ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Information-Centric Networks

Section # 7.2: Evolved Addressing & Forwarding

Instructor: George Xylomenos

Department: Informatics

European Union
European Social Fund
OPERATIONAL PROGRAMME EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Funding

- These educational materials have been developed as part of the instructors educational tasks.
- The **"Athens University of Economics and Business Open Courses"** project only funded the reformatting of these educational materials.
- The project is being implemented as part of the Operational Program "Instruction and Lifelong Learning" and is co-financed by the European Union (European Social Fund) and national funds.

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Licencing

- These educational materials are subject to a Creative Commons License.

Week 7 / Paper 2

- NIRA: A New Inter-Domain Routing Architecture
  - Xiaowei Yang, David Clark, Arthur W. Berger
  - IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 15, NO. 4, AUGUST 2007
- Main point
  - Users choose ISPs but ISPs choose routes
  - What if users could choose provider level routes?
    - How do you discover routes?
    - How do you represent routes?
    - How do you switch routes quickly?
    - How are providers compensated?
  - NIRA represents routes as a sender and receiver part
    - Each part is represented as a single address

Introduction

- Consider routes at the AS level
  - Users select ISPs only
  - ISPs interconnect independently
- Why have users control routes?
  - With cable against DSL local competition is very slim
  - Route selection introduces competition at the top level
    - Gives backbone ISPs an incentives to invest
- BGP selected routes are not always the best
  - For almost 80% of paths, better ones can be found
  - Even more opportunities for multihomed hosts
  - Users know whether they prefer (say) latency or bandwidth
- NIRA provides route selection at the domain level
  - More manageable than the router level

Design overview

- Basic concepts
  - A user can only select routes for which he pays for
  - Route consists of sender, core and receiver part
    - Source and destination part represented as a single address
- Design rationale
  - Users need to discover failure free routes
  - Routes must be encoded into packets
  - Providers must be compensated
    - Assume bilateral contracts between providers
    - Contracts are typically customer-provider or peer-to-peer
    - Tier-1 providers do not purchase transit from others
    - The Internet core consists of the Tier-1 providers

Design overview

- Route discovery
  - Users can only see their providers, recursively up to the core
    - Including the peering connections outside the core
  - This is the up-graph of the user
    - Represented as one route per domain and one link between domains
  - Topology information propagation protocol (TIPP)
  - TIPP path-vector component
    - Advertises reachability towards the core
    - Tier-1 providers advertise themselves, customers attach themselves
  - TIPP link-state component
    - Advertises network dynamics within provider hierarchy
  - A sender combines an up-graph and a reverse up-graph
    - Valley-free routes: upward, horizontal, downward

Design overview

- Efficient route representation
  - NIRA encodes a path through an up-graph into a single address
    - Both sender and receiver fit into single addresses
  - Each Tier-1 provider obtains a globally unique address prefix
    - This is subdivided to its customers, recursively
    - The final address encodes all the providers it is using
  - NIRA uses 128 bit IPv6 addresses
    - 96 bits encode the up-graph and 32 bits a host in the ISP
    - Could instead use a sequence of IPv4 addresses
  - Peering links use a private address space
    - They are also recursively allocated to customers
  - Each host eventually obtains a set of addresses
    - Each address encodes a path to the core or to a peering domain

Design overview

- Bootstrap communication
  - How do you find out the receiver's addresses?
    - Need to select a path for the receiver too!
  - NRLS: maps names to route segments
    - Similar to DNS but returns multiple up-graphs
  - Hosts are notified by TIPP about network changes
    - They may then need to notify NLRS about their routes
- Handling route failures
  - TIPP notifies the sender but not the receiver
  - If a route is unavailable, the routers return ICMP errors
    - Local errors are masked in domain level paths
    - Inter-domain errors are passed to the sender for a decision
    - The sender consults the NLRS again to choose a new route

Design overview

- Choosing routes
  - A user agent runs on each user's computer
  - The agent combines sender and receiver parts
  - Depending on preferences it chooses a combination
    - Subsequent packets can be used to switch to another route
- Forwarding
  - The up-graphs and reverse up-graphs are specified
  - The route through the core is not specified
    - Tier-1 providers retain control of these routes
    - Users are not exposed to the dense backbone connectivity
    - Each Tier-1 provider needs to advertise a single prefix
  - Any ISP (not only Tier-1's) can decide to join the core
    - It simply needs to obtain a global prefix

TIPP

- TIPP runs between domains but not in the core
  - Separate address and topology propagation
  - Propagating address information
    - A provider announces address prefixes to customers
    - Customers recursively propagate these announcements
  - Propagating topology information
    - Link-state protocol with policy controls
  - Scope enforcement: limit what neighbors know about customers
  - Information hiding: limit what neighbors know about neighbors
  - Uses the Shortest Path Topology Algorithm (SPTA)
    - Computationally more expensive than OSPF or IS-IS
    - But easier to resolve inconsistencies between different messages
    - Should be sufficient for the small scale of the upgraphs

Forwarding

- What is the next hop towards a destination address?
  - Three tables are maintained at each router
    - Uphill: points to the provider that allocated each prefix
    - Downhill: points to the customer that received each prefix
    - Bridge: points to the neighbor allocated with each private prefix
  - Separate from other routing tables (for core routers)
  - Lookup destination in downhill table
  - Lookup source in uphill table
  - Special entries in uphill table
    - Routing: forward through the core
    - Bridge: forward via bridge table (to peer)
  - Special entries in downhill table
    - Blackhole: drop packet (customer is disconnected)
    - Self: forward packet inside the domain

Evaluation

- Amount of state obtained from TIPP
  - The up-graph can grow exponentially
    - Provider hierarchy is not fully known, but inferred
  - Statistics for 90% of domains
    - Less than 20 addresses prefixes
    - Less than 30 link records
    - Less than 100 forwarding entries
- Message overhead and convergence speed of TIPP
  - Less than 1 sec to converge after link failure/recovery
  - Less than 2 messages per link failure/recovery
- Setup latency due to reactive failure detection (ICMP)
  - 80% of connections need a round trip
  - 99% need three round trips

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

End of Section # 7.2

Course: Information-Centric Networks, **Section # 7.2: Evolved Addressing & Forwarding**

Instructor: George Xylomenos, **Department:** Informatics

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND