Overlay/Underlay Interaction

1

Outline

- Motivation
  - Problems with the underlying routing system
  - Source routing, overlay networks, and hybrids
- Overlay networks
  - Pros: flexibility, limited overhead, & value-added
  - Cons: data-path overhead, probes, & feedback
- Negative interactions
  - With other overlays: the price of anarchy
  - With the underlay: influence on traffic engineering
  - With itself: bi-stability and trunk reservation
- Future directions

2

What's Wrong With Internet Routing?

- Restrictive path-selection model
  - Destination-based packet forwarding
  - Single best BGP path per prefix
  - BGP routing constrained by policies
  - Ignoring congestion and delay
  - Ignoring application requirements
- Unappealing protocol dynamics
  - Persistent oscillation (due to policy conflicts)
  - Slow convergence (due to path exploration)
  - Lost reachability (due to route-flap damping)

Need for routing to scale to millions of routers

Putting More Power in End Hosts

- Source routing (e.g., Nimrod)
  - End host selects the end-to-end path
  - Routers simply forward packets on the path
  - Requires the routers to agree to participate
- Overlay networks (e.g., RON)
  - Conventional computers act as logical routers
  - Real routers deliver packets to intermediate hosts
  - No need for cooperation from the real routers
- Hybrid schemes
  - Source routing at the AS level
  - Source routing in the overlay network

Overlay Network

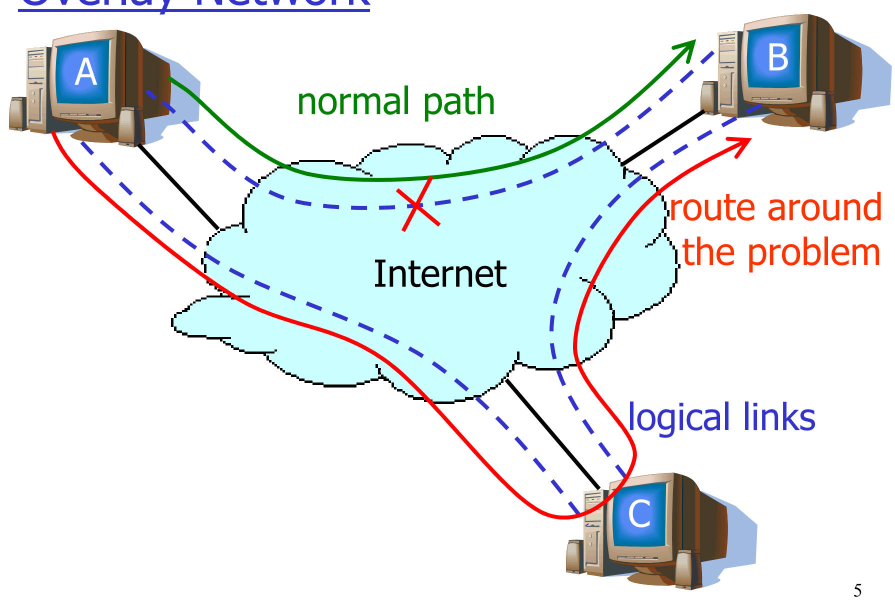

5

Advantage: Flexible Routing

- Paths that violate BGP routing policy
  - E.g., A to C goes through AT&T and Sprint
  - … and C to B goes through UUNET
  - BGP would not allow AT&T-Sprint-UUNET path
- Quick adaptation to network problems
  - Fast detection of congestion and outages
  - … by probing as aggressively as necessary
- Selecting paths based on different metrics
  - E.g., overlay selects paths based on latency
  - … whereas the underlay might try to balance load

6

Advantage: Fewer Worries About Scalability

- Small number of nodes
  - Just enough nodes to have diverse paths
  - A few friends who want better service
  - Virtual Private Network of several corporate sites
- Balancing the trade-offs
  - High probe frequency for maximum adaptively
  - Low probe frequency for minimum overhead
- Simple routing protocol
  - Link-state protocol to learn probing results
  - Selecting a good intermediate hop when needed

Deploy multiple small overlay networks, if necessary

7

Advantage: Customizing Packet Delivery

- Recovering from packet loss
  - Packet retransmission
  - Forward error correction
- Quality-of-service differentiation
  - Classify packets based on header bits
  - Schedule packet transmissions based on result
- Incremental deployment of new features
  - Multicast communication (e.g., MBone)
  - IPv6 (e.g., 6Bone)
  - Encryption of packet contents

8

Disadvantage: Traversing Intermediate Nodes

- Processing delay
  - Packets going through multiple software nodes
- Network performance
  - Propagation delay on circuitous path
  - Network congestion from extra load
- Financial cost
  - Bill for traffic going in/out of intermediate node

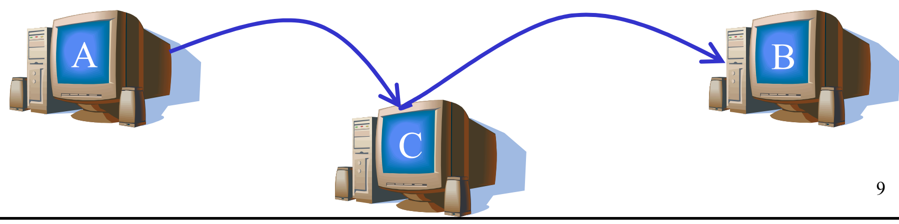

9

Disadvantage: Limitations of Active Probes

- Bandwidth overhead
  - Probe traffic between two nodes
  - Propagating probe results to other nodes
- Limited accuracy of end-to-end probes
  - Available bandwidth of logical link?
  - Losses due to congestion vs. failure?
  - Problem on forward vs. reverse path?
- Limited visibility
  - Logical links may share underlay routers/links
  - May be hard to detect the dependencies

10

Disadvantage: Feedback Effects

- Background traffic
  - Overlay traffic consumes extra resources
  - … at the expense of regular background traffic
  - But, the overlay traffic *does* get out of the way!
- Other overlays
  - Potential competition between multiple overlays
  - E.g., one overlay picks a (longer) alternate path
  - … and extra load causes another overlay to adapt
- Underlying network
  - Overlay network changes the traffic matrix
  - … forcing operators to adapt the underlay routing

Are these effects significant? Any positive effects?

11

Price of Anarchy (Roughgarden&Tardos)

- Worst-case example
  - Two paths from *s* to *d*
  - Total of one unit of load
  - Latency as function of load
- Selfish source routing
  - All traffic through bottom link
  - Mean latency of 1
- Latency-optimal routing
  - Minimize mean latency
  - Set $x = [1/(n+1)]^{1/n}$
  - Mean latency goes to 0

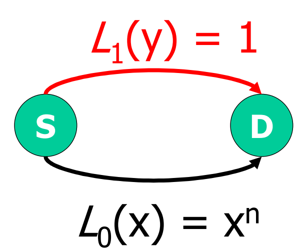

Total load: $x + y = 1$

Mean latency: $x\,L_0(x) + y\,L_1(y)$

12

Internet-like Environments (Qiu et al.)

- Realistic networks
  - Backbone network topologies
  - Link delay (propagation and queuing delay)
  - Routing set to minimize network congestion
- Realistic overlays
  - Small number of overlay nodes (limited flexibility)
  - Overlay paths chosen to minimize latency
- Practice doesn't match the worst-case theory
  - Some tension between the two different metrics
  - But, not anywhere near as bad as the worst case

13

Interaction With Traffic Eng. (Qiu et al)

- Underlay network: traffic engineering
  - Inputs: traffic matrix M_{ij} and physical topology
  - Objective: minimize overall network congestion
  - Output: routing R_{ijl}: fraction of (i,j) traffic on link l
  - Propagation and queuing delay on virtual links

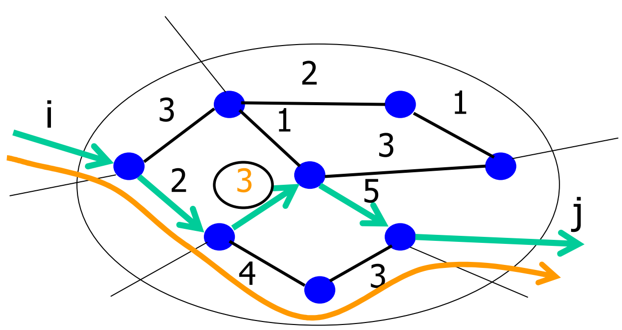

14

Interaction With Traffic Eng. (Qiu et al)

- Overlay network: selects intermediate nodes
  - Inputs: measured delay for each virtual link
  - Objective: minimizing end-to-end latency
  - Output: choice of intermediate nodes for traffic
    - Traffic matrix on the underlay network

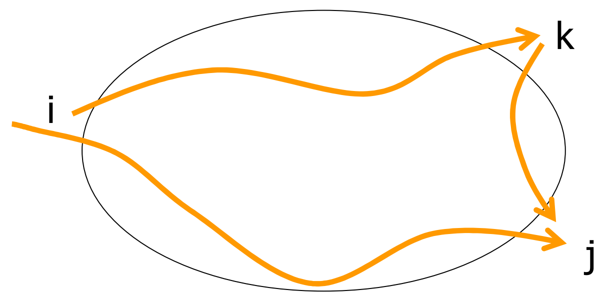

15

Interaction With Traffic Eng. (Qiu et al)

- Underlay network: traffic engineering
  - Inputs: traffic matrix and physical topology
  - Objective: minimize overall network congestion
  - Output: selection of paths in underlay network
    - Propagation and queuing delay on virtual links
- Overlay network: selecting intermediate nodes
  - Inputs: measured delay for each virtual link
  - Objective: minimizing end-to-end latency
  - Output: choice of intermediate nodes for traffic
    - Traffic matrix on the underlay network

16

Interaction TE vs. OSPF Weight Tweaks

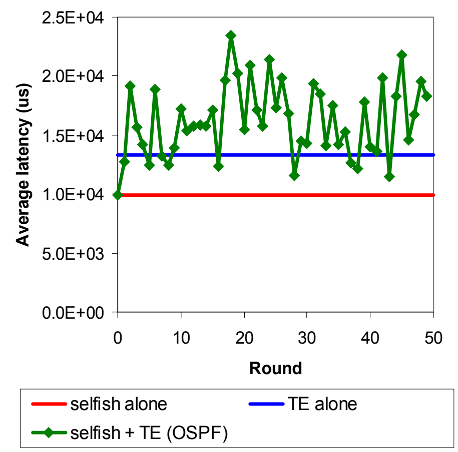

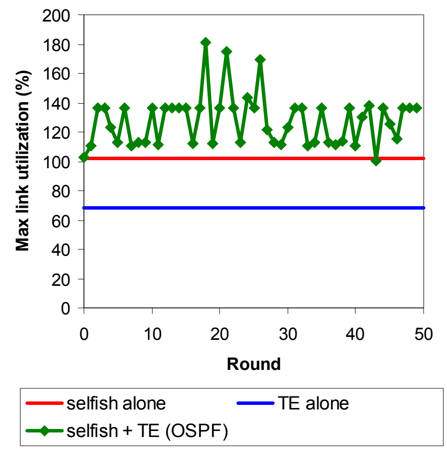

OSPF optimizer interacts poorly with selfish overlays because it only has very coarse-grained control.

17

Interaction TE vs. Multi-commodity Flow

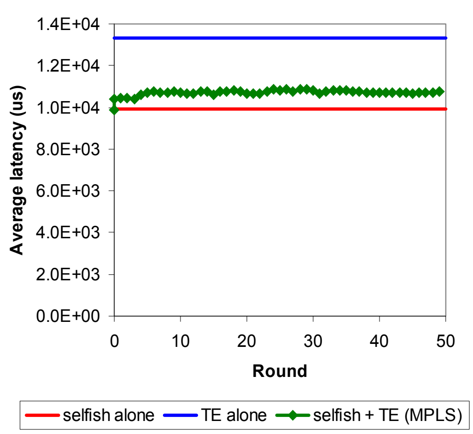

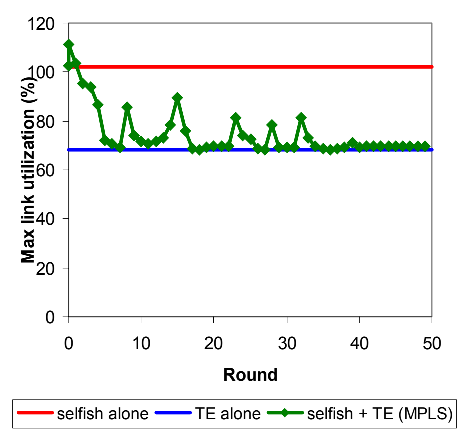

Multi-commodity flow optimizer interacts with selfish overlays much more effectively.

18

History: Bistability in Single Overlay

- Phone network is an overlay
  - Logical link between each pair of switches
  - Phone call put on one-hop path, when possible
  - … and two-hop alternate path otherwise
- Problem: inefficient path assignment
  - Two-hop path for one phone call
  - … stops another call from using direct path
  - … forcing the use of a two-hop alternate path

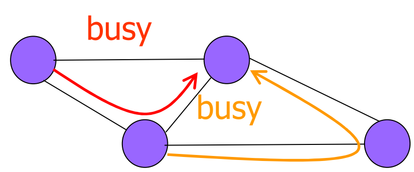

Preventing Inefficient Routes: Trunk Reservation

- Two stable states for the system
  - Mostly one-hop calls with low blocking rate
  - Mostly two-hop calls with high blocking rate
- Making the system stable
  - Reserve a portion of each link for direct calls
  - When link load exceeds threshold…
    - … disallow two-hop paths from using the link
  - Rejects some two-hop calls
    - … to keep some spare capacity for future one-hop calls
- Stability through trunk reservation
  - Single efficient, stable state with right threshold

Should ISPs Fear Overlays, or Favor Them?

- Billing
  - Con: overlays commoditize the network providers
  - Pro: overlay traffic adds traffic subject to billing
- Engineering
  - Con: traffic matrix becomes less predictable
  - Pro: TE less important because overlays can adapt
- Value-added services
  - Con: overlays become the place for new services
  - Pro: ISPs can provide overlay nodes in the core

Do the pros outweigh the cons? Beat 'em, or join 'em?

21

What Could ISPs Do to Help Overlays?

- Better visibility
  - Measurements from the underlying network
- Better control
  - Influence over underlay path selection
- Lower cost
  - Avoiding need to traverse intermediate hosts
- Better joint optimization
  - Underlay adaptation accountings for the overlays

- Idea: ISP could provide an oracle to the P2P

22

ISP Support: More Visibility Via Data

- Underlying topology
  - Dependencies: between paths for virtual links
  - Resource limits: capacity of the underlying links
- Routing-protocol update streams
  - Fast adaptation: BGP may provide early warning
  - Better adaptation: identify location of a problem
- Performance measurement
  - Statistics: per-link delay, loss, and throughput
  - Efficiency: consolidating probes for many overlays

Are there enough incentives to share the data?

23

ISP Support: Direct Control

- Influence over paths for virtual links
  - Guarantees of link/node disjoint paths?
- Control over packet forwarding
  - Overlays allowed to install forwarding entries?
- Warnings about planned changes
  - Prior notification about planned maintenance
  - Rerouting of virtual link before the outage

Are there effective ways to share control?

24

ISP Support: Lower Cost

- Measurement service to overlays
  - Reduce overhead and cost of active probing
- Deflection points inside the network
  - Avoiding packets traversing intermediate hosts
  - Encapsulation to deflect packet through landmark
- Hosting of servers for running overlays
  - Enhanced form of a hosting service
  - Servers running directly in the ISP's PoPs

Reasonable business models and technical solutions?

ISP Support: Joint Optimization

- Changes in traffic-engineering practices
  - Overlay-friendly traffic engineering?
  - Ways to ensure stability and optimality?
  - Or, just let the overlays manage the traffic?
- Cross-domain cooperation
  - Management of paths and traffic across ASes
  - Virtual links that span two (or more) ISPs
  - Ways to ensure stability and optimality?
  - Service-level agreements spanning ASes?

Reasonable optimization or game-theory models?

Discussion

- Are overlay networks a good idea?
  - Just a hack to avoid changing the underlay?
  - What if we could "fix" the underlying network?
  - Would we still have a need for overlay networks?
- Should we have overlay-friendly underlays?
  - Or underlay-friendly overlays, or both?
  - Visibility, control, economics, efficiency, …
  - Or, are the two systems inherently at odds?
- What about interactions between overlays?
  - Cooperate to reduce measurement cost and prevent suboptimality and instability?
  - Compete because that's the way life works?

27

Conclusions

- Overlays
  - Enables innovation in routing and forwarding
  - … without changing the underlying network
- Interaction effects
  - With background traffic
  - With other overlays
  - With traffic engineering
- Ongoing work
  - Right interplay between underlay and overlay?

28