takes place. There are three ISPs participating with 5 peering points. For the sake of example, a contract-link-state includes six fields: *Owner ISP*, *Link*, *Term* (i.e., the length of the offered contract link), *Offered After* (i.e., when the contract link will be available for use), and *Price* (i.e., the aggregate price of the contract link including the whole term). ISPs have the option of advertising by flooding their contract-link-states among their peering points. Each ISP has to maintain a contract-link-state routing table as shown in the figure. Some of the contract-link-states will diminish by time, e.g., the link 1-3 offered by ISP A will be omitted from contract routing tables after 5hrs and 15mins. Given such a contract routing table, computation of “shortest” QoS contracts involves various financial and technical decisions. Let’s say that the user X (which can be another ISP, CSNP, or a network entity having an immediate access to the peering point 1 of ISP A) wants to purchase a QoS contract from 1 to 5. The CSNP can offer various “shortest” QoS contracts. For example, the route 1-2-4-5 is the most cost-efficient contract path (i.e. $(10\text{Mb/s} \times 2\text{hrs} + 100\text{Mb/s} \times 3\text{hrs} + 60\text{Mb/s} \times 24)/(\$10 + \$110 + \$250) = 27.2\text{Mb/s} \cdot \text{hr}/\$$), while the 1-3-5 route is better in terms of QoS. ISPs can factor in their financial goals when calculating these “shortest” QoS contract paths. The 1-2-4-5 route gives a maximum of 10Mb/s QoS offering capability from 1 to 5, and thus the CSNP/ISP will have to sell the other purchased contracts as part of other end-to-end contracts or negotiate with each individual contract link owner. Similarly, the user X tries to maximize its goals by selecting one of the offered QoS contracts to purchase from 1 to 5. Let’s say that the CSNP/ISP offers user X two options as: (i) using the route 1-2-4-5 with 10Mb/s capacity, 2hrs term, starting in 5hrs with a price \$15 and (ii) using the route 1-3-

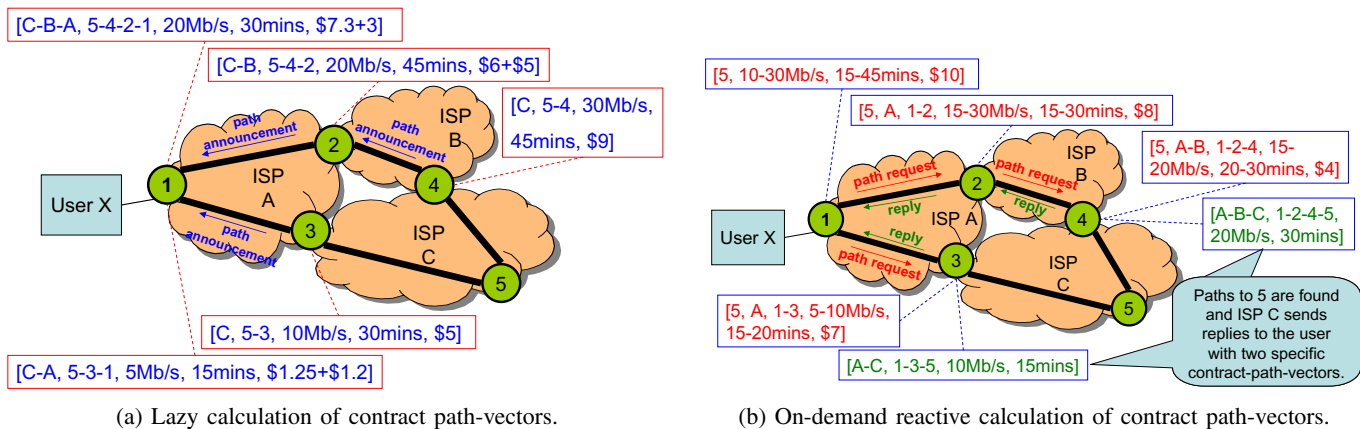


Fig. 4. Path-vector contract routing: (a) *Provider initiates* contract path-vector calculation. (b) *User initiates* contract path-vector calculation.

5 with 20Mb/s capacity, 1hr term, starting in 30mins with a price \$6. Let's say that user X selects the 1-3-5 route. Then, the CSNP/ISP starts a signaling protocol to reserve the 1-3 and the 3-5 contract links, and triggers the flooding of contract link updates indicating the changes in the contract routing tables.

One issue that will arise if an ISP participates in many peering points is the explosion in the number of "contract links", which will trigger more flooding messages into the link-state routing. But, the number of contract links can be controlled by various scaling techniques, such as focussing only on the longer-term contracts offered between the major peering points and aggregating contract-link-states as region-to-region where a region corresponds to a set of peering points. Also, a key difference between our proposed link-state contract routing and the traditional intra-domain link-state routing is that *floods only need to be performed if there is a significant change on contracting terms or in the internal ISP network conditions*. However, in traditional link-state routing, link-states are flooded periodically regardless if any change has happened.

2) *Micro-level Operations: Path-Vector Contract Routing:*

To provide enough flexibility capturing more dynamic technical and economical behaviors in the network, it is possible to design contract routing that operates at short time-scales, i.e., *tens of minutes*. This time-scale is reasonable as current inter-domain BGP routing operates with prefix changes and route updates occurring at the order of a few minutes [13]. Further, an ISP might want to advertise a spot price for an edge-to-edge contract to a subset of other ISPs and CSNPs instead of flooding it to all. Similarly, a user might want to query a specific contracting capability for short-term and involving various policy factors. Such *on-demand reactive* requests cannot be adequately addressed by the link-state contract routing.

Just like BGP composes paths, e2e contract paths can be calculated in an on-demand lazy manner. In our design, each ISP has the option of initiating contract path calculations by advertising its contract links to its neighbors. Depending on various technical, financial, or policy factors, those neighbors

may or may not use these contracts in composing a two-hop contract path. If they do, then they advertise a two-hop contract path to their neighbors. This path-vector composition process continues as long as there are participating ISPs into the contract paths. Users or ISPs receiving these contract paths will have the choice of using them or leaving them to invalidation by the time the contract path term expires.

Provider Initiates: Figure 4(a) shows an example scenario where a provider initiates contract-path-vector calculation. ISP C announces two short term contract-path-vectors at peering points 3 and 4. The ISPs B and A decides whether or not to participate in these contract-path-vectors, possibly with additional constraints. For example, ISP B reduces the capacity of the initial path-vector to 20Mb/s and increases its price to \$11. Though each ISP can apply various price calculations, in this example ISP B adds \$5 for its own contract link 2-4 on top of the price of the corresponding portion (i.e., $\$9 \times 20/30 = \6) of the contract link 4-5 coming from ISP C. Similarly, ISP A constrains the two contract-path-vector announcements from ISPs B and C at peering points 2 and 3 respectively. Then, the CSNP (or ISP A) offers the two contract-path-vectors to the user X, who may choose to use the 1-5 short-term QoS path. In this path-vector computation scheme, whenever an ISP participates in a contract it will have to commit the resources needed for it, so that the users receiving the contract path announcements will have assurance in the end-to-end contract. Therefore, ISPs will have to decide carefully as potential security and trust issues will play a role. This game theoretic design exists in the current BGP inter-domain routing. In BGP, each ISP decides which route announcements to accept for composing its routes depending on policy, trust, and technical performance.

User Initiates: Users may query for an e2e short-term contract path with specific QoS parameters which do not exist in the currently available path-vector. This kind of design can potentially allow involvement of end users into the process depending on the application-specific needs. For example, in Figure 4(b), user X initiates a path-vector calculation by broadcasting a "contract-path request" to destination 5 with a

capacity range 10-30Mb/s, term range 15-45mins with up to \$10 of total cost. This contract-path request gets forwarded along the peering points where participating ISPs add more constraints to the initial constraints identified by the user X. For example, ISP B narrows the term range from 15-30mins to 20-30mins and the capacity range from 15-30Mb/s to 15-20Mb/s while deducting \$4 for the 2-4 contract link of its own from the leftover budget of \$8. Such participating middle ISPs have to apply various strategies in identifying the way they participate in these path-vector calculations. Once ISP C receives the contract-path requests, it sends a reply back to user X with specific contract-path-vectors. The user X then may choose to buy these contracts from 1 to 5 and necessary reservations will be done through more signaling.

IV. SPOT PRICING AND RISK MANAGEMENT

The first step in financial engineering is to identify a well-defined contract and arrive at a spot price for it. Consider a single-domain capacity contract which offers a guaranteed, fixed mean effective service rate. The greatest risk in viable provisioning is the ability to recover cost of providing the service. Therefore, nonlinear pricing models designed for cost recovery to arrive at a spot price $S(0, T)$ for such elementary capacity contracts with duration T are shown to perform well [14]. Pricing of contracts that go beyond a mean service rate guarantee will involve an additional price component based upon financial derivatives pricing.

Concepts of financial derivatives are applied to several areas including telecom network services [15]. The basic goal of derivative pricing is to assign a price to a risk derived from an “underlying” source of risk. For instance, an option on stocks has risk derived from volatility in stock value. The principle behind derivative pricing is to assign a price for each future elementary outcome of the underlying risk. This is known as the price of a **contingent claim** [16]. For instance, if $s \in \Omega$ are all possible future outcomes of the underlying risk, then Ψ_s is the price of a contingent claim, which pays-off \$1 *only* if outcome s is realized, and \$0 otherwise. The underlying risk could be a traded asset or a fundamental economic factor [17], such as, interest rates, exchange rates, etc. For pricing QoS guaranteed contracts, **capacity, delay, loss** act as the **fundamental factors**.

Derivatives derive their risk from underlying risks since they provide pay-offs in a subset of elementary outcomes, i.e., a derivative would pay-off Y_s only for $s \in \Phi$ for some $\Phi \subset \Omega$. Then the price of this derivative is obtained by a linear combination of prices of contingent claims corresponding to outcomes in Φ , i.e., $\sum_{s \in \Phi} Y_s \Psi_s$. In the continuous models, discrete contingent claims are replaced by constructing a **state price density (SPD)** [18], p_s^* , which assigns a price to every future outcome, $s \in \Omega$. If the derivative yields Y_s for $s \in \Phi$, ($\Phi \subset \Omega$), then price of the derivative is obtained as: $\int_{s \in \Phi} Y_s p_s^* ds$, which is the expectation over a new measure called the risk-neutral probability P^* .

The definition of the QoS guarantee in contracts resembles certain path-dependent options. For instance, the loss process




Fig. 5. Price variations for different State Price Densities (SPDs)

averaging per unit time exceeding an upper barrier with a certain frequency will violate a guarantee or, in financial terms, knock the “option” out. Appropriate SPDs for loss/delay outcomes can be defined to evaluate the monetary “reward” for the favorable risks to the provider. Risk sharing can be achieved by assigning a zero or negative “reward” to the unfavorable outcomes as described by the QoS specifications. For instance, to price the risks in the loss processes, an SPD will capture a representative provider’s preferences for the future loss outcomes. If Y_t defines the loss guarantee at time t , the “options” price of the loss guarantee is given by

$$C(0, T) = E_{P^*} \left[\int_0^T Y_t dt \right]. \quad (1)$$

Therefore, the SPD translates into a new risk-neutral probability measure P^* . Y_t may take different forms depending on how the contract is defined. SPDs displaying different properties result in distinct price structures. An SPD may reward the zero-loss scenario and decay with larger levels of losses, i.e, exponential distribution. Alternately, under the assumption that customers may be insensitive to very small data losses up to a certain level, the provider may be able to accommodate more customers by allowing small losses to an individual customer’s data. In this case, the SPD does not reward zero-losses, but instead peaks at some small loss-level and then decays for larger loss levels, depicted by Beta distribution. These two distinct choice of SPD types result in either congestion sensitive or performance sensitive prices for loss guarantees [14] (see Figure 5). For the exponential SPDs the prices peak at the low-utilization periods of wee-hours in the night (performance sensitive) (see inset in Figure 5), while the Beta SPD prices peak at the peak of the day when the network is most utilized (congestion sensitive).

In summary, **the price of a QoS guaranteed contract** is determined as the sum of price of a capacity guarantee, obtained by the nonlinear pricing method, and appropriate options-based pricing for additional QoS guarantees provided, i.e. $S(0, T) + C(0, T)$. Once this basic contract is available, financial engineering methods can be used to bundle these contracts to compose end-to-end contracts (i.e. over space)

and/or construct forward contracts for longer terms (i.e. over time) services.

V. DEPLOYMENT ISSUES

The exact steps towards a full deployment of contract-switching need to be explored and debated. Traditionally ISPs have been very proprietary about the contracts that they offer and the notion that they would make these available to all their peers would require a significant turn in mindset. Presumably contract prices could be used to infer an ISP's cost structure and its load patterns, both of which offer huge value to the sales force of a competing provider. Thus, contract-switching raises concerns of willingness to participate in such a framework where one ISP is asked to "flood" contract terms. A practical similar example is the operations of BGP. ISPs expose their BGP connectivity information and willingness to carry traffic in order to stay connected with the rest of the Internet. BGP announcements carry a similar risk of some internal information to be inferred, e.g., topology and load patterns. However, this risk does not stop ISPs to flood BGP announcements. Our contract announcements involve similar game-theoretic issues. ISPs will gain additional rewards like e2e QoS provisioning in return of undertaking additional risks of exposing financial terms on top of the current BGP's baseline connectivity information.

We believe that the best-effort service paradigm and e2e QoS services can co-exist initially, where ISPs only incrementally participate in contract-switching with their leftover bandwidth. Given that there is a capacity glut, ISPs will not have much to lose in participating in such a new market to create new values of e2e QoS. Once users experience performance benefit of e2e QoS services (even with leftover bandwidth), ISPs will then start to see that e2e QoS provisioning via schemes like contract routing returns more than single-hop SLAs. Eventually, this will increase the amount of participation by ISPs. Though the actual stages of deployment might turn out to be different, the critical issue is to reveal the potential values of contract-switching.

Another key deployment issue is monitoring and verification of ongoing e2e contracts. Under-performing ISPs involved in an e2e QoS contract should be penalized according to these monitoring results and the terms of the contract. For long-term contracts, involvement of a neutral third party entity is not a costly choice in terms of messaging scalability. Similar to existing authentication protocols, entities engaging in a contract may choose to have a third party service for contract verification with additional costs. Verification of such crucial long-term contracts requires active involvement of software/hardware agents. Contracts with money-back guarantees and significant costs deserve such support. To address the verification of short time-scale contracts, we suggest to use e2e, mainly passive measurement techniques to monitor contracts.

VI. SUMMARY

The current Internet architecture needs more flexibilities in realizing value flows and managing risks involved in inter-domain relationships. To enable such flexibilities, we outlined the contract-switching paradigm that promotes using contracts overlaid on packet-switching intra-domain networks. In comparison to packet-switching, contract-switching introduces more tussle points into the architecture. By routing over contracts, we showed that economic flexibilities can be embedded into the inter-domain routing protocol designs and this framework can be used to compose end-to-end QoS-enabled contract paths. Within such a "contract routing" framework, we also showed that financial engineering techniques (e.g., options pricing) can be used to manage risks involved in inter-domain business relationships.

ACKNOWLEDGMENT

This work is supported in part by National Science Foundation awards 0721600, 0721609, and 0627039.

REFERENCES

- [1] D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: Defining tomorrow's internet," in *Proc. of SIGCOMM*, August 2002.
- [2] M. Yuksel and S. Kalyanaraman, "Distributed Dynamic Capacity Contracting: An overlay congestion pricing framework," *Computer Communications*, vol. 26, pp. 1484–1503, 2003.
- [3] T. G. Griffin and G. Wilfong, "An analysis of BGP convergence properties," *ACM SIGCOMM CCR*, vol. 29, no. 4, 1999.
- [4] B. Teitelbaum and S. Shalunov, "What QoS research hasn't understood about risk," *Proc. of ACM SIGCOMM Workshops*, pp. 148–150, 2003.
- [5] M. Gaynor, S. Bradner, M. Iansiti, and H. T. Kung, "The real options approach to standards for building network-based services," *Proc. of IEEE SIIT*, Oct 2001.
- [6] Mike Afergan, "Using repeated games to design incentive-based routing systems," in *In Proceedings of IEEE INFOCOM*, 2006.
- [7] R. Mahajan, D. Wetherall, and T. Anderson, "Negotiation-Based Routing Between Neighboring ISPs," in *Proc. of USENIX NSDI*, 2005.
- [8] Xiaowei Yang and David Clark and Arthur Berger, "NIRA: A New Inter-Domain Routing Architecture," *IEEE/ACM Transactions on Networking (ToN)*, vol. 15, no. 4, pp. 775–788, August 2007.
- [9] I. Ch. Paschalidis and J.N. Tsitsiklis, "Congestion-dependent pricing of network services," *IEEE/ACM Tran. on Networking*, vol. 8, no. 2, pp. 171–184, 2000.
- [10] F.P. Kelly, A.K. Maulloo, and D.K.H. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [11] A.M. Odlyzko, "Internet pricing and history of communications," AT&T labs, 2000.
- [12] I. Castineyra, N. Chiappa, and M. Steenstrup, "The nimrod routing architecture," *IETF RFC 1992*, August 1996.
- [13] R. Teixeira, S. Agarwal, and J. Rexford, "BGP routing changes: Merging views from two ISPs," *ACM SIGCOMM CCR*, 2005.
- [14] A. Gupta, S. Kalyanaraman, and L. Zhang, "Pricing of risk for loss guaranteed intra-domain Internet service contracts," *Computer Networks*, vol. 50, pp. 2787–2804, 2006.
- [15] G. Cheliotis, *Structure and Dynamics of Bandwidth Markets*, Ph.D. thesis, N.T.U. Athens, 2001.
- [16] C. Huang and R. H. Litzenberger, *Foundations for Financial Economics*, Prentice Hall, Inc., New Jersey, 1988.
- [17] V. Gorovoi and V. Linetsky, "Black's model of interest rates as options, eigenfunction expansions and Japanese interest rates," *Mathematical Finance*, vol. 14, no. 1, pp. 49–78, Jan 2004.
- [18] M. Chernov, "Empirical reverse engineering of the pricing kernel," *Journal of Econometrics*, vol. 116, pp. 329–364, 2003.