

Networking: more than the sum of its pieces

Anja Feldmann

Fakultät für Informatik
Technische Universität München

The Internet?

For Joe Doe:

- **Social phenomena:** changing/redefining
 - Human to human communication
 - Human to computer interactions
 - Business/military/society conduct

For Internet researchers:

- **Technical phenomena:** to be
 - Observed, characterized, understood, improved
 - Possibilities and capabilities
 - Limitations and dangers
 - Scale

The Internet: Advantages

- Highly engineered structure
 - Well specified and documented (RFCs)
- Unique measurement capabilities
 - In theory unlimited access to data :-)
- Exploiting available data
 - Use **invariants** not details
 - Use **network wide** data sets
 - Consider **emerging** phenomena
 - Take advantage of **structural** models

The Internet: Challenges

- Is experiencing explosive growth
- An immense moving target
- Heterogeneity any which way you look
- Complex user behavior and traffic dynamics
- Highly interacting systems
 - **Temporal**: between hosts and network
 - **Spatial**: among different connections
 - **Vertical**: across different networking layers

Challenge 1:

performance debugging

- Locating reasons for performance problem
 - Top down analysis: Searching for a cause?
 - The protocol itself
 - Missperforming link
 - Bad application design
 - Access network or client problem
 - Missconfiguration
 - DDoS attack
 - Protocol interactions
 - Scaling limits: "success disaster"

Challenge 2: service performance

- How to derive **service performance from component performance**
 - Bottom up view: searching for the relationship?
 - High utilization \Rightarrow bad Web performance
 - Link failure \Rightarrow bad/no connectivity
 - DDoS attack \Rightarrow bad response time
 - Result:
 - Never, sometimes, eventually, occasionally, maybe

Tools


- **Instrumentation and analysis**
 - Integrate measurements into the design process
 - Collect data at a variety of different locations/levels
 - Find invariants and correlate various datasets
- **Simulation**
 - Build a mirror world for "what if" studies
 - Verify explanations
- **Test-Labs**
 - Incorporate variability
 - Provide an friendly/unfriendly environment

Example: Internet Routing


Border Gateway Protocol (BGP)

- **Task:**
 - Routing between Autonomous Systems (AS)
 - Propagation of reachability information
 - Path vector protocol (to prevent routing loops)
- **Challenges:**
 - Support for routing policies (economic realities)
 - Stability vs. fast updates
 - Scale (Internet wide)
 - Limited resource (e.g. router CPUs and memory)

BGP update propagation




BGP update propagation



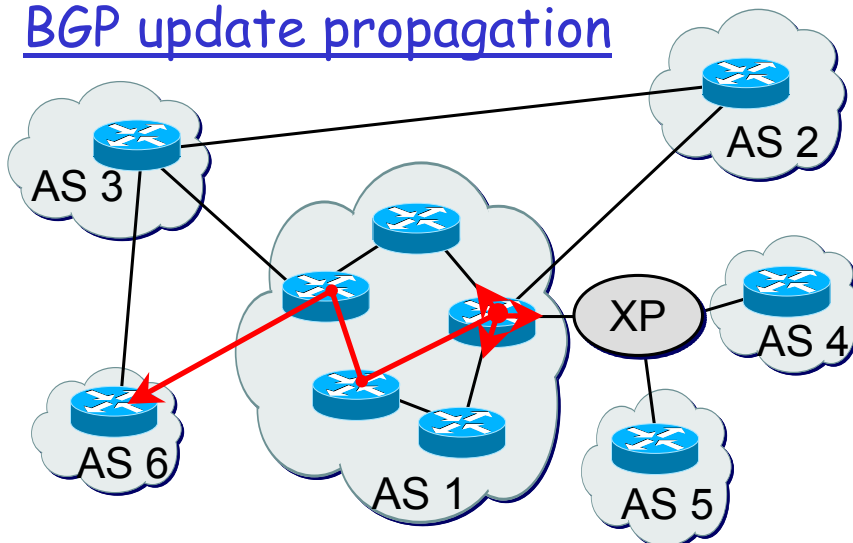
AS 2 issues an update

BGP update propagation



AS 6 may announce to reach AS 2 via AS 3

BGP update propagation

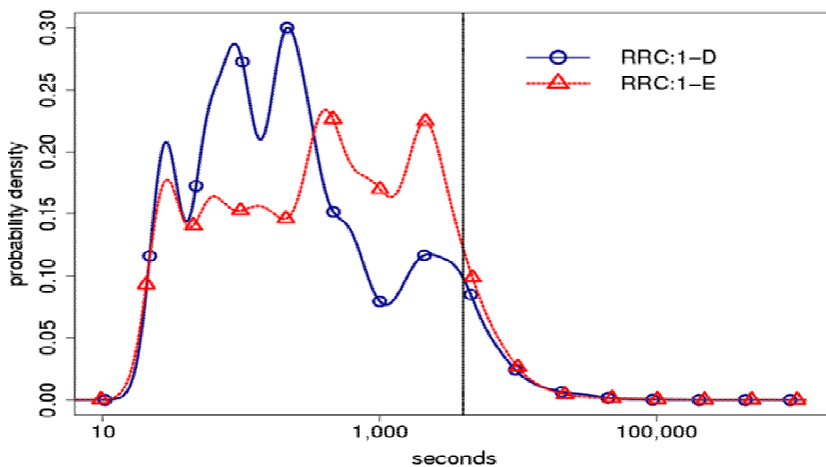


AS 6 may prefer to reach AS 2 via AS 1

Instrumentation and analysis

- BGP Data
 - German ISP
 - RIPE
 - Route-View
- Other data source
 - Packet Data
 - Flow Data
 - Reverse proxy logs
 - ...

Routing burst duration



Simulation

- Tools: SSFnet and ns-2
- Capabilities
 - Routing: BGP/OSPF/MPLS
 - Workload: Web traffic
 - Network topologies: AS and ISP topologies
- Simulation execution
 - Validation of simulation model
 - OSPF validation
 - From measurements to simulation
 - Sensibility studies
 - Interactions OSPF/BGP? Does RED help? Switching vs. Routing? MPLS vs. OSPF traffic engineering?

Test-Lab

- Internet in a Lab: **Why?**
 - Current lab tests:
 - Component tests
 - Functionality of single feature
 - Maybe limited feature interaction
 - Current real test:
 - **Deployment in the Internet!**
- Internet Test-Lab: **Ingredients**
 - Realistic workloads
 - Temporal and address variability
 - Routing

Exploring in an Internet Test-Lab

- BGP
 - Test interactions IGP vs. BGP
 - Test BGP's scalability via future workloads
- Routing protocols
 - Interactions of routing protocols: intra/inter/mobile
 - New routing protocols
- Security
 - Intrusion detection
 - Firewalls
- Measurement and analysis
 - Test instrumentation
 - Explore data correlation methods
 - Find relationship between services and components
 - Try performance debugging and management

Outlook and challenges

- Fortunate **confluence** of several research tracks
 - Network-wide measurements
 - Large-scale simulators
 - Test environments
 - Multiscale analysis/visualization analysis
- Enables **side-by-side** comparisons
 - Measurements from "real world"/"mirror world"/"test world"
 - For discovery, validation, and replication
- **Scientific Challenges**
 - Avoid drowning in measured data
 - Provide foundations for performance evaluation of large-scale networks
 - Performance debugging

Networking: more than the sum of its pieces

BGP workload ingredients

- Cause of routing instability
 - Instability creator
- Effect of routing instability
 - Instability bursts
- Baseline for prefix structure/hierarchy
 - Prefix forest
- AS topology and peering policies
 - AS path properties
- Correlations within instability
 - Attribute changes

Some projects

- ❑ Traffic flows: Mice vs. Elephants
 - Origin and impact
- ❑ Traffic characterization
 - Chat traffic, multimedia, DNS, ...
- ❑ Estimating Inter-domain traffic flows
 - Impact of routing updates on actual traffic flows
- ❑ Network Intrusion Detection
 - Distributed IDS using Netflow data
- ❑ Routing convergence
 - Understanding the delays in BGP convergence
- ❑ What if studies (e.g. using network simulator)
 - MPLS vs. OSPF traffic engineering
 - RED a performance booster or buster

Example data sets

❑ RIPE's RRC00:

- Jan 14, 2002 01:00 - Jan 20, 2002 01:10

