Logically-Centralized Control

COS 597E: Software Defined Networking

Jennifer Rexford
Princeton University
MW 11:00am-12:20pm

Software Defined Networking

Control/Data Separation

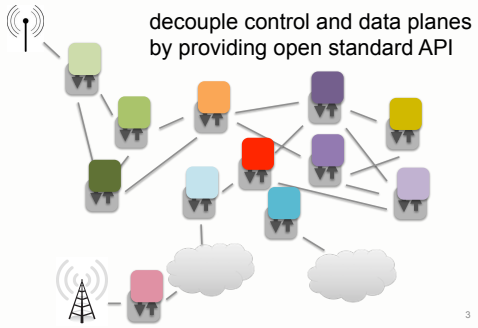

(Logically) Centralized Controller

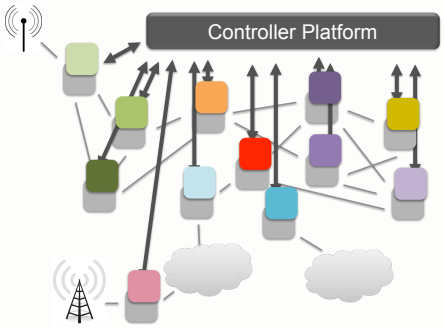

Protocols ➔ Applications

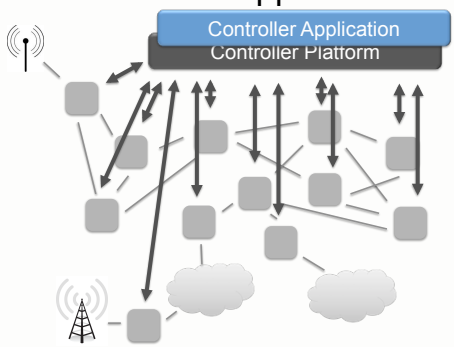

Data, Control, and Management Planes

Timescales

	Data	Control	Management
Time-scale	Packet (nsec)	Event (10 msec to sec)	Human (min to hours)
Tasks	Forwarding, buffering, filtering, scheduling	Routing, circuit set-up	Analysis, configuration
Location	Line-card hardware	Router software	Humans or scripts

7

Data and Control Planes

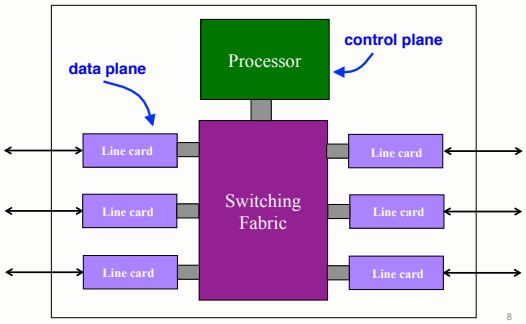

8

Data Plane

- Streaming algorithms on packets
  - Matching on some bits
  - Perform some actions
- Wide range of functionality
  - Forwarding
  - Access control
  - Mapping header fields
  - Traffic monitoring
  - Buffering and marking
  - Shaping and scheduling
  - Deep packet inspection

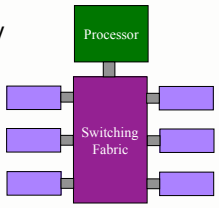

9

Switch: Match on Destination MAC

- MAC addresses are location independent
  - Assigned by the vendor of the interface card
  - Cannot be aggregated across hosts in LAN

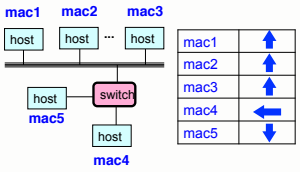

10

Router: Match on IP Prefix

- IP addresses grouped into common subnets
  - Allocated by ICANN, regional registries, ISPs, and within individual organizations
  - Variable-length prefix identified by a mask length

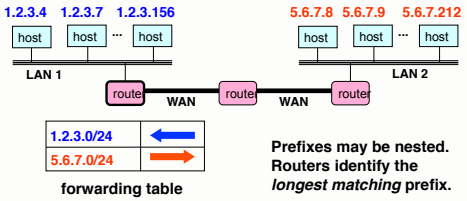

Prefixes may be nested. Routers identify the *longest matching* prefix.

11

Forwarding vs. Routing

- Forwarding: data plane
  - Directing a data packet to an outgoing link
  - Individual router *using* a forwarding table
- Routing: control plane
  - Computing paths the packets will follow
  - Routers talking amongst themselves
  - Individual router *creating* a forwarding table

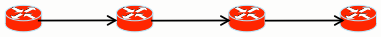

12

Example: Shortest-Path Routing

- Compute: *path costs* to all nodes
  - From a source u to all other nodes
  - Cost of the path through each link
  - Next hop along least-cost path to s

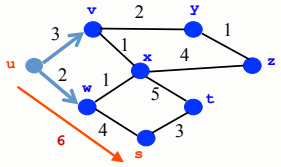

	link
v	(u,v)
w	(u,w)
x	(u,w)
y	(u,v)
z	(u,v)
s	(u,w)
t	(u,w)

13

Distributed Control Plane

- Link-state routing: OSPF, IS-IS
  - Flood the entire topology to all nodes
  - Each node computes shortest paths
  - Dijkstra's algorithm

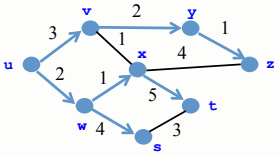

	link
v	(u,v)
w	(u,w)
x	(u,w)
y	(u,v)
z	(u,v)
s	(u,w)
t	(u,w)

14

Distributed Control Plane

- Distance-vector routing: RIP, EIGRP
  - Each node computes path cost
  - … based on each neighbors' path cost
  - Bellman-Ford algorithm

$$d_u(z) = \min\{c(u,v) + d_v(z), \; c(u,w) + d_w(z)\}$$

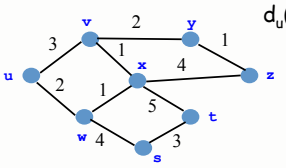

15

Traffic Engineering Problem

- Management plane: setting the weights
  - Inversely proportional to link capacity?
  - Proportional to propagation delay?
  - Network-wide optimization based on traffic?

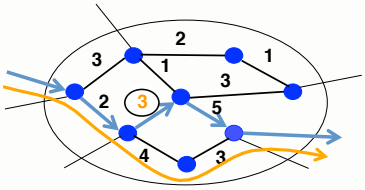

16

Traffic Engineering: Optimization

- Inputs
  - Network topology
  - Link capacities
  - Traffic matrix
- Output
  - Link weights
- Objective
  - Minimize max-utilized link
  - Or, minimize a sum of link congestion

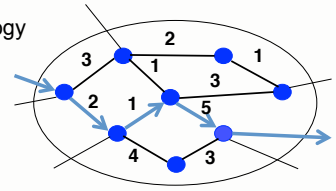

17

Transient Routing Disruptions

- Topology changes
  - Link weight change
  - Node/link failure or recovery
- Routing convergence
  - Nodes temporarily disagree how to route
  - Leading to transient loops and blackholes

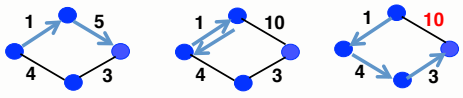

18

Management Plane Challenges

- Indirect control
  - Changing weights instead of paths
  - Complex optimization problem
- Uncoordinated control
  - Cannot control which router updates first
- Interacting protocols and mechanisms
  - Routing and forwarding
  - Naming and addressing
  - Access control
  - Quality of service
  - …

19

Discussing the Readings

20

4D

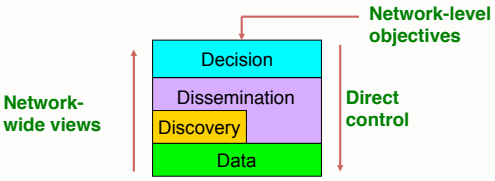

- Decision: all management and control logic
- Dissemination: communicating with routers
- Discovery: topology and traffic monitoring
- Data: packet handling

routers

21

Ethane

- Higher-level policies
  - Defined on principals, not network identifiers
  - Language for specifying policies
- Policy should dictate the paths
  - Controller should select paths based on policy
- Fine-grain control
  - Controller handles first packet of a flow
  - Data plane maintains per-flow state

22

Practical Challenges

- Scalability
  - Decision elements responsible for many routers
- Response time
  - Delays between decision elements and routers
- Reliability
  - Surviving failures of decision elements and routers
- Consistency
  - Ensuring multiple decision elements behave consistently
- Security
  - Network vulnerable to attacks on decision elements
- Interoperability
  - Legacy routers and neighboring domains

23

Next Time

- For Tuesday
  - Complete assignment #1 (MiniNet)
- For Wednesday
  - Review OpenFlow, NOX, and Open vSwitch
  - Read SDN history paper

24