ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Information-Centric Networks

Section # 4.1: Routing Issues

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

NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Co- financed by Greece and the European Union

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.

Licencing

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

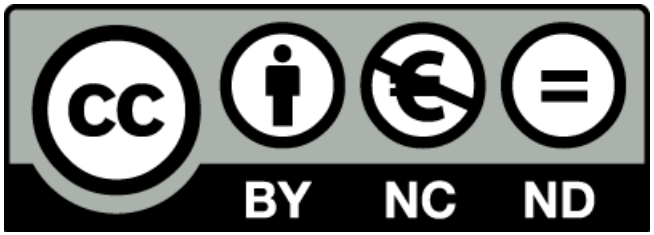

Week 4 / Paper 1

- Open issues in Interdomain Routing: a survey
  - Marcelo Yannuzzi, Xavier Masip-Bruin, Olivier Bonaventure
  - IEEE Network, Nov.-Dec. 2005, vol. 19, no. 6, pp. 49 – 56
- Main point
  - There are many challenges in interdomain routing
  - Relationships between challenges
  - Review of problems and proposed solutions
  - Reasons for non-adoption of solutions

Introduction

- Limitations of BGP
  - Hard to replace due to widespread deployment
  - Growth of Internet strains the protocol
- Routing among distinct domains
  - Independent management of domains
  - Active competition between providers
- Basics of interdomain routing
  - Autonomous System (AS): a distinct network on the Internet
    - Managed by a single authority
    - Common internal routing policy
    - Intradomain routing via IS-IS or OSPF
  - More than 20000 ASes in the Internet

Basics of interdomain routing

- Types of ASes
  - Single homed stub ASes
    - The “leaves” of the network
  - Multi homed stub ASes
    - Load balancing or failure resiliency
  - Transit ASes
- Types of AS relationships
  - Costumer-provider
  - Peer-peer
- The AS hierarchy
  - Tier-1 (no upstream provider) composing the internet core
  - Tier-2 are customers of Tier-1 and providers of Tier-2
  - ASes also use peering links to directly exchange traffic

Basics of interdomain routing

- Interior and exterior BGP
  - BGP routers exchange reachability information via eBGP
  - This information is distributed internally via iBGP
- Route discovery and selection
  - BGP routers advertise AS level paths to IP prefixes
    - There is no global topology view
    - The path is abstracted into AS numbers
  - Route selection is heavily policy influenced
    1. Choose the route with the highest local preference
    2. Choose the route with the shortest AS path
    3. Choose the route with the lowest Multi Exit Discriminator
    4. Prefer the route with the lowest IGP metric
    5. Run tie-breaking rules

Research challenges

- Limited traffic engineering capabilities
  - Only reachability is advertised
    - No way to propagate multiple routes
  - Inbound traffic is hard to control with BGP
    - It is the sender and not the receiver that chooses paths
    - This is due to the uncoordinated routing on the Internet
- Lack of QoS support
  - Many ASes have deployed differentiated services
  - These allow coarse grained intradomain QoS
  - But BGP does not allow interdomain QoS
  - Most ASes prefer over-provisioning instead of QoS

Research challenges

- How has the Internet changed?
  - AS numbers have swelled
  - Connections per AS have increased
  - Additional applications
    - Application requirements are not reflected into BGP
- Slow convergence and chattiness
  - BGP messages: Open, Update, Notification, Keepalive
  - A failure causes large amounts of BGP updates
    - Path exploration takes place until convergence
  - BGP routers wait for MRAI before updating a route
  - BGP routers often employ route flap damping
    - Ignore routes that change too often
  - Improve stability at the expense of convergence

Research challenges

- Slow convergence and chattiness
  - Proposal: faster propagation of updates due to failures
    - Two methods, ghost flushing and reporting the root cause
    - Limits path exploration
    - Requires BGP modification to indicate a failure cause
    - Very hard to pinpoint a failure due to route aggregation
    - Route disaggregation impacts scalability, so there is a tradeoff
  - Proposal: infer source of failures by correlating data
    - Requires multiple vantage points and offline processing
  - Each solution adversely affects some objectives
    - Solutions: root-cause, MRAI timer, flap damping, aggregation
    - Objectives: scalability, convergence, message load

Research challenges

- Scalability problems due to multihoming
  - Many stub ASes are multihomed
    - Resilience and load balancing
  - Each such AS has multiple IP prefixes (one per provider)
    - Each prefix is advertised to all providers
    - Different paths are advertised to indicate route preference
    - Depending on provider aggregation rules, there may be problems
    - Eventually, disaggregation may be required to achieve policy goals
    - Also, each provider can only aggregate its own prefixes
    - All these lead to even more routes being advertised
  - Proposal: route filtering
    - Avoid propagating very long prefixes (very specific routes)
    - This inhibits load balancing (it hides some routes)
  - What we really need is better support for multihoming

Research challenges

- Expresiveness and safety of policies
  - Each domain independently chooses its policies
  - The result is suboptimal due to lack of coordination
    - Global routing anomalies arise
    - Global divergence of routing policies
  - BGP policies are not that expressive
    - Rich enough to express intricate routing policies
    - Not rich enough to allow discovery of problems
  - Each AS does not want to disclose its internal details
    - Many problems are hidden inside the AS's network

Research challenges

- Robustness of BGP sessions
  - BGP routers communicate via TCP
  - Network congestion can lead to failures
    - Hard to resolve congestion due to routing problems
  - Need to distinguish routing messages from ordinary traffic
- Security issues
  - Spoofed TCP RSTs can bring down BGP sessions
    - Also spoofed TCP messages can be inserted
    - Filter spoofed packets
    - Use authentication between BGP routers
  - BGP advertisements are not authenticated
    - S-BGP certifies the validity of routes by signatures
    - Processing cost and need for a PKI

Research challenges

- Lack of multipath routing
  - BGP routers only advertise a single path per prefix
    - Even if they have received many alternatives
  - BGP routers may only use a single path per prefix
    - Some implementations use many for load balancing
  - Proposed extensions raise scalability concerns
- Transit through an AS: iBGP issues
  - A large AS has trouble propagating routes inside it
    - Ideally all internal routers should communicate with each other
    - This not scalable for large ASes
    - Different routers may treat traffic differently
  - Encapsulation: guarantees that packets will use preferred routes
    - Operates between ingress and egress BGP routers

End of Section # 4.1

Course: Information-Centric Networks, **Section # 4.1: Routing Issues**

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