Volume 36, Issues 17–18, November/December 2013 ISSN 0140-3664

ELSEVIER

computer communications

Computer Communications 36 (2013) 1708–1725

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Bailout forward contracts for edge-to-edge internet services ☆

Hasan T. Karaoglu [a,1], Aparna Gupta [b], Murat Yuksel [a,*], Weini Liu [b,2], Koushik Kar [c]

[a] Computer Science and Engineering, University of Nevada - Reno, Reno, NV 89557, USA
[b] Lally School of Management and Technology at Rensselaer Polytechnic Institute, Troy, NY 12180, USA
[c] Department of Electrical, Computer and Systems Engineering at Rensselaer Polytechnic Institute, Troy, NY 12180, USA

ARTICLE INFO

Article history:
Received 27 October 2012
Received in revised form 28 May 2013
Accepted 3 August 2013
Available online 23 August 2013

Keywords:
Routing economics
Service level agreements
Quality of service
Intra-domain routing

ABSTRACT

Despite the huge success of the Internet in providing basic communication services, its economic architecture needs to be upgraded so as to provide end-to-end guaranteed or more reliable services to its customers. Currently, a user or an enterprise that needs end-to-end bandwidth guarantees between two arbitrary points in the Internet for a short period of time has no way of expressing its needs. To allow these much needed basic services, we propose a single-domain edge-to-edge (g2g) dynamic capacity contracting mechanism, where a network customer can enter into a bandwidth contract on a g2g path at a future time, at a predetermined price. For practical and economic viability, such forward contracts must involve a bailout option to account for bandwidth becoming unavailable at service delivery time, and must be priced appropriately to enable Internet Service Providers (ISPs) manage *risks* in their contracting and investments. Our design allows ISPs to advertise point-to-point different prices for each of their g2g paths instead of the current point-to-anywhere prices, allowing discovery of better end-to-end paths, temporal flexibility and efficiency of bandwidth usage. We compute the *risk-neutral prices* for these g2g *bailout forward contracts* (BFCs), taking into account correlations between different contracts due to correlated demand patterns and overlapping paths. We apply this multiple g2g BFC framework on network models with Rocketfuel topologies. We evaluate our contracting mechanism in terms of key network performance metrics like fraction of bailouts, revenue earned by the provider, and adaptability to link failures. We also explore the tradeoffs between complexity of pricing and performance benefits of our BFC mechanism.

1. Introduction

The Internet is a commercial environment embodying multiple service providers competing with each other. Provisioning end-to-end (e2e) quality-of-service (QoS), thus relies on the viability and flexibility of single-domain edge-to-edge (g2g) contracting capabilities. Current single-domain contracts or service level agreements (SLAs) are typically *point-to-anywhere* settlements happening in peer-to-peer or customer-provider ISP relationships. This point-to-anywhere nature of SLAs carry all the way to the end users, and thus the current Internet services are packaged in a typically flat-rate and point-to-anywhere deals without any specific end-to-end performance guarantees, except the access bandwidth guarantees. Though such best-effort point-to-anywhere contracting has the convenience of making the customer not worry about per-destination prices (i.e., different prices for the traffic destined to different locations instead of a single price for all possible destinations), the tradeoffs are (i) lack of e2e QoS and (ii) the lost opportunity for discovering potentially better value flow paths both economically (e.g., cheaper) and technically (e.g., higher capacity).

Another key characteristic missing in the current SLAs is the economic flexibility to manage risks involved in the inter-ISP settlements. For example, the time-scale of SLAs is too long (e.g., months to years) and there is typically no way of bailing out of an SLA if the ISP finds a better deal. Further, SLAs are arranged for immediate service (or in the very near future such as a few days/weeks) and an ISP typically cannot easily close deals for its future investments to reduce risks involved in its investment. It is a pressing need to have such economic instruments for enabling the ISPs to *manage risks* in their investments.

We consider an Internet architecture that allows flexible, finer grained, dynamic contracting over multiple providers. We propose a new family of single-domain contracting mechanisms based on edge-to-edge (g2g) dynamic capacity contracting [2] involving

☆ An initial version of this work appeared in IEEE IWQoS [1]. On top of this conference publication, we have extended the paper with a large amount of work. Our extensions are mainly in Section 7 and include (i) revenue analysis of bailout forward contracts (BFCs) concept, (ii) exploration of simplifying the pricing complexity of BFCs, and (iii) the potential benefit of predictable future demand on the robustness of BFCs.

* Corresponding author. Tel.: +1 775 3272246; fax: +1 775 7841877.
E-mail addresses: karaoglu@cse.unr.edu (H.T. Karaoglu), guptaa@rpi.edu (A. Gupta), yuksem@cse.unr.edu (M. Yuksel), winnieliuzzz@gmail.com (W. Liu), koushik@ecse.rpi.edu (K. Kar).
[1] Dr. Karaoglu is now with Cisco Systems, Inc.
[2] Weini Liu is now with Deutsche Bank.

0140-3664/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comcom.2013.08.004

forward contracts with bailout options, also called *g2g bailout forward contracts* (BFCs). Our design allows ISPs to advertise point-to-point different prices for each of their g2g paths instead of the current point-to-anywhere prices. Such g2g contracts enable composition of end-to-end higher quality paths given that inter-domain relationships are made over such g2g contracting [3]. Breaking the point-to-anywhere contracts into point-to-point g2g contracts allows more tussle points [4] (between multiple network service providers and content providers) in the system and thus opens the door for discovering better end-to-end paths [5]. This phenomenon is also illustrated in Fig. 1, where end-to-end QoS paths can be composed by concatenating single-domain g2g contracts.

A forward contract for a g2g path, as the name suggests, offers a service on that g2g path which will be delivered at a future time, but at a predetermined price, called the "forward price". We enhance this forward contract with a bailout clause and establish a bailout forward contract (BFC). BFCs allow the provider to bail out from offering the service at a future time, if the available capacity or resources on the g2g path is not sufficient to support the service. Offering such g2g BFCs on all the chosen g2g paths in a domain increases the *spatial tussle* by enabling point-to-point economics rather than the current pure point-to-anywhere approach, and this provides mechanisms for more efficient use of bandwidth. Further, such g2g BFCs create *temporal tussle* points for network management where risks involved in future investments can be tackled better. Taking this to the inter-domain level, multiple g2g BFCs between multiple network service providers will create a platform for higher spatial and temporal flexibility and efficiency for end-to-end bandwidth services.

The forward contracting mechanism introduces a time frame between the time of agreement and the time of actual service delivery. In that sense, a bailout forward contract term can be separated into three stages. The first stage represents the current time (now) where the forward contract advertisements and consequent agreements are made. An important distinction that comes with the forward contracting approach is that there is no payment taking place now within the first stage. The second stage is the time period between the time when the customers lock in the deal and the actual service delivery time. During this period, the provider may bail out of the deal if the conditions arise as long as they are specified (and also agreed) with the bailout clause of the forward contract. This bailout mechanism provides a means to exit from the contract when troubles of extra-ordinary network conditions emerge. However, such bailouts, if frequent, will adversely affect the providers' reputation and the customers' demand for BFC in future. Since *the main innovation behind the bailout forward contracting depends on sharing the risk of unpredictable future between customers and providers of Internet services*, the robustness and reliability of such contracts are crucial for building the trust for wide-acceptance of these tools. In this work, we develop mechanisms for the robustness of BFC definitions so as to minimize the frequency of bailouts. Finally, the third stage of the BFC begins when the delivery of the contract initiates. Once the actual delivery of the service starts, the bailout terms become irrelevant and the main contract terms should be honored. If they are not met, then it will be a breach of the contract and a penalty must apply similarly to the case of today's SLA practices. One can attach penalties to the bailout terms as well; however, that is not a typical practice.

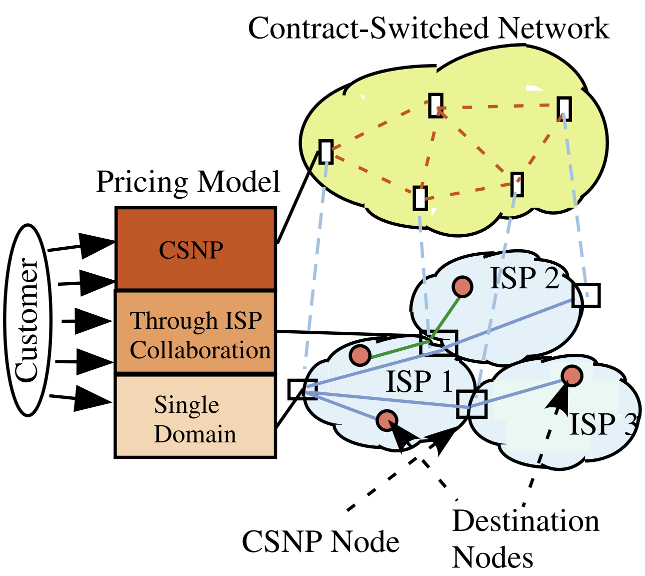

Fig. 1. Major components of an inter-network leveraging edge-to-edge (g2g) contracts. An overlay network provider (called Contract-Switched Network Provider (CSNP)) can concatenate several g2g single-domain contracts to compose an end-to-end path. Distributed end-to-end path composition is also possible with appropriate updates to inter-domain protocols [3].

It is worth emphasizing that contracting mechanisms studied in this paper are discussed and analyzed mainly from the perspectives of ISPs selling and buying edge-to-edge services to/from each other. Thus in this context, the terms 'provider' and 'customer' both refer to ISPs. Also, an ISP can be a provider on some BFCs and a customer on other BFCs. Analysis of the contracting between end users and their provider ISPs, which is a major topic in itself, is beyond the scope of this work. As argued above, BFCs will allow ISPs better manage their risks individually, as well as enable better sharing of risk among ISPs. It is conceivable however that use of BFCs between ISPs would also lead to more flexible terms of service, and better risk sharing, between end users and ISPs. In particular, use of BFCs may help in the realization of differentiated services for end users, offered at different price points. End users who desire guaranteed services would pay more, and their ISPs would have to ensure that their services are least affected in a bailout scenario (often a result of unexpected congestion). On the other extreme, end users that desire best effort service will pay a low price, but would be the first to be affected in a bailout scenario.

The rest of the paper is organized as follows: We first detail a few motivating scenarios and our contributions in the rest of this section. We, then, discuss architectural considerations and implementation issues for BFCs in Section 2. Next, we cover the related literature in Section 3. In Section 4, we formally define bailout forward contracts (BFCs). Section 5 details our proposed method of composing edge-to-edge prices for multiple BFCs for an ISP domain. In Section 6, we build our experimental setup using Rocketfuel topologies and describe our network performance analysis methodology on the robustness of BFCs under stress and link failures. Later in Section 7, we make an economic performance analysis of BFCs in comparison with two other pricing alternatives with specific focus on the ability to manage risks and derived benefits. We summarize our findings in Section 8.

1.1. Motivating scenarios and use cases

To explain how the BFC mechanism works, we now describe a particular market case involving a local ISP and a regional ISP. Golden Gate Telecommunication (GGT), which is a San Francisco-based local ISP, investigates the feasibility of offering value-added IPTV services to its subscribers (end users) in addition to commodity data services. To offer this service, GGT has to embark a risky investment project and upgrade its connection to its Chicago hub, where most of the IPTV broadcast channels are headquartered. The advertisement and high-speed data connection costs are the two big components of this investment project.

During the same period, West Side Telecom (WST), which is a regional ISP, is exploring and evaluating various options to upgrade one of its backbone links amongst a few national hubs: Chicago, San Francisco, Seattle, and Los Angeles. To collect demand figures of customers, WST begins offering forward contracts (i.e. BFCs) at 20% discounted prices over spot market, concerning connectivity services between these major hubs starting next year.

This scenario presents three clear benefits to entities involved:

- *The customer can buy cheaper service via BFCs:* This example scenario presents an opportunity for GGT (the potential customer) for locking some cheap contracts for its investment on improving the performance of connection from San Francisco to Chicago. Under the traditional SLAs, GGT would have to pay high spot prices for purchasing the same level of service from WST.
- *The provider can reduce the risk in costly investments via BFCs:* In the sample scenario, WST will have more accurate demand projections ahead of time, i.e., a year ahead. WST will choose to upgrade backbone link connecting San Francisco and Chicago as a result of higher demand indicated by forward contracts in place with many customers, including GGT. More accurate user demand projection on this link will also provide a better negotiating position to WST with banks and investors in seeking loans/credits to finance its upgrade operations.
- *Bailout clauses can be utilized for early management of congestion and demand–supply balancing:* After a year, suppose that WST begins operating its upgraded San Francisco–Chicago link and starts providing data services advertised previously based on the forward contracts. However, due to congestion (or a failed router) WST is only able to provide a portion of the capacity that was promised to GGT under the forward contracts put in place. Due to this service degradation, WST invokes the bailout clause for some of the BFC contracts in place. Under this condition, GGT (a) would serve somewhat fewer IPTV subscribers, or (b) may buy required capacity on spot market on a higher price to be able to serve all of its IPTV subscribers without a disruption.

Achieving robustness via over-provisioning or low-level reliability are costly, but yet, service quality degradation or disconnections may be costly as well due to opportunity costs or SLA violations. Though the possibility that either entity in a contract may bail out sounds disturbing, the bailout probability provides a means for the contracting entities to express their risk-taking position. Thus, the bailout mechanism in BFCs is a tool provided for ISPs to find a particular answer for each specific case by quantifying the risks and attached costs of particular circumstances.

For example, managing applications requiring more than a basic always-on connectivity will be easier to do when economic tools like BFCs are available. When a crucial link fails in the network, those applications or SLAs traversing the link with tight QoS thresholds will need to be rerouted as soon as possible. This rerouting could be needed either to satisfy the already-promised QoS targets or to take precautions for possible secondary or tertiary failures in the network. In the current practice, the ISP would use *internal* methods like MPLS to reroute the traffic to keep up with its promises. Having the capability of bailout enables the ISP to look for other *external* options like inter-domain paths in addition to the internal possibilities and better manage the risks involved.

A relevant issue is the viability of such bailout offerings within the eyes of customers. In general, reliability of the service is reinforced by additional penalty terms in the cases of breakage or faults in the service. Overlay content distribution services, for instance, offer services with clearly defined penalty (e.g., one day of service for free of charge[3] or return of 25% of the monthly service charge[4]) for the cases when the service falls below a well-defined performance threshold. This is, in general, a concept very similar to the bailouts, except that the provider may quit from the contract before the actual service duration starts and pay the pre-defined bailout penalty. BFCs are explicit risk management tools, where each ISP in the value chain can participate in BFC contracting to manage its demand and service quality risk. Therefore, ISPs will not just contract using BFCs, but treat BFCs as a part of a contract portfolio via which they participate in end-to-end service creation and delivery. Each ISP must assess the demand risk and service quality preference arising from downstream customers/ISPs, and accordingly create a contract portfolio with upstream ISPs where BFCs are one of the contract types. On an aggregate, end users requiring a highly reliable access to bandwidth will be supported using guaranteed contracts. And, flexible end users will be perfectly satisfied with being served using BFC-based bandwidth. Since the utilization of BFCs can reduce bandwidth costs for ISPs, it will be passed onto the flexible end users in terms of cheaper service.

As in the existing SLA practices, the BFCs will also be subject to enforceability requirements. Similar to the existing SLA verifications [6], a third party entity could verify whether or not the ISP did indeed encounter problems to justify the bailout. For instance, lets assume that the bailout clause was to allow the ISP to bail out of the contract when the availability of the contracted link drops below 85%. If the ISP invokes the bailout clause and breaks the contract, it should provide evidence that the links availability did indeed go below 85%. If the customer wants assurances that the ISP is not deceiving because the contracted link at that price is so crucial for him/her, then a third party verifier could be paid at the time of contracting. Even if such third party entity is not utilized, the ISP will still be motivated to be truthful since its reputation will be hindered. The existing SLAs from ISPs are handled in this manner and the reputation of the providers matter a lot in determining the value of the service to the customer. Currently, the ISPs post their services reliability values publicly (e.g., loss or delay in their networks) [7] and the market essentially forces the providers to be truthful [8]. Similarly, we can imagine such reliability values on bailout clauses offered by an ISP, e.g., this ISP's rate of breaking out of a BFC is x%. In general, the ISPs do not have revenue motivation to discriminate in the long run, since deceptive or discriminative behavior in the short-term is not beneficial for the long-term. Thus, in this paper, we assume ISPs are rational and do not addressing malicious or deceptive behavior of ISPs or customers as it is a larger issue and relates to all SLA types, not just BFCs.

1.2. Contributions and major findings

We present an edge-to-edge (g2g) contracting mechanism and pricing framework that opens the doors for end-to-end technologies and services that are otherwise impractical to implement in the current Internet. Our g2g contracting and pricing framework gives the providers and customers flexible and simple knobs to better manage inter-domain economics and value flows, which are essential in attaining end-to-end guarantees lacking in the current Internet. Specifically, the key novelties and contributions of the proposed framework are as follows:

- *Point-to-point demand-sensitive non-linear pricing:* In contrast to the existing static point-to-anywhere schemes, our demand-sensitive pricing schemes can incentivize the providers to invest

[3] Akamai SLA. "http://contracts.onecle.com/websidestory/akamai.svc.2000.03.30.shtml".

[4] Azure SLA. "http://www.microsoft.com/download/en/details.aspx?id=18571".

in innovative services. This will drive more renovation of the Internet's core infrastructure. In realistic topology settings, we show that despite its complexity **point-to-point non-linear pricing** could be deployed efficiently by exploiting price segmentation of services.

- *Forward contracting of edge-to-edge bandwidth guarantees:* We propose and analyze a forward contracting scheme that enables providers to capture demand–supply dynamics on advertised services between different edge-points of its network.
- *Risk sharing between provider and customer through bailout clauses:* Through the use of bailout clauses in the forward contracts, we provide a mechanism for: (i) insuring the provider from the risk of the promised bandwidth being undeliverable due to unavailable capacity, and (ii) assuring the customer that the promised bandwidth will be obtainable at a discount when it is available.
- *Correlation modeling between g2g contracted services with overlapping resources:* Our pricing analysis and risk evaluations show that using simplistic models providers can capture complex interactions between g2g contracts.
- *Robustness of contracting schemes against network disruptions:* We demonstrate that services provided via our g2g contracts deployed on realistic Rocketfuel topologies are **robust** against link failures and topology changes.

2. Architectural considerations and implementation issues

In existing Internet practices *the spatial granularity of SLAs is too coarse.* In the current inter-domain architecture, each ISP domain (or autonomous system) is abstracted as a single node, and thus appears as a single entity in space. This limits the economic flexibility of the network services, as an aggregate price has to be used to express a multi-tiered value in the background [9]. Instead, we abstract each ISP domain as a "set of links", and consider an ISP's domain as an abstraction of multiple g2g contracts involving buyer's traffic flowing from an ingress point to an egress point, i.e., from one edge to another. Our design abstracts 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. Previous work showed that such g2g abstraction of ISP domains can significantly improve end-to-end path qualities [10,11], and this kind of g2g dynamic contracting can be done in a distributed manner with low costs [2]. Fig. 2 illustrates the big picture of such a distributed framework. Customers can only access the 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, where locally computed prices can be advertised with information received from other stations.

Furthermore, in current Internet practices *the temporal granularity of the SLAs is too coarse and they are too rigid to change over time.* If we consider the process that leads to the establishment of an inter-ISP relationship under an SLA, usually this tedious process itself would take weeks and require significant effort from both parties. As a first step of this process, both parties measure the amount of data traffic over a test period to determine the characteristics and amount of the particular traffic flows as it concerns a future agreement. Usually two measurements, namely average and 95th percentile level of traffic will be considered by the SLA as the terms of the traffic. Finally, both parties reach an agreement on an SLA by fixing the terms over reachability, packet loss, and other technical and economic items in addition to these measurements. As a result of this long and cumbersome process, current SLAs are made on long-term basis and their terms are usually not flexible or negotiable over time.

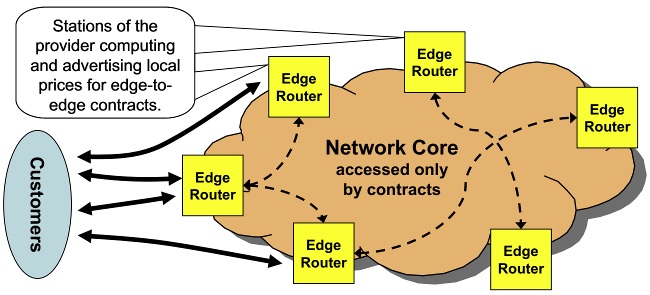

Fig. 2. Dynamic Capacity Contracting (DCC) framework.

BFCs aim to improve the SLAs' temporal flexibility by providing forwarding contracting and bailout options, which allow customers and providers to plan ahead as well as providers to share some of the risk with the customers. In contrast to the current SLA process, our BFC approach offers multiple flexibilities and opportunities to both negotiating parties through its risk sharing mechanisms. In the BFC scheme, parties negotiate over a service which will be delivered in the future instead of now. So, required infrastructure or service capacity may or may not exist at the time of negotiation. The selling party (provider) may deviate from the agreement at the actual delivery time of the service by invoking the bailout clause of the BFC agreement. The buying party (customer) gets the advantage of getting the service at cheaper terms than would be afforded under resource scarcity or high spot price conditions. Therefore, in such a scheme, the customer buys the service at a discounted price in return of sharing the risk of future service delivery and unexpected market conditions with the provider. The provider may also benefit from BFC scheme by capturing the user demand for future services, and also reducing the risk of its investments, by analyzing and managing its portfolio according to more accurate user demand projections.

Tools to handle inter-domain economics are necessary to compose end-to-end QoS guarantees. However, instrumenting financial mechanisms such as BFCs in SLAs for future communication services introduces several architectural research challenges. First of all, definition of contracts imply the consideration of both routing and economic constraints all together in an intra-domain setting. Once such contracts are defined, they must be disseminated through the Internet so as to allow construction of end-to-end multi-domain QoS paths. Constructions of such end-to-end paths on top of contracts makes the negotiation mechanisms between ISPs necessary. Such negotiation mechanisms have been proposed earlier for coordinated route selection and improving resource utilization [12–14] considering best-effort traffic. In case of network services with QoS guarantees, such negotiations should also involve more complicated compliance and monitoring mechanisms for service quality levels. Although we anticipate application of BFCs mainly for elephant flows (which are of considerable traffic volume) between ISPs and enterprise companies, even at this aggregated flow granularity, automated multi-party multi-metric negotiation mechanisms should be developed [15].

Further, composition of better end-to-end paths requires flexibilities in single-domain contracting capabilities. In our design, we consider point-to-point edge-to-edge (g2g) contracting capabilities with unidirectional prices, which pose the question of "How should an ISP price its g2g contracts?" Though similar questions were asked in the literature [16,2] for simple contracts, BFCs require new pricing methodologies. Availability of such flexible g2g contracting provides the necessary building blocks for composing end-to-end QoS paths if inter-domain contracting is performed at sufficiently small time-scales. This distributed contracting architecture gives more flexibilities to users as well, e.g., users can potentially choose various next-hop intermediate

ISPs between two peering points that are involved in users' end-to-end paths [5].

These mentioned challenges have been addressed within the context of future Internet architecture design and have attracted great interest in recent years from researchers [17–21], service providers, and market vendors [22,23]. More needs to be done to implement g2g contracting mechanisms like BFCs on the Internet. In our recent work, we showed how to implement g2g contracts as overlays on top of the existing BGP-based inter-domain routing architecture [24]. Further, we also showed how to reduce the overhead of disseminating (i.e., advertising) the g2g contracts to other ISP domains by augmenting the existing path-vector techniques [25]. However, further exploration and standardization of the g2g contracting concepts are needed. In this paper, our focus is to investigate pricing and risk management aspects of intra-domain edge-to-edge contracts by means of bailout forward contracts. However, we refer interested readers to our work on "Contract-Switching Architecture" which we position contracts as "building blocks" of inter-domain routing and introduce necessary forwarding mechanisms in realization of Contract-Switched Internet [14].

3. Related work

Traditionally network QoS involved the study of different queueing, scheduling and buffer management mechanisms to provide bandwidth and delay guarantees to flows at a statistically multiplexed resource [26]. Several QoS mechanisms have been adopted within single ISP domains, while inter-domain QoS deployment has not become a reality. Arguably the reasons for this include the highly fragmented nature of the ISP market and the glut in core optical capacity due to over-investment and technological progress of the late 1990s. BGP routing convergence and routing instability issues [27] also contribute to inter-domain performance uncertainties. Recent QoS research (e.g., [28,29]) clearly identified a lack of inter-domain business models and financial settlement methods (i.e. monetary flows to compensate for QoS traffic flows), and a need for flexible risk management mechanisms (including insurance, money-back-guarantees). Related QoS pricing issues have been considered as well [30,31]. Consideration of forward contracts with bailout clauses, and the use of risk-neutrality ideas in the pricing of contracts, are some of the key aspects in which our work differs from previous work.

Our work also relates to the Internet pricing research, which has focused on pricing for a single link or for traffic within a single provider's domain [32,33]. Several flavors of pricing schemes have been debated, ranging from flat rate, time-of-the-day, usage-based pricing, to dynamic pricing, and congestion-sensitive pricing [32,2,34]. The consensus is that price elasticity of **individual consumers** is low, and that whatever be the pricing mechanism, it has to be simple, efficient and scalable [35,36]. This simplicity requirement does not translate as is to customers in our context, since customers in our contract-switching paradigm are mainly business enterprises, ISPs, content-providers, Service Overlay Networks (SONs), and other network service providers. We do, however, illustrate that it is possible to design simple pricing of our edge-to-edge contracts to leverage most of the potential economic benefits as was recently shown in the general context of Internet service pricing [9].

In most dynamic pricing schemes, the duration of contract is very short (in milliseconds). Designing and delivering spot and derivative contracts on such time scales is in general quite difficult [37]. In [38], the authors propose a two-component spot pricing framework for intra-domain expected bandwidth contracts with a loss based QoS guarantee. Auction-based pricing models are studied for bandwidth pricing, including combinatorial auctions, progressive second price auction, and simultaneous Dutch auctions [39–43].

Game theoretic analysis of competitive behavior between ISPs as well as interaction between users and ISPs from a pricing standpoint was also analyzed [44–46]. From a cooperative game theoretic perspective, Shapley Value solutions are also considered as they are known for certain properties of stability, efficiency and balanced compensation structures [47,48]. In [47], authors proposed Shapley Settlement Mechanism which offers globally optimal routes coinciding with Nash Equilibrium solutions which encourage ISP connectivity and maximize aggregate profits for selfish ISPs. According to [48], non-monetary peer-to-peer settlements are no longer well-aligned with Shapley Value solutions as service providers are not homogeneous as they were in past.

Forward contracting mechanisms have been implemented for various markets (i.e., electricity, agriculture) regardless of the diverse characteristics of products and services of interest [49–51]. The idea of instrumenting forward contracts in communication market is rather new as recent proposals by several researchers have pointed out the promising benefits of such mechanisms [45,14,46]. Forwards are considered as valuable tools in reducing price fluctuations and providing stable revenue structures for a market where service providers offer bandwidth guarantees on a temporary basis considering congestion-based pricing and congestion control [45]. In the existence of forward contracting options, capacity provisioning and pricing issues for complementary QoS-aware offerings were considered in [46] by a two-stage pricing setting from a game-theoretical stand point in order to investigate the feasibility of forward contracting for communication markets. In this work, we propose risk-neutral pricing mechanisms for managing the risk involved in such forward contracts mainly due to topology changes and demand surges. We also investigate the practical value of introducing forward contracts for risk management and overcoming the unpredictability of future market conditions. Another aspect of our work that differentiates itself from previous work is preservation of dynamic intra-domain routing while considering pricing and risk related issues at the same time. Technical aspects of preserving dynamism in the case of strict SLA compliance previously have been also considered mostly from a network routing and resilience perspective [52].

4. Bailout forward contracts (BFCs)

In general, one can define a contract as an embedding of three major flexibilities in addition to the contracting entities (i.e., buyer and seller): (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 a 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 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 duration the contract will expire in, and the time left for the insured term when the money-back guarantee will expire. Notice that all these three components operate over an aggregation of several packets instead of a single packet. Given the potential scalability issues, this is the right granularity for embedding economic tools into the network protocols instead of finer granularity at the packet level, e.g., per-packet pricing.

In this section, we first formally define a forward contract and a bailout forward contract, and then explain the steps for determining the price of a bailout forward contract. A bailout forward is useful for a provider since it eliminates the risk of demand for bandwidth in the future without imposing a binding obligation to meet the contract if the network cannot support it. A customer of a bailout forward contract locks in the bandwidth required in future, but obtains the bandwidth at a discount. The discount is provided since the customer shares the risk of the scenario that if the network is congested at the future time, the contracted bandwidth may not be delivered due to the bailout clause. The customer may choose not to purchase a forward, but in that case runs the risk of not being able to obtain the necessary bandwidth at the future time due to congestion or price reasons. Therefore, constructing and offering bailout forwards is beneficial for both providers and customers.

Another key benefit of contracting with bailouts and forwards is to encompass practical issues spanning planning, provisioning and operation of a network into a generic framework. Though these operational issues may take place at different time-scales, our aim is to provide an economic tool that can help ISPs to manage risks or opportunities existing in all layers of ISP management. So, the issue of time-scale will be up to the network manager who will decide when to bailout or forward depending on the sense of risk involved in the network dynamics. Issues like network reliability, over-provisioning, actions taken against failures are all economic decisions as much as technical. Though failures or demand spikes are random events by their nature, the ISP manager has to respond to them as they occur. Further, network management might require proactive actions when opportunities arise without necessarily waiting for failures or demand spikes. Likewise, network security is a risk management or cost optimization problem from planning to its execution [53].

4.1. A spot contract

A spot contract is the most basic form of contract. The spot prices reflect present utilization of the network and price the contract using either linear or non-linear pricing kernels to promote utilization and cost recovery. The characterization of risks underlying the spot contract prices is key to formulating the pricing framework. Appropriate modeling abstractions are necessary.

4.2. A forward contract

A forward contract is an obligation for delivering a (well-defined) commodity (or service) at a future time at a predetermined price – known as the 'Forward Price'. Other specifications of the contract are Quality Specification and Duration (start time – T_i, and end time – T_e, for the delivery of a timed service).

Forward contracts, and functionally similar contracts traded in exchanges called 'futures contracts', have been common simple derivatives for risk management for eliminating price risk. The underlying asset or instrument to these contracts are varied, ranging from agriculture products, metals, to energy, electricity and weather. Forwards and futures contracts on agricultural products, electricity and weather are closely related with the context of wired bandwidth, where agricultural products are storable, but are likely to be perishable, electricity has limitations on storability, and weather non-storable and perishable. Specific pricing and analysis mechanism must be developed for these cases [49–51]. A bandwidth contract of certain duration is storable through the duration of the contract, but decays with the passing time to the point of perishing by the end of the contract term.

4.3. A bailout forward contract (BFC)

In the case of a capacitated resource underlying a forward contract, restrictions may be necessary on what can be guaranteed for delivery in future. A key factor that defines the capacity of the resource is used to define the restriction. A bailout clause added to the forward contract releases the provider from the obligation of delivering the service if the bailout clause is activated, i.e. the key factor defining the capacity rises to a level making delivery of the service infeasible. A set up is essential for the two contracting parties to transparently observe the activation of the bailout clause in order for the commoditization of the forward contract and elimination of moral hazard issues. The forward price associated with a bailout forward contract takes into account the fact that in certain scenarios the contract will cease to be obligatory.

Considering multi-provider end-to-end SLA verification problem, both active and passive measurement techniques can be applied as they have been widely used for best-effort services [54]. However, in case of bailout forward contracts, verification protocols which authorize customers and third-party entities to execute temporary verification tasks within provider domain would become instrumental to address potential transparency concerns. Third parties can objectively identify whether or not there was enough congestion in the providers network. By running software tools at the ingress-egress points, they can probe whether different g2g contracts (ingress-egress pairs) are congested or not. As part of the BFC settlement, the customer can require that invoking of the bailout clauses are reported to this third party. Then, the third party will be able to verify whether there has been a violation or not. Crowd-sourcing mechanisms integrated with multiple customer networks and cloud network paradigm also make it possible for a customer to observe and verify network conditions of a provider domain from multiple perspectives. Recent trends have shown great tendency in supporting such capabilities at router level [22].

Further, various bailout "protocols" are possible to implement the bailouts in practice. For example, the provider might decide to bailout and the customer might be given a predetermined amount of time to verify that the bailout conditions really hold or not. Alternatively, the provider might have to get the approval of the customer before bailing out. These protocol possibilities can be negotiated and settled in the bailout clause before the BFC is established.

4.4. Risk segmentation

Creation and pricing of a bailout forward contract on a capacitated resource allows for risk segmentation and management of future uncertainties in demand and supply of the resource. Contracts are written on future excess capacity at a certain price, the forward price, thus guaranteeing utilization of this capacity; however if the capacity is unavailable at the future time, the bailout clause allows a bailout. Therefore, it hedges the precise segment of risk. The price of the bailout forward reflects this.

4.5. Formalization for pricing a BFC

For pricing a bailout forward, we first need to define the price of spot contracts on which the forward is defined. First, we make an edge-to-edge (g2g) "contract" abstraction of the network, where this contract is defined unidirectionally from an ingress edge point to an egress edge point instead of the traditional point-to-anywhere contracting scheme of the Internet. We model the time-dependent demand for the spot contract, μ_t (in Mbps), and the available capacity on this g2g path, A_t (also in Mbps), where $\mu_t < A_t$ by virtue of admission control. Price of the spot contract is obtained as a non-linear transformation, $S_t = P(\mu_t, A_t)$ (in \$/Mbps/s). A predictive model for A_t is used as the bailout factor to define the bailout condition and price the BFCs on this contract abstraction of the network. Therefore, the forward price, F_t, is a function of the spot contract price, predicted future available capacity, and parameters that define the bailout term.

We model the time-dependent demand for the spot contract as follows,

$$d\mu_t = \gamma(m - \mu_t)dt + b_1\mu_t dW_t^1. \tag{1}$$

Therefore, the demand for g2g spot contracts is taken to be generally around a long-term mean level, m, with deviations caused from the long-term mean due to the volatility in demand. The volatility in demand is captured using the Wiener process, W_t^1, where b_1 is the volatility parameter scaling the Wiener process of the extent of risk in demand. When the demand deviates too far above or below the long-term mean, the model captures the tendency of demand to revert to the long-term mean. γ is the rate of mean-reversion parameter. As the Internet traffic demand is known to exhibit fractal behavior at different scales [55,56], we choose to use simplistic Wiener processes to model such behavior [57].

The available capacity on the g2g path is also modeled as:

$$dA_t = \beta(\bar{A} - A_t)dt + b_2A_t dW_t^2. \tag{2}$$

Similarly, we assume that the available capacity is also mean-reverting driven by a second Wiener process, W_t^2. Therefore, the available capacity is in the long-run a level, $\bar{A}$, and experiences deviations from this level captured by the increments in a second Wiener process, W_t^2, scaled by volatility parameter, b_2. After a significant deviation away from the long-run mean, the available capacity is pulled back to the long-run mean at a rate, β. As such the long-run mean, $\bar{A}$, can be time-dependent, perhaps displaying a diurnal or weekly pattern [58], even though it is shown as a constant in (2). Since BFCs are considered for ingress-to-egress (edge-to-edge) paths within an ISP network, the available edge-to-edge capacity will be the residual capacity from the background traffic. Assuming that the background traffic follows a Wiener pattern and there is large enough capacity to capture the background traffic, the residual capacity will be very similar to a Wiener process as well.

Further, for simplicity, we take the two Wiener processes, W_t^1 and W_t^2 to be independent. Although incorporating a correlation between the demand and supply is possible, it is not necessary for the cases here since the time-scale we consider is in the order of many hours. At this time-scale, the dependency between the demand and supply processes does not contribute much to the overall behavior, which is mainly driven by the daily and weekly human activity. Observations have shown that the Internet traffic demand average exhibits an on–off behavior depending on the time-of-day or day-of-week [58]. Although observations at sub-flow or multi-second levels are presenting different pictures [59], the aggregate and multi-flow traffic demand behavior is shown to be driven mainly by time-of-day or day-of-week irrespective of the available capacity [57]. The independence of the demand and supply processes also allows us to look at a scenario where the bailout conditions are less predictable, and thus presents an unfavorable evaluation of our BFC model. If a more correlated model of the traffic demand and capacity supply is assumed, the provider could leverage that for better predicting the future and use it for refining the bailout terms and eventually reduce the bailout frequencies. However, we avoid such demand–supply correlation assumption to make a more rigorous and adversarial evaluation of our BFC concept.

A specific choice of demand profile, described in greater detail in Section 7.2, results in the spot price to be the following function of μ_t and A_t,

$$S_t = P\left(\frac{\mu_t}{A_t}\right) = \int_0^{\mu_t/A_t} p^*(q)dq, \tag{3}$$

where the price of a spot contract responds to the amount of available capacity and the fraction the demanded bandwidth, μ_t, is of the available capacity. Lower the available capacity, the same level of demanded bandwidth will be a higher fraction of the available bandwidth. Therefore, this pricing scheme is congestion-sensitive, and utilizes an optimal nonlinear marginal price structure, $p^*(q)$. $p^*(q)$ is the optimal price schedule in the nonlinear pricing methodology, obtained as

$$p^*\left(\frac{\mu_t}{A_t}\right) = \frac{c + \left(1 - \frac{\mu_t}{A_t}\right) \times \alpha}{1 + \alpha}, \tag{4}$$

which is obtained by maximizing the total surplus for a specific demand profile. Parameters c (in \$/Mbps/s) and α are the marginal cost and the Ramsey number, respectively, used to construct the total surplus and maximization of the same. We apply Ito's formula [60], a well known result of the chain rule in stochastic calculus, to describe the change in the spot price due to change in demand and available capacity.

$$\begin{aligned} dS_t &= \frac{c + \alpha\left(1 - \frac{\mu_t}{A_t}\right)}{1 + \alpha} d\left(\frac{\mu_t}{A_t}\right), \\ &= \frac{c + \alpha\left(1 - \frac{\mu_t}{A_t}\right)}{1 + \alpha} \left[\left(\frac{\gamma(m - \mu_t)}{A_t} + \frac{\beta(\bar{A} - A_t)\mu_t}{A_t^2} - \frac{\mu_t}{A_t^3}\right)dt \right. \\ &\quad \left. + \frac{b_1\mu_t}{A_t} dW_t^1 - \frac{b_2\mu_t}{A_t} dW_t^2\right]. \end{aligned} \tag{5}$$

If $f(S_t, t; T)$ is the price of a BFC at time t, maturing at a future time, T, then the standard derivative pricing derivation, under the corresponding assumptions of existence of delta-hedging strategy, for any derivative defined on the spot contract, S_t, gives that $f(S_t, t)$ should satisfy the following partial differential equation.

$$\frac{\partial f}{\partial t} + \frac{1}{2}p^2\left(\frac{\mu_t}{A_t}\right)\left(b_1^2\frac{\mu_t^2}{A_t^2} + b_2^2A_t^2\right)\frac{\partial^2 f}{\partial S^2} + \frac{\partial f}{\partial S_t}rS_t - rf = 0, \tag{6}$$

along with the end condition,

$$f(S_T, T) = (S_T - F)\mathbf{I}_{\{A_T > Th\}}, \tag{7}$$

where T is the time of delivery of service in future, F is the forward price of bandwidth agreed at time t, and $\mathbf{I}$ is the indicator function for no bailout defined occurring at time T in terms of a threshold level, Th. Therefore, the terminal pay-off of the forward contract is how much benefit having the contract will have over being exposed to the spot price risk, in cases when the bailout clause is not activated. The r in the partial differential equation is the short-term, risk-free interest rate. The derivation entails that a risk-free trading strategy is designed combining the spot and forward (derivative) contract, which in an arbitrage-free world should match a risk-free asset, resulting in the above equation and end-condition that the forward price should satisfy [60]. Using the Feynman–Kac theorem [60], the solution of the above partial differential equation with end-condition is obtained as follows,

$$f(S_0, 0) = E[e^{-rT}(S_T - F)\mathbf{I}_{\{A_T > Th\}}]. \tag{8}$$

Since in the design of a forward contract, there are no payments made at the outset, and only the forward price is determined, we obtain the forward price, F, by equating the above equation to zero and solving for F.

$$F = \frac{1}{P(A_T > Th)} E[S_T \mathbf{I}_{\{A_T > Th\}}], \tag{9}$$

where S_t in the risk-neutral world evolves by the process,

$$dS_t = rS_t dt + p\left(\frac{\mu_t}{A_t}\right)\frac{b_1\mu_t}{A_t} dW_t^1 - p\left(\frac{\mu_t}{A_t}\right)\frac{b_2\mu_t}{A_t} dW_t^2, \tag{10}$$

as inferred from the partial differential equation in (6) for the price of a derivative contract defined on the spot price, S_t.

Had there been no bailout clause in the BFC, which would make it a regular forward contract, the end condition would be

$$f_{regular}(S_0, 0) = E[e^{-rT}(S_T - F_{regular})], \tag{11}$$

and again, since there are no payments initially, the regular forward price simplifies to,

$$F_{regular} = e^{rT} S_0, \tag{12}$$

noting that $e^{-rT}S_T$ is a Martingale in the risk-neutral world, i.e., $E[e^{-rT}S_T] = S_0$.

Therefore, in the regular forward contract, the customer pays full price and gets the service versus in BFC, the customer pays a discounted price and gets the service so long as the provider does not become over-committed (determined by the indicator function in the formula) at the time of delivery. Otherwise, the customer does not get the service and does not have to pay.

To evaluate the performance tradeoffs attained by the BFC definition and pricing framework, we first need to expand it to multiple g2g paths within a single domain. We will then test it in a realistic network topology setting.

5. Single-domain architecture with multiple edge-to-edge BFCs

We 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. Thus, for an ISP with N edge points, there can potentially be $N(N-1)$ advertisable g2g contracts. ISPs would in practice choose to have only a portion of their network capacity available for market, which means only a subset of these $N(N-1)$ g2g contracts will be advertised by the owner ISP. ISPs will have to find a balance between being able to advertise more contracts and manage the additional overhead due to these contracts, which we will look at next.

5.1. Multiple edge-to-edge BFC definition and management

To capture a realistic network topology, we will need to generalize from a single g2g contract abstraction of the network to a set of g2g contracts. For this, appropriate formalization will be necessary, along with information to support the formalization. A "multiple contracts" abstraction will be created, where the available capacity of each g2g contract is modeled, along with a pair-wise interaction of available capacities of g2g contracts to denote the intensity of overlap between the contracts. Definition and pricing of the bailout forward contract terms will be implemented for each contract. However, the multiple contracts abstraction will require the ISP's network topology information, where interactions between the g2g paths are known. Thus, an intensity of overlap will help determine how much g2g capacity is available for a g2g contract C^i, given the existing committed contracts crossing (i.e., overlapping) C^i's g2g path.

5.1.1. Multiple g2g BFCs

In a multiple contracts abstraction of the network, two basic terms need to be modeled: demand and available capacity on each contract link. In this paper, we consider a multiple g2g contracts abstraction of the network, where a contract link ij corresponds to a contract between edge nodes i and j in the network. The time-dependent demand for g2g spot contract on the ij contract link is denoted by μ_t^{ij}, which is modeled for each g2g link by

$$d\mu_t^{ij} = \gamma^{ij}(m^{ij} - \mu_t^{ij})dt + b_1^{ij}\mu_t^{ij} dW_t^{1ij}, \tag{13}$$

taken as a direct adaptation from the original model in (1) and the available capacity on ij contract link, denoted by A_t^{ij}, is also defined by

$$dA_t^{ij} = \beta^{ij}(\overline{A}^{ij} - A_t^{ij})dt + b_2^{ij}A_t^{ij} dW_t^{2ij}, \tag{14}$$

also taken as a direct adaptation from (2). As seen in the above equations, demand and available capacity models are formulated in the form of mean-reverting stochastic processes with driving Wiener processes; long-term mean and volatility parameters which are similar to those described earlier in Section 4.5 for the case of single g2g spot contracts (see Eqs. 1 and 2).

All the parameters described thus far will need to be estimated, as will be discussed later. More complex demand models based on the methods in Gupta et al. [61] and Gupta et al. [38] can be developed, however we have considered a Wiener process based model for maintaining simplicity in our present analysis. The Wiener processes are considered to be sufficient for modeling the fractal behavior of the Internet traffic at large time-scales as they capture the large frequency components of a long-range dependent (LRD) process resembling the Internet traffic [57]. The weakness of Wiener processes is that they may not be able to capture short frequencies in an LRD behaving process. Since our analysis is taking place at large time-scales, we anticipate that the lost accuracy due to our assumption of Wiener processes is minimal.

The Wiener processes underlying the evolution of available capacity of each link are correlated, since different links integrally share network resources. We use an **intensity of overlap** term to model the correlation between available capacity on g2g links:

$$dW^{2ij} dW^{2kl} = \rho^{ijkl} dt, \tag{15}$$

where ρ^{ijkl} is the intensity of overlap describing correlation between link ij and link kl. Therefore, the change in available capacity on a pair of g2g links evolves with a degree of correlation capturing the intensity of overlap of the two g2g abstractions. The details of determining ρ will be provided next.

As described in the previous section, the price of the spot contract is a non-linear transformation, $S_t^{ij} = f(\mu_t^{ij}, A_t^{ij})$ (in \$/Mbps/s). An intensity of overlap, ρ_t^{ijkl}, models the correlation between the contracts, and a predictive model for A_t^{ij}, are used to define and price the bailout forward contracts on each g2g path of the network.

Maintaining the same choice of demand profile for each g2g path gives the spot price to be the following function of μ_t^{ij} and A_t^{ij},

$$S_t^{ij} = P\left(\frac{\mu_t^{ij}}{A_t^{ij}}\right) = \int_0^{\mu_t^{ij}/A_t^{ij}} p^*(q)dq, \tag{16}$$

where the formulation for spot price is adapted from (3). $p^*(q)$ is the optimal marginal price schedule in the nonlinear pricing methodology as before, and hence, requiring an integration to obtain the contract price in above equation. Applying rest of the derivation as in the single g2g path case, the price of bailout forward for contract link ij, $f^{ij}(S_t^{ij}, t; T^{ij})$, is

$$F^{ij} = \frac{1}{P(A_T^{ij} > Th^{ij})} E[S_T^{ij} \mathbf{I}_{\{A_T^{ij} > Th^{ij}\}}], \tag{17}$$

where S_t^{ij} in the risk-neutral world evolves by the process,

$$dS_t^{ij} = rS_t^{ij}dt + p\left(\frac{\mu_t^{ij}}{A_t^{ij}}\right)\frac{b_1^{ij}\mu_t^{ij}}{A_t^{ij}}dW_t^{1ij} - p\left(\frac{\mu_t^{ij}}{A_t^{ij}}\right)\frac{b_2^{ij}\mu_t^{ij}}{A_t^{ij}}dW_t^{2ij}. \tag{18}$$

T^{ij} is the time of delivery of service in future for the ij^{th} g2g path, F^{ij} is the corresponding forward price, and $\mathbf{I}$, as before, is the indicator function for no bailout defined in terms of a threshold level, Th^{ij}. The bailout clause is defined by the available capacity on the g2g link, i.e., if the available capacity, A_T^{ij}, at the time of delivery of the forward contract, T, is below a threshold level, Th^{ij}, then the provider can bailout on the forward contract.

Therefore, as opposed to the single g2g path case, when contracts are defined on multiple g2g paths, the forward price of a g2g path is modified by the extent that evolution characteristic of A_t^{ij} is affected by variability in the available capacity in other g2g paths (e.g., A_t^{kl}), dictated by the intensity of overlap between paths.

5.1.2. Intensity of overlap – ρ^{ijkl}

To evaluate the risk involved in advertising a particular g2g contract, knowledge of the interactions among crossing flows within the underlying network is crucial. As shown in the previous subsection, we develop our multiple g2g BFC terms based on the assumption that an intensity of overlap, ρ^{ijkl}, abstractly models the correlation between flows ij and kl. High correlation means that flows ij and kl are tightly coupled and share more of the network resources on their paths. In other words, an increase in flow ij's traffic will adversely affect the available g2g capacity for flow kl and vice versa.

We construct the correlation information among the g2g contracts as a square matrix of overlapping links. Each entry ρ^{ijkl} reflects the overall effect of flow ij on flow kl, which is the result of the contention that takes place on common links that two flows overlap on their e2e paths. We model this contention as being dominated by the *severity of contention* at the bottleneck link on the g2g path. Thus, we pick the severity of contention on the most utilized common link as the indicator of the correlation between the two overlapping flows. In our calculation, we also reflect the utilization level of bottleneck link as an indicator of severity of race condition among the flows.

Also, we consider the *asymmetric* characteristic of the overlaps arising due to the amount of traffic along different contract paths being not necessarily equal. So, the effect of flow ij on flow kl, is not necessarily equal to the effect of flow kl on flow ij. In that sense, the effect of flow ij on kl is proportional to the ratio of traffic that flow ij generates to the overall traffic generated by this flow pair.

Thus, we model the correlation between flows ij and kl as:

$$\rho^{ijkl} = U_{link} \times \left(\frac{\tau_{ij}}{\tau_{ij} + \tau_{kl}}\right)$$

where τ_{kl} is the portion of bandwidth that flow kl can have according to max–min fair share among all flows passing through the common bottleneck link, and U_{link} is the utilization of the bottleneck link. To calculate τ_{kl} for flow kl, first we calculate bandwidth distribution over every single link using e2e demand for flow kl (i.e., μ_t^{kl}) and the available link capacities (i.e., $C_{N\times N}$). More specifically, on the common bottleneck link, we distribute the available capacity to all passing flows according to max–min fair share. Then, we distribute the excess capacity evenly across all the flows until no excess capacity is left on the link. This strategy makes τ_{kl} being the minimum capacity allocated to flow lk over all links it passes through.

6. Robustness analysis of BFCs

The frequency of bailouts for BFCs will be highly determined by the nature of internalities and dynamics of the ISP's network. Since link/node failures are the major internality for a conventional ISP network, we focused our analysis to the robustness of BFCs against link failures. Among other internalities is the scheduled maintenance of the network, which would be a reason for designing forward contracts. Since realistic modeling of such future events is hard, we chose to use failures as the basis for our robustness analysis, and tried to investigate the forward contracting performance via stochastic modeling of demand in Section 7 later. This section answers the following questions before we deal with details of alternative pricing models: How viable is our edge-to-edge virtual link definitions for constructing BFC definition on top them? What is the probability that a g2g BFC will break due to a link/node failure in the ISP's network?

In order to make performance evaluations of BFCs based on these questions, we need to construct a framework which can house realistic network topologies and a working model of the market. In this section, we first explain the steps for constructing our proposed framework from a network perspective and then, use this framework to examine the robustness of our BFC service definition in question.

6.1. Network model

In our experimental setup, we first devise a realistic network model with Rocketfuel's ISP topologies [62], shortest-path intra-domain routing, and a gravity-based traffic matrix estimation. We assume that the QoS metric of BFCs is the g2g capacity. We focus on developing our network model to reflect a typical ISP's backbone network. Crucial components of a network model include (i) a realistic *topology* (i.e., adjacency matrix, link weights, link propagation delays, link capacities) and (ii) a realistic *traffic matrix*. We first calculate a routing matrix R for the ISP network from the link weight information. With a realistic traffic matrix T, we can then calculate the traffic load pertaining to individual links by taking the product of T and R. We use this realistic network model to identify a demand (i.e., μ) and supply (i.e., A) model, which we use to develop multiple g2g BFCs.

6.1.1. Methodology

For a network with N nodes, L links, and $F = N(N-1)$ flows, let $T_{N\times N}$ be the traffic matrix. If there exists a positive flow from ith node to jth node, then $T_{i\times j}$ is the traffic rate in Mb/s from ith node to jth node; if not, then $T_{i\times j}$ is 0. Let $\lambda_{F\times 1}$ be the traffic vector, which is the vectorized version of $T_{N\times N}$ such that $\lambda_{(i-1)N+j} = T_{i\times j}$ where $i, j = 1 \ldots N$. Let $R_{F\times L}$ be the routing matrix, where $R_{i\times j}$ is 1 if the ith flow traverses the jth link. If not, $R_{F\times L}$ is 0. The network model requires the following inputs:

- *The traffic matrix:* $T_{N\times N}$
- *Topology information:* Adjacency matrix $Adj_{N\times N}$, link weight matrix $W_{N\times N}$, link propagation delay matrix $S_{N\times N}$, link capacity matrix $C_{N\times N}$

Our network model takes the following steps to calculate a baseline where all the g2g traffic can be served feasibly:

- *Step 1:* Construct the *routing matrix* $R_{F\times L}$ based on shortest path first (Dijkstra's) algorithm using the topology information $Adj_{N\times N}$ and $W_{N\times N}$.
- *Step 2:* Form the *traffic vector* $\lambda_{F\times 1}$ from $T_{N\times N}$.
- *Step 3:* Calculate the *traffic load on each link* by performing the matrix operation $Q = R^T\lambda$, where $Q_{L\times 1}$ is the link load vector (in Mb/s).
- *Step 4:* Check the *feasibility* of the traffic load and routing. If any link's capacity is less than the load onto that link, then we fix the infeasibility by increasing the capacity of that link.

6.1.2. Topology

To obtain some of the topology information, we used the Rocketfuel [62] data repository which provides router-level topology data for six ISPs: Abovenet, Ebone, Exodus, Sprintlink, Telstra, and Tiscali. Specifically, it provides *Adj*, *W*, and *S* for the six ISPs, but an estimation of *C* is not provided. Table 1 shows a summary of the topology information for the six Rocketfuel topologies. We updated the original Rocketfuel topologies such that all nodes within a PoP (assuming that a city is a PoP) are connected with each other by adding links to construct at least a ring among routers in the same PoP.

6.1.3. BFS-based link capacity estimation

In order to assign estimated capacity values for individual links of the Rocketfuel's topologies, we use a technique based on the Breadth-First Search (BFS) algorithm. We, first, select the maximum-degree router in the topology as the center node for BFS to start from. After running BFS from the max-degree router, each router is assigned a *BFS distance* value with respect to the center node. The center node's distance value is 0.

Given these BFS distances, we apply a very simple strategy to assign link capacities: Let the BFS distances for routers *i* and *j* be d_i and d_j respectively. For the links (i,j) and (j,i) between the routers *i* and *j*, the estimated capacity $C_{i,j} = C_{j,i} = \kappa[\max(d_i, d_j)]$ where κ is a decreasing vector of conventional link capacities. In this paper, we used: $\kappa[1] = 40$ Gb/s, $\kappa[2] = 10$ Gb/s, $\kappa[3] = 2.5$ Gb/s, $\kappa[4] = 620$ Mb/s, $\kappa[5] = 155$ Mb/s, $\kappa[6] = 45$ Mb/s, and $\kappa[7] = 10$ Mb/s. So, for example, a link between the center router and a router with BFS distance 5 will be assigned 155 Mb/s as its estimated link capacity. Similarly, a link between routers with distances 1 and 3 will be assigned with a capacity estimation of 2.5 Gb/s. The intuition behind this BFS-based method is that an ISP's network would have higher capacity and higher degree links towards center of its topology. This intuition is well-supported by the recent study [63] showing that router technology has been clearly producing higher degree-capacity combinations at core routers in comparison to the edge routers.

6.1.4. Traffic model

A crucial piece in modeling an ISP network is the workload model, i.e., a traffic matrix. In addition to being realistic in size, each traffic flow in the network model must reflect the traffic from edge router to another edge router. Thus, there are two important steps in constructing a reasonable traffic matrix. First, we identify the edge routers from the Rocketfuel topologies by picking the routers with smaller degree or longer distance from the center of the topology. To do so, for each of the Rocketfuel topologies, we identified *Degree Threshold* and *BFS Distance Threshold* values so that the number of edge routers corresponds to 75–80% of the nodes in the topology.

Second, we use gravity models [64,65] to construct a feasible traffic matrix composed of edge-to-edge (g2g) flows. The essence of the gravity model is that the traffic between two routers should be proportional to the product of the populations of the two cities where the routers are located. We used CIESIN [66] dataset to calculate the city populations. We construct an initial traffic matrix based on the gravity model using populations of the cities, and then adjust the BFS-bases link capacity estimations (see Section 6.1.3) so that traffic load on individual links are feasible. This method of generating traffic matrices based on gravity models yields a power-law behavior in the flow rates as was studied earlier [12,64]. We assume that this final traffic matrix reflects the state of the network in a steady state condition. During the simulation, we base our work on this initial condition and analyze the transitions from this initial state of the network.

6.2. Network analysis

In this part, to test the viability of our BFC definitions, we evaluate the performance of our BFCs when a failure occurs in the underlying networks that we compose using six Rocketfuel topology maps, i.e. Abovenet, Ebone, Exodus, Sprint, Telstra, and Tiscali. We aim to examine to what extent existing network topologies are capable of supporting BFC robustness promises. Specifically, we take the baseline BFC definition and identify the fraction of g2g BFCs getting invalidated (i.e., to be bailed out) due to a link failure in each underlying network topology. Notice that this analysis *conservatively* assumes no a priori knowledge of the failure scenarios.

To perform the analysis we take down each link of the topology one by one. During this analysis, we only consider links with non-zero traffic load. After each link failure, we reroute (i.e., compute the new shortest path for) each g2g path and determine the effective g2g capacity each BFC will be able to get based on max–min fair share or equal share when there is excess capacity. For a g2g link, threshold value for bailout is determined as the 15^{th} percentile of the set consisting of bandwidth capacity values observed for that particular g2g link on all possible cases arisen in the aftermath of all simulated single link failures. According to these threshold values, we count the number of bailing out BFCs after each link failure.

Table 2 summarizes the number of link failures simulated, number of g2g traffic flows investigated and the statistics regarding what is the fraction of times a BFC bails out on average due to these link failures within each topology. As the simulation results point out in Table 2, the average fraction of times for BFCs to bail out is well bounded by 15% level for each Rocketfuel topology. For Sprint topology mean value is as low as 12.785% and for Exodus topology it is as high as 14.985%. Fig. 3 shows these results with more details as a histogram of bailout fraction for each

Table 1
Rocketfuel-based router-level ISP topologies.

ISP	# of routers	# of links	Degree (avg/max)	BFS dist. (avg/max)	Degree threshold	BFS dist. threshold	# of edge routers	# of g2g flows
Abovenet	141	922	6.6/20	2.3/4	9	3	108	11,556
Ebone	87	404	4.7/11	3.3/7	6	4	66	4290
Exodus	79	352	4.5/12	3.0/5	6	4	60	3540
Sprintlink	315	2334	7.4/45	2.7/7	9	5	254	64,262
Telstra	108	370	3.8/19	3.5/6	5	4	84	6972
Tiscali	161	876	5.6/31	2.6/5	8	4	125	15,500

Rocketfuel topology. The vertical axes on these histogram graphs show the number (i.e., the count) of g2g BFCs that had to bail out due to the failures in the underlying topology. For example, out of 1484 g2g BFCs/flows in the Tiscali topology (Fig. 3f), about 1200 of them had to bail out because of the 15% of the link failures, but none of them bailed out more than 15% of the time. This means the g2g BFCs on Tiscali will not have to bail out for about 85% of the potential link failures. However, for Abovenet (Fig. 3a), about 65 out of 454 g2g BFCs had to bail out for 27% of the link failures, which indicates a more vulnerable g2g topology for BFCs. Overall, Fig. 3 shows that BFCs over Rocketfuel topologies are pretty robust against single link failures and that BFCs will not have to be bailed out more than 85% of the time on average.

7. Economic benefit analysis of BFCs

For the different contracting schemes with increased computational complexity, although added sophistication in pricing allows more accurate capacity allocation and improves the efficiency of value creation, there is a tradeoff between revenue and computational cost. In order to justify the deployment of more sophisticated contracting, we evaluate and compare the different pricing schemes for their overall economic benefit. In this section, we develop mathematical formulation and derivation of different contracting schemes with different pricing complexity for their benefit analysis. We consider three contracting scenarios (three cases), from the simplest to the most complex in terms of pricing: (i) point-to-anywhere contracts with linear pricing, (ii) point-to-point contracts with non-linear pricing, and (iii) point-to-point, nonlinear with BFCs. In each scenario, we determine the price of contracts, starting with point-to-anywhere spot contracts, followed by point-to-point spot, and finally we model and price bailout forward contracts.

7.1. Baseline case 1: Pt-to-anywhere, linear

In the first scenario, we define the simplest contract scheme. The contracts at each edge node are point-to-anywhere spot contracts, and a flat (linear) pricing scheme is employed, representing the status quo. For a demand profile described below, which is used in all other scenarios, we obtain the optimal flat marginal to determine the point-to-anywhere contract prices. The point-to-anywhere spot prices can be determined for each edge node of a chosen network topology. The pricing scheme is then utilized to compute the revenue generated from this approach during a certain fixed planning horizon.

In order to implement point-to-anywhere pricing consistent with the point-to-point contracts, demand stream μ_t^{ij} is aggregated for all j, defined as aggregate demand M_t^i for each ingress point-to-anywhere node i.

$$M_t^i = \sum_j \mu_t^{ij}. \tag{19}$$

The available capacity for point-to-anywhere traffic is similarly defined as A_t^i. The marginal cost of service of all links is similarly aggregated to compute the marginal cost of point-to-anywhere service at node i. We choose a demand profile of $N(p,q) = 1 - p - q$, defined as the number or fraction of the customer-base that will buy at least q units at the marginal price $p(q)$ [37]. This demand profile with linear relationship between p and q represents a customer-base that is highly sensitive to price changes. Although other choices of demand-profiles with medium and low sensitivity can be considered [37], we use this sample demand profile to represent a population that sees the best-effort service as a good substitute to the higher quality of service access to bandwidth, i.e., if the price of higher QoS bandwidth goes up, there is a strong tendency of customers to revert to best-effort service. The corresponding demand function is obtained, by integrating the demand-profile over demand q, as,

$$D(p) = \frac{(1-p)^2}{2}. \tag{20}$$

We seek the profit maximizing constant marginal p^*, which maximizes the profit function,

$$\text{Profit} = D(p)(p - c_i), \tag{21}$$

where c_i is the marginal cost of service at the node i. The linear spot price for point-to-anywhere contract is obtained as,

$$B_t^i = p^* \frac{M_t^i}{A_t^i}. \tag{22}$$

The spot price for a fixed duration of time is used to estimate the aggregate revenue generated from all the point-to-anywhere traffic at node i.

7.2. Baseline case 2: Pt-to-pt, non-linear

Moving one step up from Baseline Case 1, in this case the provider offers non-linear pricing based point-to-point spot contracts. The comparison of this case with Baseline 1 allows measuring the benefit of added complexity of non-linear pricing of g2g (point-to-point) spot contract. Adding complexity means that now for every g2g link, we conduct more effortful computation in terms of demand modeling and capacity estimation, irrespective of using linear or non-linear pricing.

In a multiple g2g contracts abstraction of the network, we modeled the time-dependent demand for spot contract on each contract link by μ_t^{ij}, and the available capacity on each contract link is modeled by A_t^{ij}, defined previously. In non-linear pricing schemes, the demand characteristics of a population or customer-base are often described by a demand profile. The demand profile $N(p(q), q)$ for a commodity is defined as the number or fraction of the customer-base that will buy at least q units at the marginal price $p(q)$. We apply the Ramsey pricing model [67,68] to determine the optimal price schedule. The guiding principle of the Ramsey pricing model is to develop tariffs that maximize an aggregate of customers' benefits, subject to the constraint that the provider's revenues at least recover its total costs, fixed and variable. With our choice of demand profile, $N(p(q), q) = 1 - p(q) - q$, the optimal price schedule $p^*(q)$ in the non-linear pricing methodology by Ramsey Rule is:

$$p^*\left(\frac{\mu_t^{ij}}{A_t^{ij}}\right) = \frac{c_{ij} + \left(1 - \frac{\mu_t^{ij}}{A_t^{ij}}\right) \times \alpha}{1 + \alpha}. \tag{23}$$

Parameters c_{ij} and α are the marginal cost of service for link ij and the Ramsey number, respectively. Ramsey number captures the extent of monopoly power the provider can exert, where $\alpha = 1$ corresponds to a profit-maximizing monopolistic setting and $\alpha = 0$ is the perfect competition. We will work with an α lying in the interval

Table 2
Contract link robustness.

ISP	# of simulations	# of g2g flows	Fraction of Bailouts (avg/std)
Abovenet	290	454	0.15/0.07
Ebone	170	390	0.15/0.05
Exodus	160	372	0.15/0.03
Sprintlink	494	1456	0.13/0.04
Telstra	134	742	0.12/0.05
Tiscali	333	1484	0.14/0.03

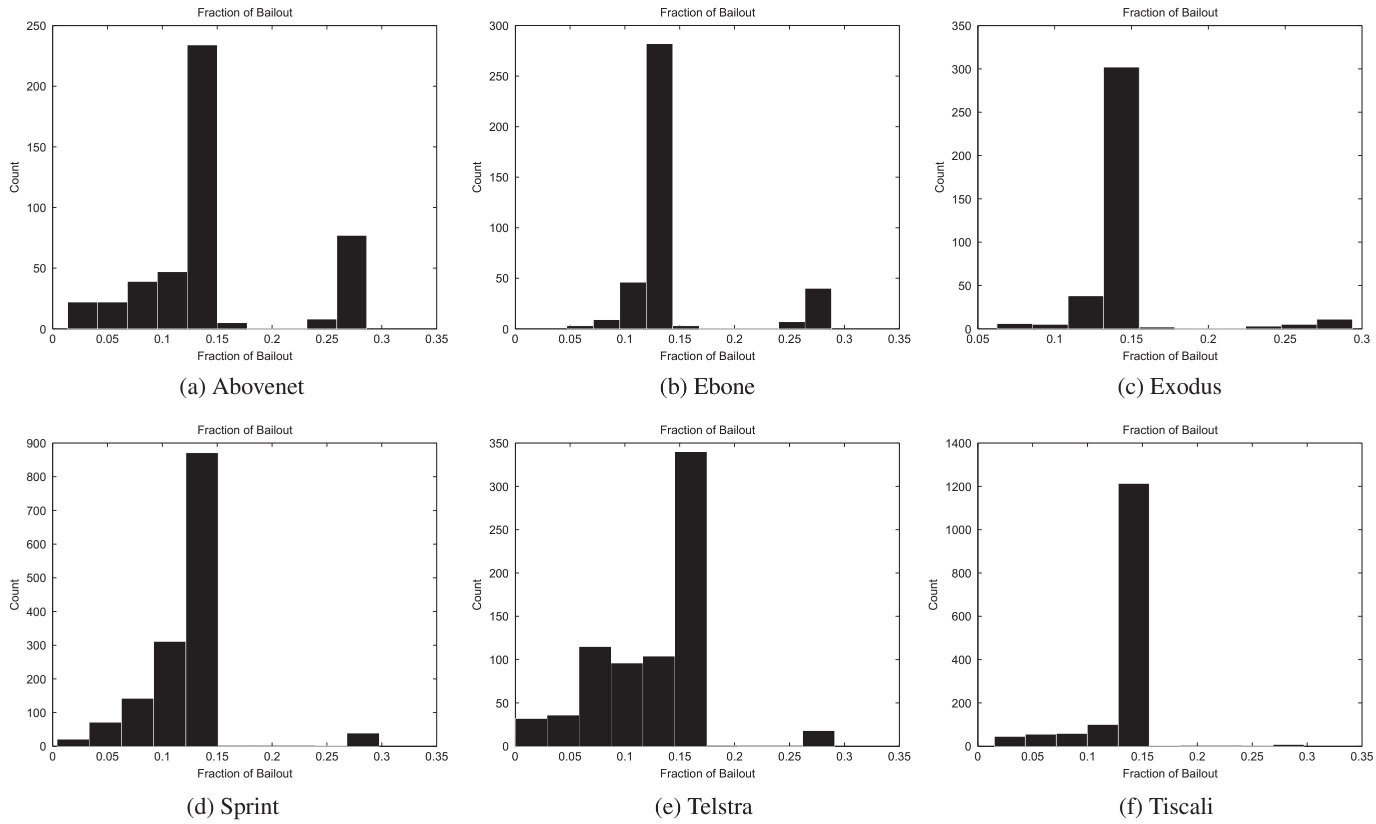

Fig. 3. Histogram of the fraction of g2g paths bailing out after link failures.

$(0,1)$ representing an oligopolistic competition. Price of the spot contract is an integral of optimal price schedule:

$$S_t^{ij} = P\left(\frac{\mu_t^{ij}}{A_t^{ij}}\right) = \int_0^{\mu_t^{ij}/A_t^{ij}} p^*(q)dq, \tag{24}$$

where integral is taken up to the demanded bandwidth expressed as a fraction of the available bandwidth, μ_t^{ij}/A_t^{ij}. This definition acknowledges the capacitated nature of the resources as well as encourages high utilization.

Using the g2g spot prices, S_t^{ij}, for the fixed duration of time, estimates of aggregate revenue are generated from all the point-to-point spot contracts for comparison with Baseline Case 1.

7.3. Bailout forward contract (BFC): Pt-to-pt, non-linear w/bailout forward

For a set of g2g spot contracts of the provider's network, the provider can define and price bailout forward contracts for a chosen set of maturities during the fixed duration of time. It is reasonable to assume that when the provider offers BFCs, a fraction of future demand for spot g2g contracts migrates to g2g BFCs. Under this assumption, the total revenue from g2g spot contracts and forwards is evaluated for the planning period.

To define the bailout feature of the BFC, the model described earlier in Section 5.1, for available capacity, A_t^{ij}, is used to define the bailout clause and price the BFC on each g2g link of the network.

Following the steps described in Section 5.1 for solving F^{ij}, we obtain

$$F^{ij} = \frac{1}{P(A_T^{ij} > Th^{ij})} E[S_T^{ij} \mathbf{I}_{\{A_T^{ij} > Th^{ij}\}}], \tag{25}$$

where S_t^{ij} evolves in the risk-neutral world.

Based on the fraction of the future demand for spot service that migrates to the set of BFCs offered, we obtain the aggregate revenue generated from all the point-to-point spot and forward contracts in the fixed duration of time.

7.3.1. Preliminary analysis of the BFC model

Next we analyze our BFC service model from an economic perspective. We start with analyzing and displaying the model characteristics of a single g2g link BFC. For a single g2g link in the Exodus topology, we display sample paths for the evolution of available capacity, the bandwidth demand and the price of spot contracts in panels (i)–(iii) of Fig. 4. As shown in Fig. 4(i) and (ii), our BFC model uses a more volatile bandwidth demand behavior in comparison to the available capacity, which reflects the comparative behavior of the two in real ingress–egress contracts. The spot price, in turn, behaves in between and shows a moderate volatility as shown in Fig. 4(iii). The probability of bailout, plotted in Fig. 4(iv), shows an increasing trend with an increasing threshold level, as expected. The threshold levels, Th^{ij}, for the bailout clause are defined as a low percentile of the distribution of available capacity, therefore as the threshold level is increased, it is easier and more likely for the provider to bailout from the forward contract.

We report the price of the BFC for a sample of 5 g2g links in Fig. 5, determined within the single g2g framework of Section 5. The BFC delivers service five days in the future with the threshold for bailout set at 15th percentile of the available capacity distribution, i.e., if the available capacity is less than 15th percentile of its distribution at maturity, the BFC will bailout. The objective in this display is to indicate how the forward prices compare with the spot contract prices at maturity. The forward prices remain slightly above the average spot prices at maturity, $E[S_T]$; however, the risk in future spot prices entails that the forward prices will be below future spot prices by a probability, $Prob\{S_T > F\}$, exceeding 45%

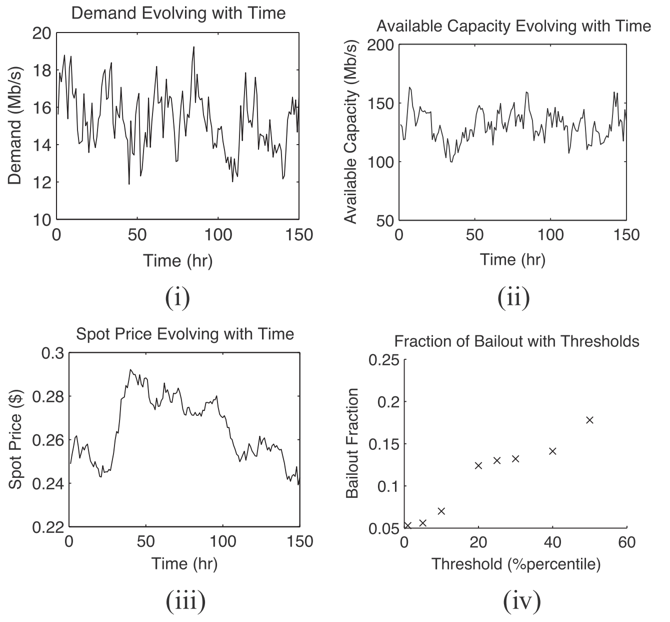

Fig. 4. EXODUS: (i) One sample demand evolution for the next five days; (ii) Available capacity evolution for next five days; (iii) Price of spot contract; and (iv) Probability of bailout as function of Threshold.

(see the fourth column in Fig. 5). For these 5 g2g links, we also indicate the probability of BFCs to bailout in the last column, $Prob\{A_T < Th\}$, i.e., the probability that available capacity, A_T, is less than the 15th percentile threshold, Th. For these links, the probability of bailout is well-bounded by 10%.

We next implement the multiple g2g link framework for BFC pricing of Section 7 to analyze the effect of the interaction between the links that is captured in terms of the intensity of overlap, ρ^{ij}. The forward price of the set of 372 links in the Exodus topology are determined and plotted as a histogram in Fig. 6. As the histogram suggests, although there is variability in the forward prices across the set of links, many of the links pick a forward price in a similar range, in this case approximately around $0.25. This suggests that a distinct forward price for each of the thousands of g2g links in a topology may be an overkill, and hence, directs us to a much desired simplicity in the forward pricing structure. This analysis also is in accordance with the recent findings on tiered pricing of wholesale Internet transit services which indicate that only a couple of pricing tiers are effective enough to ripe near-optimal profits when traffic demand and forwarding costs are considered [9]. Fig. 7 shows a similar analysis for the Abovenet topology. Similar to Exodus, the forward prices for Abovenet g2g contract links (shown in the top graph in Fig. 7) are grouped in a small range around $0.15. So, a similar conclusion can be made about the potential simplicity of pricing the g2g BFCs.

Bailout characteristics are the next important feature to study to evaluate the BFC framework. For Exodus, we plot the fraction

Link	Forward Prices	E[S_T]	Prob{S_T>F}	Prob{A_T<Th}
1	0.20609	0.20305	0.502	0.09
2	0.27162	0.24982	0.449	0.065
3	0.21293	0.21213	0.486	0.079
4	0.25039	0.24825	0.477	0.094
5	0.22177	0.21211	0.465	0.093

Th = 15%

Fig. 5. EXODUS: Sample BFC prices for five g2g links.

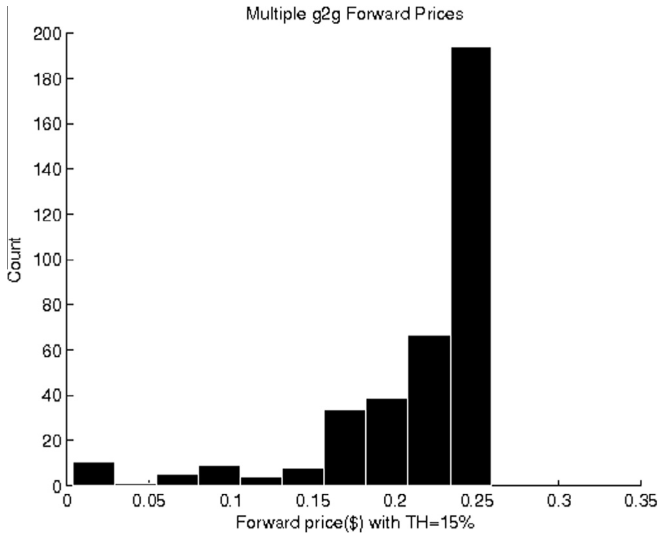

Fig. 6. EXODUS: Histogram of the forward prices for the 372 g2g links.

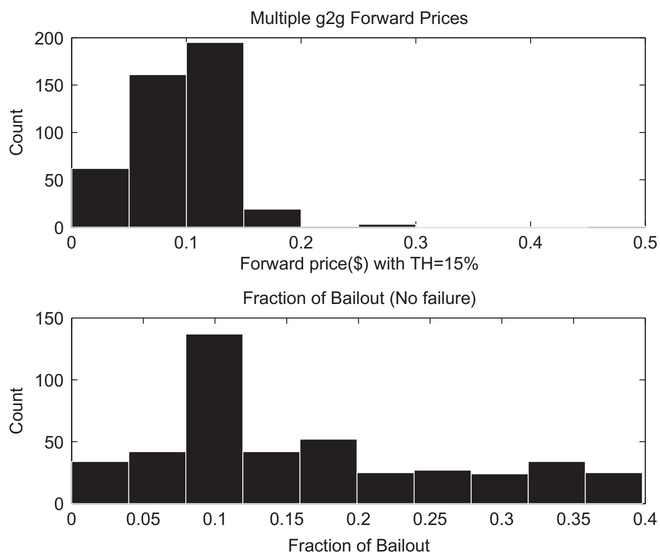

Fig. 7. ABOVENET: *Top* – histogram for forward prices across the 441 g2g links.*Bottom* – histogram of fraction of g2g links bailing out in 1000 runs of simulation.

of 372 g2g links bailing out in 1000 runs of simulation during the planning period in a histogram in Fig. 8. The mean fraction of g2g links bailing out from this histogram is 0.16403, or 16.4%. The bottom graph in Fig. 7 shows the same histogram for the Abovenet topology, where we observe a similar maximum bailout frequency (i.e., $\approx 40\%$), but a more even distribution of the bailout frequency with a peak at 10%. To highlight which specific links bail out in these simulation runs, we also plot the number of times each link bails out in the 1000 runs of simulation during the planning period in Fig. 9 for the Exodus topology. There are a few links that clearly stand out in bailing out most frequently, marking the 'skyline,' while most of the links cluster in the bottom. A similar picture arises for the individual g2g contract links in the Abovenet topology as well.

Another important measure of performance is how much revenue is lost when the BFC on a g2g link bails out. This is shown by each g2g link in the top graph of Fig. 10, while the bottom graph of the figure shows the forward prices on each g2g link. The revenue

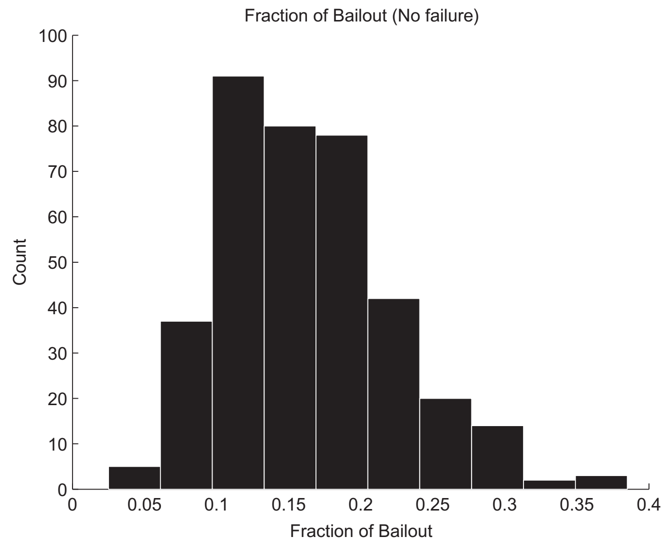

Fig. 8. EXODUS: Histogram of fraction of g2g links bailing out in 1000 runs of simulation.

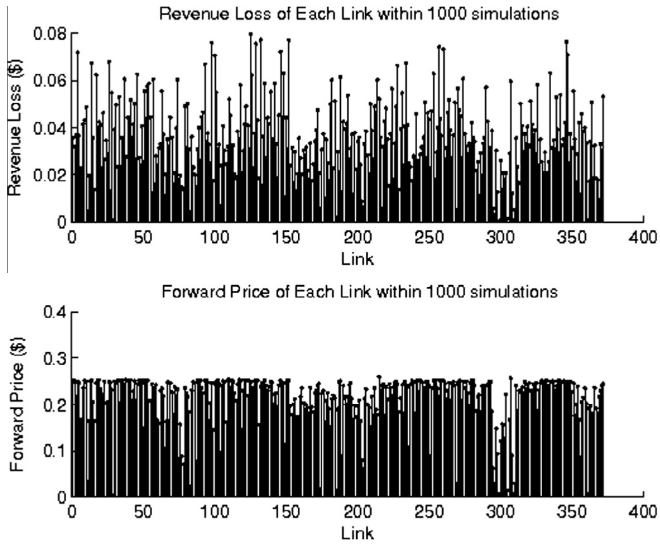

Fig. 10. EXODUS: *Top* – amount of revenue lost for each of the 372 g2g BFCs upon bail out in the 1000 runs of simulation. *Bottom* – forward prices in g2g BFCs.

loss on the BFCs is capped with \$0.08. This is a promising value given that the forward prices reach as high as \$0.27. Clearly, the pattern of clusters here is similar to Fig. 9, however the height of the bars is a function of the forward price of each g2g link and how frequently it bailed out in the runs of simulation. We observe a more even distribution of revenue loss across the links in comparison to the few spikes in the bailout frequencies shown in Fig. 9. This means that bailout risks on the individual BFCs can be successfully suppressed by adjusting forward prices accordingly, and hence, attain a lower revenue loss due to bailouts. To gain more insight into the revenue benefits of BFCs and their tradeoff with other dimensions such as pricing complexity and market predictability, we embark a full-scale analysis next.

7.4. Benefit analysis of BFCs

With our pricing framework formulated earlier in this section, we conduct a series of experiments for a full-scale benefit analysis to assess the added benefit from a more complex contract and pricing structure beyond the currently used point-to-anywhere spot pricing. In these experiments, we examine performance of three contracting schemes with increasing levels of complexity on realistic ISP network topologies. The benefit analysis addresses the following important questions:

- *Complexity trade-off:* How much complexity can be beneficially introduced into the financial component of our contracting schemes?
- *Revenue comparison:* How does the total revenue gained in the three cases compare?
- *Assessing value of information:* What is the value of improved predictability of future demand from the perspective of the provider?

7.4.1. Revenue comparison and complexity tradeoff

We begin our analysis with focusing on the linear pricing scheme as in Baseline Case 1 (henceforth BC1), which is

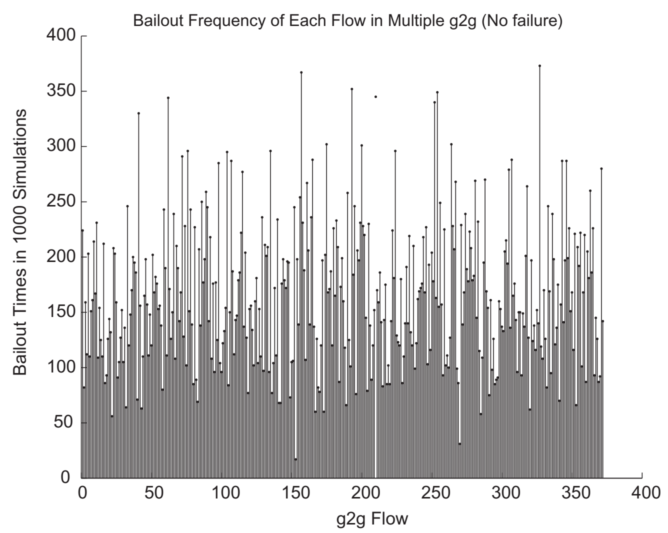

Fig. 9. EXODUS: number of times each of the 372 g2g BFCs bailing out in the 1000 runs of simulation.

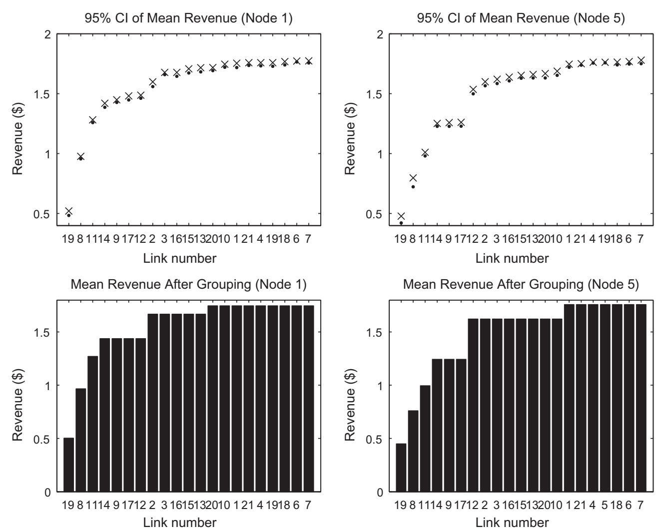

Fig. 11. ABOVENET: Top panel: 95% Confidence intervals for mean revenue for each link of Node 1 and 5 in BC2. Bottom panel: Mean revenue for each group of links of Node 1 and 5 in Reduced BC2.

implemented for each node of the Abovenet and Exodus topologies. We identify a period of seven days to track the revenue comparisons between the three scenarios. Since decision of price complexity needs to be addressed for each link emanating from a node, for illustrative purposes, we track a couple of nodes here. Results obtained for single nodes can be aggregated to give the overall network summary.

We next implement the Baseline Case 2 (henceforth BC2) pricing for each node. In order to address the first research question of complexity, we need to assess the value of treating each point-to-point contract emanating from a node distinctly. In the top panel of Fig. 11, we plot 95% confidence level of revenue obtained from each of the links emanating from two arbitrarily picked nodes in the Abovenet topology. The level of revenue obtained from a set of links is indistinguishable, which merits that they be merged for pricing purposes. We make six groups of links, each group is priced separately, while each link in the group inherits the pricing for the group. This is the reduced Baseline Case 2 (henceforth Reduced BC2). The revenue level of each link is plotted by its group in the bottom panel of Fig. 11. When comparing the top and the bottom panels of Fig. 11, the grouping of the g2g links does not noticeably change the revenue distribution across the g2g links. This is a promising result and attests that we can ripe virtually the same revenue by using only six different prices rather than 441.

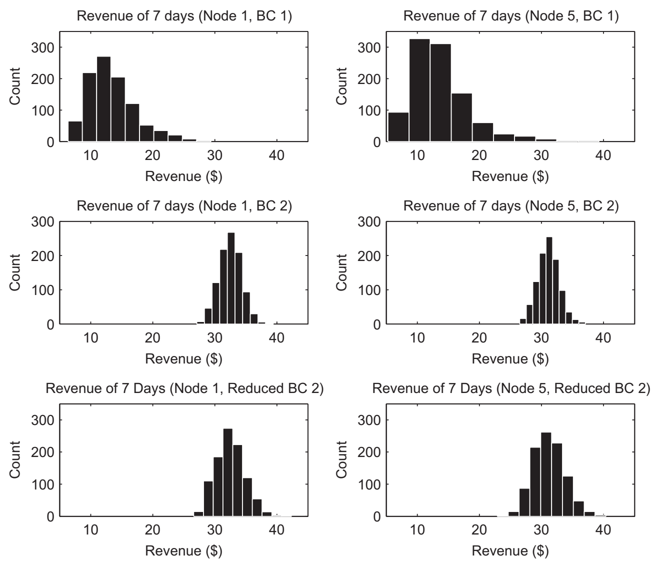

Fig. 12. ABOVENET: Histograms of 7 Day Total Revenue from BC1, BC2, and Reduced BC2 for Nodes 1 and 5.

Revenue from BC2 is clearly expected to be higher than BC1, but we need to investigate the impact of reduction in complexity going from BC2 to Reduced BC2. Histograms for total revenue for the seven days period are plotted in Fig. 12. The top, middle and bottom rows in the figure show the total revenue for BC1, BC2 and Reduced BC2 cases, respectively. Even after a significant reduction in the number of distinct prices being determined in Reduced BC2, the total revenue is almost the same as in BC1. For Node 1, the most frequent peak revenue was about \$12 for BC1, and it moved to about \$33 when utilizing BC2 or Reduced BC2. A similar pattern is observable for Node 5's total revenue, which clearly shows that Reduced BC2 can collect virtually the same amount of revenue as BC2 with only six different prices.

Based on the Reduced BC2, we can now introduce the temporal innovation for BFCs. We define and price the BFCs based on the Reduced BC2 for the distinctly priced group of links. Fig. 13 plots the 95% confidence interval of the mean revenue obtained from the BFCs over the seven days period, for both Abovenet and Exodus topologies. Here for the Exodus topology, it is possible to merge some of the BFCs due to the almost indistinguishable revenue they generate. The level of mean revenue generated after the merging is shown in the left panel of Fig. 13 (four groups of BFCs were made). This is the Reduced BFC implementation for the nodes.

A BFC sells bandwidth at a discount, but locks into future deterministic revenue. We now study the level of reduction in revenue from Reduced BFC when compared to Reduced BC2. While Reduced BFC may be inferior than Reduced BC2 in mean revenue terms, we also want to see how much Reduced BFC is able to exceed the baseline setting of BC1. The total revenue histograms in Fig. 14 are revealing. Reduced BFC significantly dominates the BC1 scenario in terms of total revenue generated.

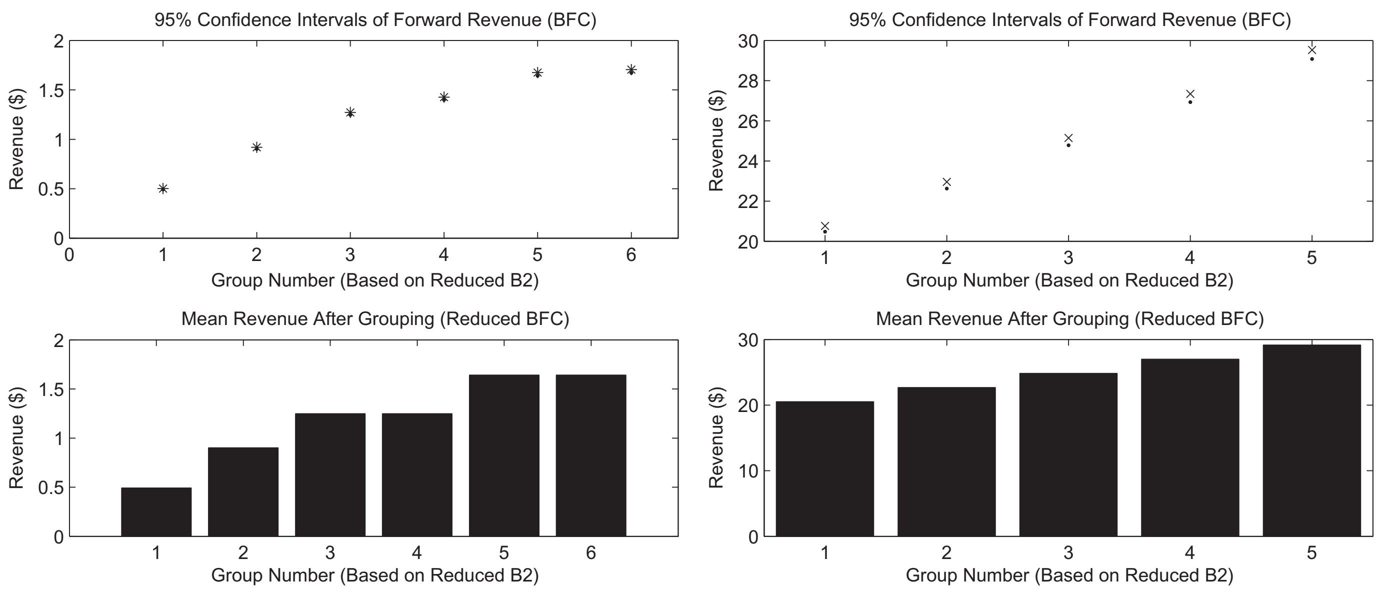

Fig. 13. Left panel: ABOVENET: 95% Confidence intervals for mean revenue in BFC and Reduced BFC (Node 1). Right panel: EXODUS: 95% Confidence intervals for mean revenue in BFC and Reduced BFC (Node 2).

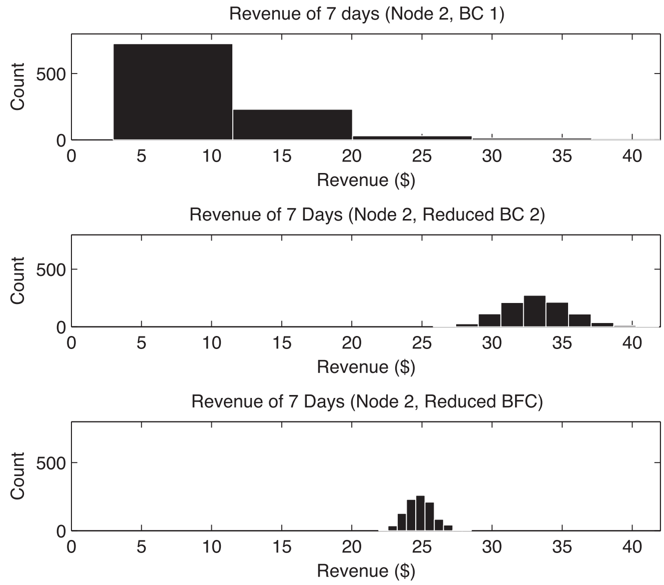

Fig. 14. ABOVENET: Histograms of 7 Day Total Revenue for BC1, Reduced BC2, and Reduced BFC with demand conversion at 40% (Node 1).

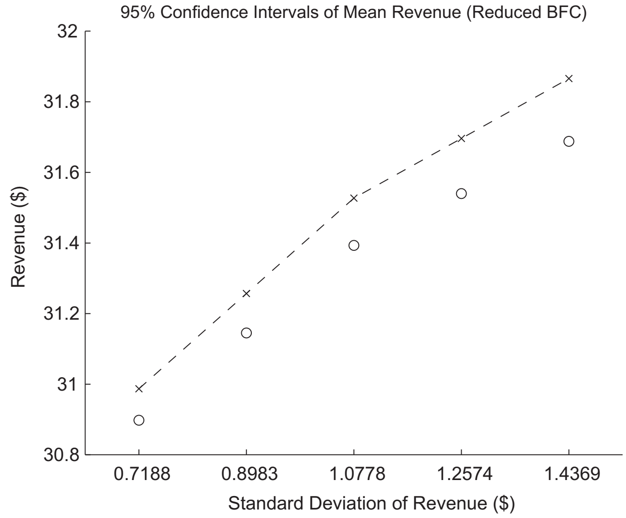

Fig. 16. ABOVENET: 95% Confidence Intervals for Mean Revenue for Reduced BFC plotted by the Standard Deviation of Revenue for changing demand conversion rates (CR = 50, 40, 30 and 20%) (Node 1).

7.4.2. Value of information

A provider can sell all its bandwidth in forward contracts, thus lock into deterministic future revenue, thereby possessing perfect information for future demand for bandwidth. Alternately, the provider may not sell any bandwidth in forward contracts, thus obtaining higher revenues from spot contracts, but with higher variability in total revenue. We now analyze the impact of points in the middle of the two extremes of all or no forward contracts. For all the BFC results displayed so far, the provider lets 40% of its demand for spot contracts migrate or convert to forward contracts. In Fig. 15, we vary this conversion rate and plot the histogram of revenue from Reduced BFC with different conversion rates of demand from spot to forwards. These are compared with the Reduced BC2 (top-left). It is clear that as the demand conversion rate increases, not only does the total revenue decrease on average, the spread of revenue decreases as well. In other words, the provider is trading-off the mean revenue for the variability or risk in the revenue. The cost of lower variability in revenue is paid in terms of reduction in mean revenue – this is the standard risk-return trade-off.

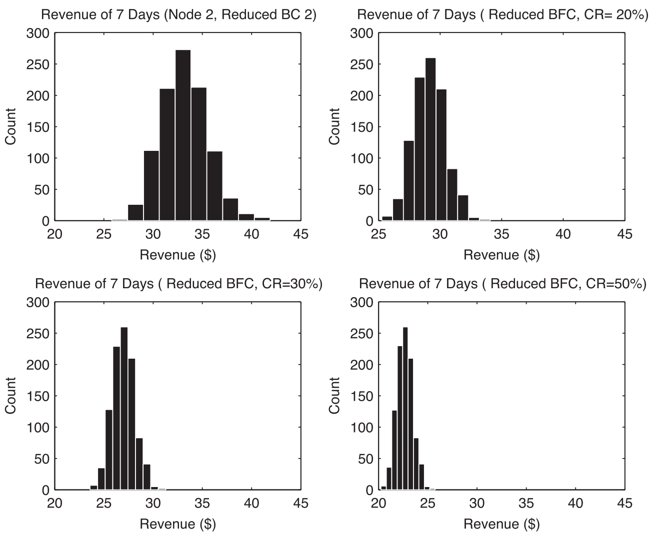

Fig. 15. ABOVENET: Histograms for 7 Day Total Revenue from Reduced BC2 (top-left) and Reduced BFC with different demand conversion rate (CR = 20% (top-right), 30% (bottom-left) and 50% (bottom-right)) (Node 1).

We further highlight this trade-off in Fig. 16. For a range of demand conversion rates between spot and forward contracts, we plot the standard deviation of total revenue versus the confidence interval of mean total revenue. Therefore, for each standard deviation level, there is an associated demand conversion rate, and the two points in the plot indicate the lower and upper limits of the confidence interval of the corresponding mean revenue. With decreasing conversion rate, hence increasing standard deviation (or risk) of revenue, the mean revenue also increases. For reducing the risk in revenues or obtaining better predictions of future demand for bandwidth, the provider has to give up some of its revenue on average. How much return a provider will be willing to give up for reduction in the risk depends on its risk-aversion. A very risk-averse provider will traverse to the left end of the curve, giving up substantial mean revenue for a significant reduction in risk. This will correspond to selling a high fraction of bandwidth in forward contracts. On the other hand, a less risk-averse provider may function at the right end of the curve, selling only a small fraction of its bandwidth in forward contracts. This risk-return trade-off curve provides a clear view for assessing the value of information for a provider, and determining the optimal mode of operation depending on the provider's risk-preference.

8. Summary

In this paper, we compare our dynamic contracting mechanism with traditional pricing models from a service provider perspective by laying out tradeoffs between introduced complexity, level of risk taken in revenue predictions and realized revenues.

We propose a single-domain edge-to-edge (g2g) dynamic capacity forward contracting mechanism with bailout options. In this contracting mechanism, a network service provider can enter into forward bandwidth contracts with its customers, while reserving the right to bail out (for a pre-determined penalty) in case capacity becomes unavailable at service delivery time. We show how the risk-neutral contract prices can be derived for a domain with multiple g2g contracts, which may have correlated demand processes and capacity variations, potentially due to overlapping paths. The proposed risk-neutral contract pricing mechanism allows the ISPs to appropriately manage risks in offering and managing these contracts.

In the proposed architecture, providers can advertise different prices for different g2g paths, thereby providing significantly increased flexibility over the current point-to-anywhere prices. Experiments on a Rocketfuel-based realistic topologies show that our g2g bailout contracting mechanism is quite robust to individual link failures in terms of the bailout fraction and revenue lost.

Several dimensions for future work are possible. It is possible to improve the pricing model for BFCs by consideration of forwarding costs on the g2g paths the BFCs are traversing. Further, it would be interesting to observe the pricing complexity and revenue tradeoff when the BFCs are aggregated based on their destination points, i.e., egresses. A more detailed model for demand and available capacity, potentially supported by empirical data, could be used to analyze BFCs and further provide resolution on more dynamic scenarios to see applicability at finer time-scales. Such applicability at finer time-scales than hours will also need to be studied with a correlated demand–supply model. Under such dynamic situations, it will be useful to explore shorter time-scale contracting metrics such as delay or loss.

Another line of fruitful future work is to model and explore the potential benefits of BFCs to the customers under a multi-provider market where a mixture of BFCs and traditional spot contracts is offered to the customers. With such market analysis, the tradeoff between users' satisfaction and the ISPs risk management benefits could be fine-tuned and lessons for better constructing of bailout clauses could be derived.

Acknowledgment

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

References

[1] W. Liu, H.T. Karaoglu, A. Gupta, M. Yuksel, K. Kar, Edge-to-edge bailout forward contracts for single-domain internet services, in: Proceedings of IEEE International Workshop on Quality of Service (IWQoS), Enschede, Netherlands, 2008.

[2] M. Yuksel, S. Kalyanaraman, Distributed dynamic capacity contracting: An overlay congestion pricing framework, Computer Communications special issue on Internet Pricing and Charging, 2003, pp. 1484–1503.

[3] Contract-Switching: Value Flows and Risk Management Architecture for Future Internet, http://www.cse.unr.edu/yuksem/contract-switching.htm.

[4] D.D. Clark, J. Wroclawski, K.R. Sollins, R. Braden, Tussle in cyberspace: Defining tomorrow's internet, in: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), 2002.

[5] X. Yang, An internet architecture for user-controlled routes, http://www.nets-find.net/Funded/InternetArchitecture.php.

[6] P. Bhoj, S. Singhal, S. Chutani, SLA management in federated environments, Computer Networks 35 (1) (2001) 5–24.

[7] AT&T Global IP Network, http://ipnetwork.bgtmo.ip.att.net/pws/index.html.

[8] F.C.C. (FCC), Measuring Broadband America, July 2012, http://www.fcc.gov/measuring-broadband-america.

[9] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, V.V. Vazirani, How many tiers? pricing in the internet transit market, in: Proceedings of ACM SIGCOMM, Toronto, Canada, 2011, pp. 194–205.

[10] P. K. Muthuswamy, K. Kar, A. Gupta, H.T. Karaoglu, M. Yuksel, ISPs as nodes or sets of links?, in: Proceedings of IEEE International Conference on Communications (ICC), Ottawa, Canada, 2012.

[11] P.K. Muthuswamy, A. Gupta, M. Yuksel, K. Kar, Path-vector contracting: Profit maximization and risk management, Computer Networks 56 (4) (2012) 1286–1302.

[12] R. Mahajan, D. Wetherall, T. Anderson, Negotiation-based routing between neighboring isps, in: Proceedings of USENIX NSDI, 2005.

[13] Y. Shavitt, Y. Singer, Trading potatoes in distributed multi-tier routing systems, in: Proceedings of the NetEcon, 2008, pp. 67–72.

[14] M. Yuksel, A. Gupta, S. Kalyanaraman, Contract-switching paradigm for internet value flows and risk management, in: Proceedings of IEEE Global Internet, Symposium, 2008.

[15] H.T. Karaoglu, M. Yuksel, Effectiveness of multi-hop negotiation on the internet, in: Proceedings of IEEE Global Communications Conference (GLOBECOM), Houston, TX, 2011.

[16] M. Yuksel, Architectures for congestion-sensitive pricing of network services, Ph.D. Dissertation, Rensselaer Polytechnic Institute, Troy, NY, 2002.

[17] GENI exploring future of the networks, http://www.geni.net/.

[18] X. Yang, D. Clark, A. Berger, NIRA: A new inter-domain routing architecture, IEEE/ACM Transaction on Networking 15 (2007) 775–788.

[19] W. Xu, J. Rexford, MIRO: Multi-path interdomain routing, ACM CCR 36 (4) (2006) 171–182.

[20] K.K. Lakshminarayanan, I. Stoica, S. Shenker, J. Rexford, Routing as a service, Technical Report UCB/EECS-2006-19, 2006.

[21] P.B. Godfrey, I. Ganichev, S. Shenker, I. Stoica, Pathlet routing, in: Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication, ACM, 2009, pp. 111–122.

[22] TM Forum IPsphere Framework, http://www.tmforum.org/ipsphere.

[23] The OpenFlow Switch Con. http://www.openflowswitch.org/.

[24] H.T. Karaoglu, M. Yuksel, Path-vector contract routing, in: Proceedings of IEEE International Conference on Communications (ICC), Ottawa, Canada, 2012.

[25] H.T. Karaoglu, M. Yuksel, M.H. Gunes, On the scalability of path exploration using opportunistic path-vector routing, in: Proceedings of IEEE International Conference on Communications (ICC), 2011.

[26] S. Giordano, S. Salsano, S.V. den Berghe, G. Ventre, D. Giannakopoulos, Advanced qos provisioning in ip networks: The european premium ip projects, IEEE Communications Magazine 41 (2003) 30–36.

[27] T.G. Griffin, G. Wilfong, An analysis of BGP convergence properties, ACM SIGCOMM Computer Communication Review CCR, 29(4) (1999) 277–288 [Online]. Available: http://citeseer.nj.nec.com/griffin99analysis.html.

[28] B. Teitelbaum, S. Shalunov, What QoS research hasn't understood about risk, in: Proceedings of the ACM SIGCOMM 2003 Workshops, 2003, pp. 148–150.

[29] M. Gaynor, S. Bradner, M. Iansiti, H.T. Kung, The real options approach to standards for building network-based services, in: Proceeding of IEEE Conference on Standardization and Innovation in Information Technology, 2001.

[30] R. Hampshire, W. Massey, Q. Wang, Dynamic pricing for on-demand bandwidth services, in: Sixth INFORMS Telecommunications Conference, 2002, [Online]. Available: http://www.princeton.edu/ wmassey/PAPERS/pricing.pdf.

[31] R. Hampshire, W. Massey, D. Mitra, Q. Wang, Provisioning for bandwidth sharing and exchange, in: Sixth INFORMS Telecommunications Conference, 2002, [Online]. Available: http://www.princeton.edu/wmassey/PAPERS/provision.pdf.

[32] J.K. MacKie-Mason, H.R. Varian, Pricing the congestible network resources, IEEE Journal on Selected Areas of Communications 13 (1995) 1141–1149.

[33] F.P. Kelly, A.K. Maulloo, D.K.H. Tan, Rate control in communication networks: Shadow prices, proportional fairness and stability, Journal of Operations Research Society 49 (1998) 237–252.

[34] I.C. Paschalidis, J.N. Tsitsiklis, Congestion-dependent pricing of network services, IEEE/ACM Transactions on Networking 8 (2) (2000) 171–184.

[35] A.M. Odlyzko, Internet pricing and history of communications, AT & T Research Lab, Technical Report, 2000.

[36] R.J. Edell, P.P. Varaiya, Providing Internet access: What we learnt from the INDEX trial, University of California, Berkeley, Technical Report 99-010W, 1999.

[37] A. Gupta, S. Kalyanaraman, L. Zhang, A spot pricing framework for pricing intra-domain assured bandwidth services, International Journal of Information Technology & Decision Making 4 (1) (Mar 2005) 35–58.

[38] A. Gupta, S. Kalyanaraman, L. Zhang, Pricing of risk for loss guaranteed intra-domain Internet service contracts, Computer Networks 50 (2006) 2787–2804.

[39] C. Courcoubetis, M. Dramatinos, G. Stamoulis, An auction mechanism for bandwidth allocation over paths, in: 17th International Teletraffic Congress (ITC), 2001, [Online]. Available: http://www.m3i.org/papers/finalITC.pdf.

[40] B. Rupp, V. Grimm, A proposal for an automated bandwidth exchange, in: The Third Berlin Internet Economics, Workshop, 1999.

[41] F. Kelly, R. Steinberg, A combinatorial auction with multiple winners for universal service, Management Science 46 (4) (2000) 586–596.

[42] N. Semret, Market mechanisms for network resource sharing, Ph.D. Dissertation, Columbia University, 1999.

[43] B. Tuffin, Revisited progressive second price auction for charging telecommunication networks, Telecommunication Systems – Modeling, Design, Analysis and Management 20 (3–4) (2002) 255–263.

[44] X.-R. Cao, H.-X. Shen, R. Milito, P. Wirth, Internet pricing with a game theoretic approach: Concepts and examples, IEEE/ACM Transactions on Networking 10 (2) (April 2002) 208–216.

[45] E. Anderson, F. Kelly, R. Steinberg, A contract and balancing mechanism for sharing capacity in a communication network, Management Science 52 (1) (January 2006) 39–53.

[46] M. Reiter, R. Steinberg, Forward contracts for complementary segments of a communication network, in: Proceedings of IEEE INFOCOM, 2010, pp. 2705–2713.

[47] R.T.B. Ma, D.-M. Chiu, J.C.S. Lui, V. Misra, D. Rubenstein, Internet economics: The use of shapley value for isp settlement, in: Proceedings of the 2007 ACM CoNEXT Conference, 2007, pp. 1–12.

[48] R.T.B. Ma, D.-M. Chiu, J.C.S. Lui, V. Misra, D. Rubenstein, On cooperative settlement between content, transit and eyeball internet service providers, in: Proceedings of ACM CoNEXT Conference, 2008, pp. 1–12.

[49] H. Bessembinder, M.L. Lemmon, Equilibrium pricing and optimal hedging in electricity forward markets, The Journal of Finance 57 (3) (2002) 1347–1382.
[50] L. Golden, M. Wang, C. Yang, The pricing of electricity futures: Evidence from the european energy exchange, The Journal of Futures Markets 27 (4) (2007) 319–346.
[51] S. Wilkens, J. Wimschulte, Handling weather related risks through the financial markets: Considerations of credit risk, basis risk, and hedging, The Journal of Risk and Insurance 74 (2) (2007) 387–410.
[52] Z. Cai, F. Dinu, J. Zheng, A.L. Cox, T.S.E. Ng, Contract: Incorporating coordination into the ip network control plane, in: International Conference on Distributed, Computing Systems, 2010, pp. 189–198.
[53] J.-C. Bolot, M. Lelarge, A new perspective on internet security using insurance, in: Proceedings of IEEE INFOCOM, 2008, pp. 1948–1956.
[54] Peering Policy Database, peeringdb, https://www.peeringdb.com.
[55] M.E. Crovella, A. Bestavros, Self-similarity in World Wide Web traffic: Evidence and possible causes formulation and example, IEEE/ACM Transactions on Networking 5 (6) (December 1997) 835–846.
[56] V. Paxon, S. Floyd, Wide area traffic: The failure of Poisson modelling, IEEE/ACM Transactions on Networking 3 (3) (1995) 226–244.
[57] W.-B. Gong, Y. Liu, V. Misra, D. Towsley, Self-similarity and long range dependence on the internet: A second look at the evidence, origins and implications, Computer Networks 48 (3) (2005) 377–399.
[58] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, C. Diot, Packet-level traffic measurements from the Sprint IP backbone, IEEE Network (2003).
[59] T. Karagiannis, M. Molle, M. Faloutsos, A. Broido, A nonstationary Poisson view of Internet traffic, in: Proceedings of IEEE INFOCOM, 2004, pp. 1558–1569.
[60] P. Wilmott, Derivatives: The Theory and Practice of Financial Engineering, John Wiley, Hoboken, NJ, 1999.
[61] A. Gupta, B. Jukic, M. Li, D. Stahl, A. Whinston, Estimating internet users demand characteristics, Computational Economics 17 (2001) 203–218.
[62] N. Spring, R. Mahajan, D. Wetherall, Measuring isp topologies with rocketfuel, in: Proceedings of Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM), 2002.
[63] L. Li, D. Alderson, W. Willinger, J. Doyle, A first principles approach to understanding the internet's router-level topology, in: Proceedings of ACM SIGCOMM, 2004.
[64] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, C. Diot, Traffic matrix estimation: Existing techniques and new directions, in: Proceedings of ACM SIGCOMM, 2002.
[65] Y. Zhang, M. Roughan, C. Lund, D. Donoho, An information-theoretic approach to traffic matrix estimation, in Proceedings of ACM SIGCOMM, 2003.
[66] The Center for International Earth Science Information Network (CIESIN), http://www.ciesin.columbia.edu.
[67] R. Wilson, Nonlinear Pricing, Oxford University Press, Inc., New York, NY, 1993.
[68] R. Wilson, H.-P. Chao, Design of Wholesale Electricity Markets, Electric Power Research Institute, Stanford, CA, 1999.