CSE-561
Evolving Network Protocols

Yatin Chawathe
yatin.chawathe@intel.com

Evolution of the Internet

- Internet started with best-effort point-to-point communication
- TCP for reliable delivery
- Extensions
  - IP Multicast for many-to-many communication
  - TCP-SACK to optimize TCP performance
  - Mobile IP for seamless mobility

Why haven't they been deployed yet

- Multicast need router upgrades
  - Router vendors must implement them
  - ISPs must deploy the new routers
- Many technical and business concerns
  - Protocol complexity: keep the network simple
  - Network heterogeneity
  - Layering complex transport protocols

What about end-to-end protocols

- TCP-SACK is technically superior
  - Yet not widely deployed
- Backward compatibility
- Upgrade hell
  - Windows 98/98/NT still account for 7.5% of deployed systems

Examples of successful evolution

- Domain Name System
  - Replacement for *hosts.txt*
- TCP Congestion Control
  - [Jacobson '88]
- Cell-phone networks
  - Upgrading from analog to digital systems

An alternative approach

- Keep the network unchanged
  - Layer new protocols on top
- Application-layer multicast
  - An alternative evolutionary approach
  - Instead of IP-layer multicast, deploy it entirely at the end-points
  - E.g., End-system Multicast, Scattercast, Overcast

Application-layer multicast

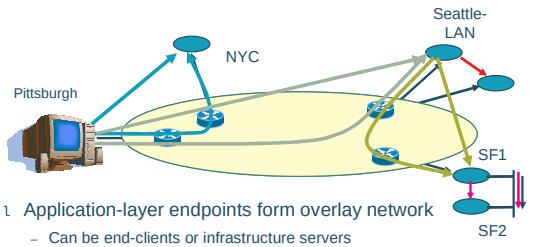

- Application-layer endpoints form overlay network
  - Can be end-clients or infrastructure servers
- Customizable for app-specific needs

Other overlay networks

- RON: Resilient Overlay Network
  - Move routing control toward end-systems
  - Route using overlays; bypass ISP policies
- Akamai: Content delivery network
  - Intelligent distribution of content

Why are overlays attractive

- Deploying new network services is hard
  - The standards challenge
  - Upgrade inertia
- Get the service provider out of the loop
  - Makes deployment "trivial"
  - The overlay IS the infrastructure

i3: A general-purpose routing overlay

- Generalized rendezvous/routing service
  - Sender sends packets to logical id
  - Receiver inserts "trigger" indicating interest in the logical id
- Can be used to implement multicast, anycast, mobility, service composition and more

Protocol description

- IDs have *m* bits
- Matching algorithm
  - Must match at least *k* bits
  - Pick longest prefix match from the IDs that pass above test
  - Forward to all matching triggers
  - Chain of forwarding using stacks of identifiers

Usage

- Multicast
  - All receivers subscribe to the same ID
  - In practice, build a hierarchy of IDs
- Anycast
  - All anycast receivers pick the same first *k* bits for their ID and random (*m*-*k*) bits
  - Sender picks an ID whose first *k* bits match the anycast group ID
- Mobility
  - Trivial: receiver updates its trigger when it moves

Service composition

- Sender-driven or receiver-driven
- Use stacks of identifiers
- Can be used to implement heterogeneous multicast

Discussion

- Enables new services by decoupling addressing from routing
  - Similar example is Intentional Naming System
- But, is it really trivial to deploy?
  - Benefits only when enough people start using it
  - Who will control/manage/pay for the infrastructure?

i3-like overlays can cause problems

- Service level agreements
  - Routing overlays violate SLAs
- Traffic engineering
  - Becomes harder in the face of overlays
- Churn
  - Potentially dynamic membership
  - Management overhead may overwhelm the network

Should network providers care?

- Providing just connectivity is a losing proposition
- Overlays enable new services ⇒ potential for more $$$
- Improved performance, reduced traffic in core ⇒ potential for saving $$$
- Control over monitoring and engineering the overlay