



Network Diagnosis

Jennifer Rexford

Fall 2010 (TTh 1:30-2:50 in COS 302)

COS 561: Advanced Computer Networks

<http://www.cs.princeton.edu/courses/archive/fall10/cos561/>

Networks Break (In Weird Ways)



- Bad things happen
 - Reliability: link, router, firewall, DNS server, Web server
 - Performance : congestion, long paths, overloaded server
- Not straight-forward
 - Selective failure (e.g., MTU mismatch, server replica)
 - Application problems (e.g., receive window)
 - Short-lived problems (e.g., convergence, incast)
 - Problems in other domains (e.g., downstream loss)
 - Unexpected causes (e.g., hot weather, software bugs)
- Yet, we can approach diagnosis in a rigorous way

Detecting and Diagnosing Problems




- Do nothing
 - Rely on the network to adapt to failures
 - E.g., dynamic routing protocols, TCP congestion control
 - Doesn't help in detecting and fixing persistent problems
- Direct observation
 - Detailed measurement to observe problem directly
 - E.g., route monitoring, fault logs, ...
 - High overhead and works only for problems you know
- Inference
 - Infer the root causes from indirect observations
 - Common attributes of the observed failures, and uncommon attributes of the things that don't fail

Fault Localization in a Single Domain




- Failures are often correlated
 - Links connected to same router or traversing same fiber
 - Routers using same power supply or software version
- Inputs
 - Shared risk link groups
 - Group of failed components
- Output
 - Most likely root cause
- Practical challenge: dirty data
 - Lost failure-reporting messages
 - Inaccurate model of risk groups




Fault Localization in Path-Vector Routing

- Routing changes are correlated
 - A single link failure causes multiple routing changes
 - ... for all paths that traverse the failed edge
- Inputs
 - No knowledge of the underlying topology
 - Path changes viewed from several vantage points
- Output
 - Link(s) responsible for the changes
- Practical challenges
 - Incomplete data, multiple failures
 - Complex routing policies



Link-Level Parameter Estimation

- Path performance is correlated
 - Path performance is affected by each link in the path
 - Many paths have (some) common links
- Inputs
 - Network topology and routes
 - Path-level observations of packet loss, delay, ...
- Outputs
 - Estimate of link parameters
- Practical challenges: noise
 - Time-varying link properties



Path-Level Traffic Intensity Estimation



- Link loads are correlated
 - Each ingress-egress pair imparts load on all the links along a path
- Inputs
 - Network topology and routes
 - Total traffic load on each link
- Outputs
 - Offered load for each ingress-egress pair
- Practical challenge
 - Under-constrained inference problem

