Design and Benefit Analysis of Edge-to-Edge Bailout Forward Contracts for Single-Domain Internet Services

Aparna Gupta
Lally School of Management and Technology
Rensselaer Polytechnic Institute
Troy, NY

Collaboration with
K.Kar , W.Liu (RPI), H.T.Karaoglu, M.Yuksel (UNR)

About FIND (Future Internet Design) Program

- A major new long-term initiative of the NSF NeTS research program.
- Engages a research community to consider what the requirements should be for a global network of 15 years from now *if we could design it from scratch.*
- It solicits research across the broad area of network architecture, principles, and mechanism design, aimed at answering questions as:
  - How can we design a network that is fundamentally more secure and available than today's Internet? How would we conceive the security problem if we could start from scratch?
  - How might such functions as information dissemination, location or identity management best fit into a new network architecture?
  - What will be the long-term impact of new technologies such as advanced wireless and optics?
  - How will economics and technology interact to shape the overall design of a future network?
  - How do we design a network that preserves a free and open society?

General Motivation

- Current architectural problems:
  - Users cannot express value choices at sufficient granularity – only at access level
  - Providers do not have economic knobs to manage risks involved in
    - investing innovative QoS technologies and
    - business relationships with other providers

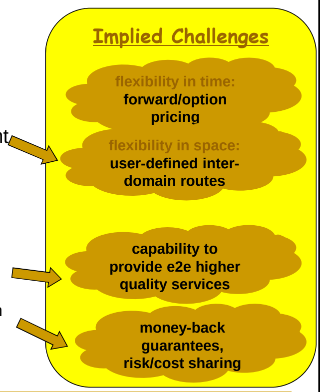

Inter-domain struggles…

- When crossing domains, all bets are off..
- End-to-end reliability or performance-criticality requires
  - assurance of single-domain performance, i.e., "contract"s
  - efficient concatenation of single-domain contracts
- Inter-domain routing needs to be aware of economic semantics
  - contract routing + risk management
- We address translation of these struggles to architectural problems

Contract-switching: A paradigm shift…

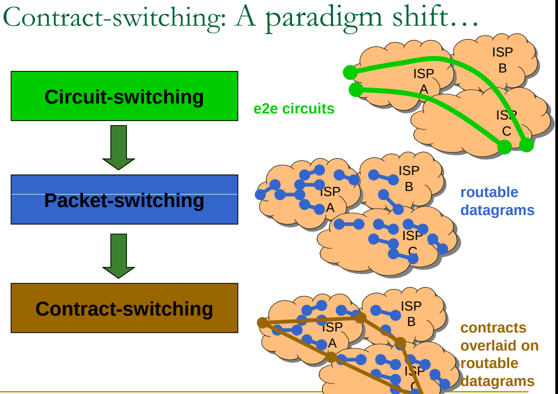

A Contract-Switched Network Core

- Contracts: a practical way to manage "value flows"
- Technologies to Support QoS
- Economic considerations for service definition and delivery
  - Scalability, Efficiency and Fairness
  - Contract timescales
  - Cost recovery
  - Pricing the risk in QoS guarantees
  - Single-domain and end-to-end contracts

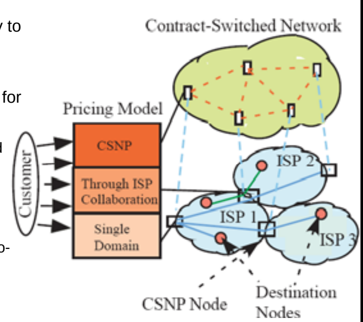

The Contract-Switching Paradigm (CSP)

- Utilize overlay contract links between edge nodes (peering points) at domain boundaries
  - To indicate wider range of service choices.
- Contracts are the building block
- Contracts include performance, financial and time duration specification

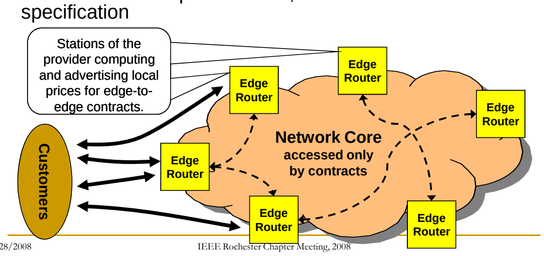

Defining the Contracts in the CSP

- Time Duration for Contracts

- Financial Component – Price discovery
  - Pricing in medium and long timescale
  - Pricing for bandwidth and allowing contracts to be composed dynamically in time
  - Pricing for cost recovery and risk management
- Financial Component – Complexity trade-off
  - Introduce measured sophistication justifying the economic benefit
  - Evaluate 3 scenarios of increased complexity.

Increasingly Complex Contracting Scenarios

- Baseline Case 1:
  - point-to-anywhere
  - linear price schedule designed for cost recovery
  - responsive to demand
- Baseline Case 2:
  - point-to-point
  - nonlinear price schedule designed for cost recovery
  - responsive to demand profile
- Bailout Forward Contract Case:
  - point-to-point, nonlinear price schedule
  - bailout forward for dynamic temporal composing of bandwidth services and risk management

Baseline Case 1 (BC1)

- Contracts at each edge are point-to-anywhere spot contracts
- Flat (linear) pricing scheme
- Demand profile $N(p,q)$ – Number or fraction of customers who purchase q-th unit of product at p. We choose a demand profile:

$$N(p,q) = 1 - p - q$$

- The linear spot price for point-to-anywhere at node i is:

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

M is the aggregate flow through node i and A is the available capacity at node I

- p* is the optimal marginal price obtained from price optimization for cost recovery for the above demand profile

Baseline Case 2 (BC2)

- Price of the spot contract is a non-linear transformation of time-dependent demand and available capacity.

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

$p^*(q)$ is the optimal nonlinear price schedule obtained from price optimization for cost recovery (demand profile from BC1)

μ_t^i A_t^i are the demand and available capacity modeled by two Ito processes

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

- Ito's formula describes the change in the spot price due to changes in demand and/or available capacity.

Link Demand and Capacity Models

- The time-dependent demand for spot contracts on each g2g link:

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

- The available capacity on each g2g path:

$$dA_t^i = \beta^i (\overline{A}^i - A_t^i)dt + b_2^i A_t^i dW_t^{2i}$$

- The intensity of overlap between links
  - The correlation between the driving Wiener processes

$$dW^{2i} dW^{2j} = \rho^{ij} dt$$

Bailout Forward Contract (BFC) Case

- Bail-out Forwards Contracts on advertisable spot contracts
  - between peering/edge points i and j of an ISP
  - flexibility of advertising different **forward prices** for edge-to-edge (g2g) intra-domain paths
  - forwards contracts with provision for **Bail-out** conditioned on network congestion
  - spot and forwards concatenated to create long-term contracts
  - use to realize revenue stability and guaranteed network utilization
  - tool for demand prediction and network upgrades

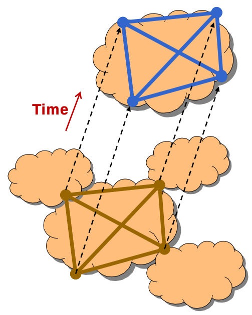

Pricing of Bailout Forward Contract (BFC)

- Based on option pricing derivation, the price of the bailout forward satisfies:

$$\frac{\partial f^i}{\partial t}+\frac{1}{2}p^2(\frac{\mu_t^i}{A_t^i})((b_1^i)^2(\frac{\mu_t^i}{A_t^i})^2+(b_1^i)^2(A_t^i)^2)\frac{\partial^2 f^i}{\partial S^2}+\frac{\partial f^i}{\partial S_t^i}rS_t^i-rf^i=0$$

- With the end condition:

$$f(S_T^i,T)=(S_T^i-F)\mathrm{I}_{\{A_T^i>Th^i\}}$$

- The solution is obtained as:

$$F=\frac{1}{P(A_T{}^i>Th^i)}E[S_T{}^i\mathrm{I}_{\{A_T>Th^i\}}]$$

T is the time of delivery of service in future, F is the forward price, and I is the indicator function for no bailout defined in terms of a threshold level, Th.

Implementation Setup

- Network topologies
  - Two of the Rocketfuel ISP topologies with different network characteristics:
    - Abovenet - well-engineered, stable
    - Exodus - hub-and-spoke
- Experimental Specification
  - Inputs: A, M, μ, ρ (Get for the two topologies), Th(15% percentile), time duration(7 days)
  - Simulate each process and determine prices for a 7 day period
  - Use 1000 replications of simulation

Implementation Setup

- Realistic Simulation requires
  - Realistic ISP Topology
    - Adjacency Matrix (Given by Rocketfuel Data)
    - Link Delays & Weights (Given by Rocketfuel Data)
    - Link Capacities (we model)
    - Edge and Backbone Router Classification (we model)
    - Routing Matrix (Path calculated by Shortest Path Algorithms, as the OSPF and BGP protocols' do for real world)
  - Realistic Traffic Model
    - Traffic Demand (we model)

Implementation Setup

How to assign link capacity?
1) Distance from network center (BFS)
2) Connectivity Degree

Assign higher capacity to links between routers with low distance and high connectivity

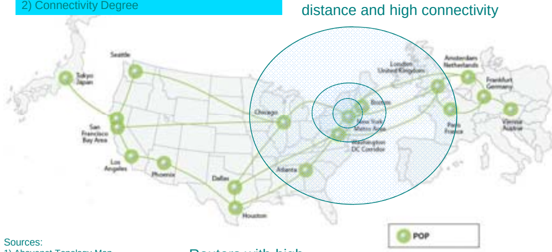

Sources:
1) Abovenet Topology Map
http://www.above.net/products/datasheets/IPTransit.pdf
2) CIESIN Population Data
http://www.above.net/products/datasheets/IPTransit.pdf

Routers with high connectivity and low distance are backbone routers

Implementation Setup

Gravity Model
Associate Regional Populations with Edge Routers and model demand size according to that

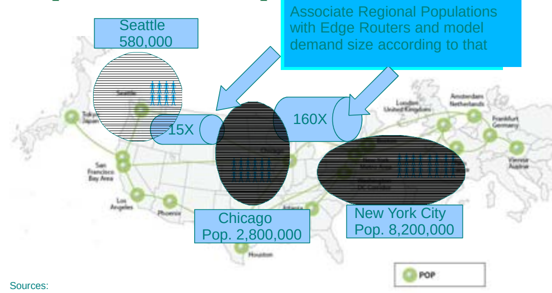

Sources:
1) Abovenet Topology Map
http://www.above.net/products/datasheets/IPTransit.pdf
2) CIESIN Population Data
http://www.above.net/products/datasheets/IPTransit.pdf

Implementation Setup

San Francisco – London
Seattle – London
Chicago – Paris

All flowing through common links between NY and London, what will be the consequence of that ?

Sources:
1) Abovenet Topology Map
http://www.above.net/products/datasheets/IPTransit.pdf

Implementation Setup

$$\rho^{ij} = U_{link} \times \left(\frac{\tau_i}{\tau_i + \tau_j} \right)$$

- Intensity of Overlap ρ
  - Models the severity of competition impact from edge i on edge j for available bandwidth capacity of link
  - Indicator of congestion risk
  - U_{link} is the highest utilization value among common links on g2g path
  - τ_k is the minimum of bandwidth share that $flow_k$ can get over links on the g2g path according to min-max fair share.

Revenue Comparison between BC1 and BC2

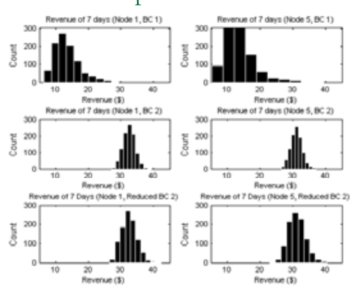

Use ABOVENET Data

- 7 Day Total Revenue Histograms for BC 1 and BC 2
- Total revenue is much more favorable for BC2
  - At the cost of additional complexity

Reduced BC2 -- Reduce Complexity in BC 2

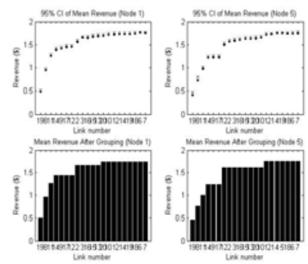

Use ABOVENET Data

- Top panel: 95% CI for Mean Revenue for each link of Node 1 and 5 in BC2.
- Bottom panel: Mean Revenue level for 6 group of links of Node 1 and 5 in Reduced BC2

Revenue Comparison between BC1, BC2, and Reduced BC2)

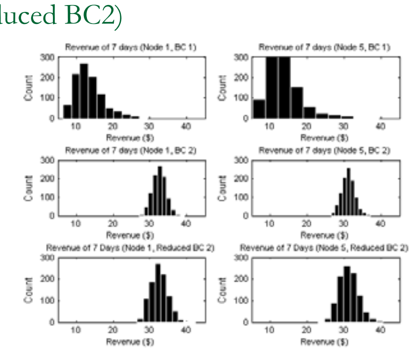

Introducing the BFC (with reduced complexity)

- Reduced BFC:
  - obtained by similar principle as Reduced BC2,
  - with links grouped by similar forward prices

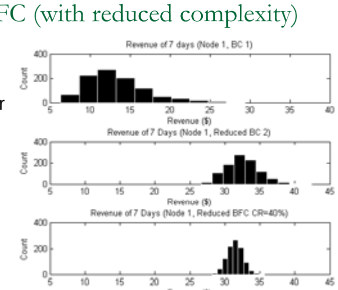

- 7 Day Total Revenue comparison for BC 1, Reduced BC 2, and Reduced BFC with demand conversion rate (CR) at 40%
- Reduced BFC significantly dominates BC1, but slightly inferior to Reduced BC 2

Revenue Comparison of Reduced BFC with varying demand conversion

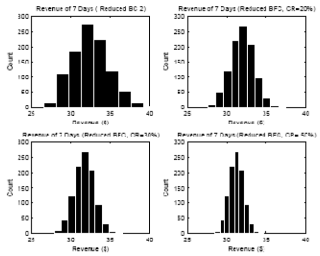

- The provider is trading-off the mean revenue for the variability or risk in the revenue.

Revenue Comparison of Reduced BFC with varying CR

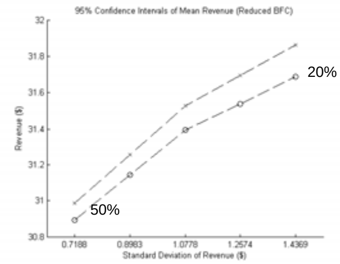

- 95% CI on Mean Revenue vs. Standard Deviation of Revenue for Reduced BFC with different demand conversion rates
- The provider gives up mean return for reduction in the risk, depending on her risk-aversion

Take away from the Economic Benefit Analysis

- Nonlinear, point-to-point pricing of contracts significantly improves revenue over linearly priced point-to-anywhere contracts.
- Grouping of links along with nonlinear pricing retains the benefits over linear pricing, with considerable reduction in computational complexity.
- Bail-out Forward contracts, with controlled complexity, give:
  - nice tradeoff between risk and return
  - flexibility of prediction of future demand
  - possibility of concatenation for longer duration service

How often do BFC Bail-out? - Robustness of g2g BFC

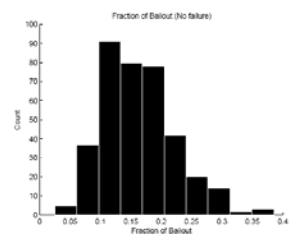

- Use Rocketfuel's ISP topology - Exodus.
- Histogram of fraction of BFCs bailing out
  - Under normal network conditions

How often do BFCs bail out in network failures?

- Three failure modes created by failing specific high load links for this analysis.
- The failures change the network characteristics in the model by changing
  - intensities of overlap between links,
  - means of available capacity, and
  - standard deviations of available bandwidth.

Bailout Frequency of 3 Failure Cases

- Bail out of BFCs on 372 g2g paths on Exodus under Failure mode 1-3

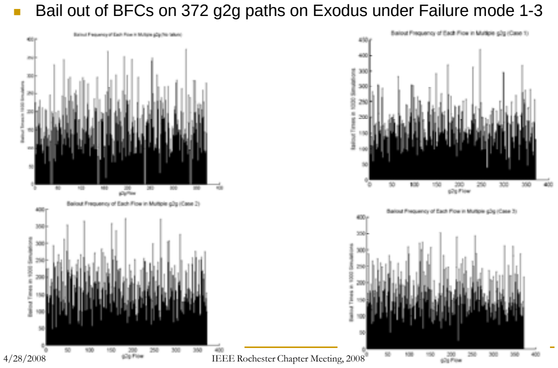

Revenue Impact of BFC, with and w/o failure

Case	Expected Total Revenue	Mean Bailout Fraction
Artificial No Bailout or Failure Case	**95.7464**	0
Base Case Bailout Scenario	80.43655	0.16403
Bailouts in Failure Mode 1	78.98833	0.16505355
Bailouts in Failure Mode 2	81.34074	0.163980954
Bailouts in Failure Mode 3	80.98213	0.16676308

- There is a small increase in the fraction of paths bailing out in the failure modes
- There is a small reduction in revenue in the failure modes

Network Analysis

- Conservative Assumption
  - Although for real world more failures occur at edge routers, we fail every link in our network, including high capacity backbone links one by one.
  - As link fails, shortest path calculations and routing matrix change accordingly
  - Traffic previously passing over failed links shifts to other links following updated routes
  - According to the changed link loads and capacity figures, even under this conservative failure scenario **73% of BFCs** still achieve their promise, on average
  - These results underline the robustness of the BFC model

Network Analysis

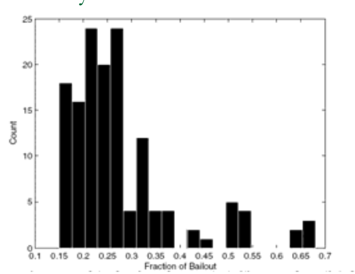

This is more conservative since we are considering all links failing

Summary

- Nonlinear, point-to-point pricing of contracts significantly improves revenue over linearly priced point-to-anywhere contracts.
- Grouping of links along with nonlinear pricing retains the benefits over linear pricing, with considerable reduction in computational complexity.
- Bail-out Forward contracts, with controlled complexity, give:
  - nice tradeoff between risk and return
  - flexibility of prediction of future demand
  - possibility of concatenation for longer duration service
- Experimentations shows that the g2g BFC mechanism is robust to link failures, both in terms of the bailing out behavior and revenue lost.