


## Operating Systems and Networks

### Network Lecture 7: Network Layer 2

Adrian Perrig  
 Network Security Group  
 ETH Zürich

## Where we are in the Course

- More fun in the Network Layer!
  - We've covered packet forwarding
  - Now we'll learn about routing



Computer Networks

2

## Routing versus Forwarding

- Forwarding is the process of sending a packet on its way
- Routing is the process of deciding in which direction to send traffic




Computer Networks

3

## Improving on the Spanning Tree

- Spanning tree provides basic connectivity
  - e.g., some path B → C
- Routing uses all links to find "best" paths
  - e.g., use BC, BE, and CE



Computer Networks

4

## Perspective on Bandwidth Allocation

- Routing allocates network bandwidth adapting to failures; other mechanisms used at other timescales


Mechanism	Timescale / Adaptation
Load-sensitive routing	Seconds / Traffic hotspots
Routing	Minutes / Equipment failures
Traffic Engineering	Hours / Network load
Provisioning	Months / Network customers

Computer Networks

5

## Delivery Models

- Different routing used for different delivery models



Computer Networks

6

## Goals of Routing Algorithms

- We want several properties of any routing scheme:

Property	Meaning
Correctness	Finds paths that work
Efficient paths	Uses network bandwidth well
Fair paths	Doesn't starve any nodes
Fast convergence	Recovers quickly after changes
Scalability	Works well as network grows large

Computer Networks

7

## Rules of Routing Algorithms

- Decentralized, distributed setting
  - All nodes are alike; no controller
  - Nodes only know what they learn by exchanging messages with neighbors
  - Nodes operate concurrently
  - May be node/link/message failures



Computer Networks

8

## Topics


- IPv4, IPv6, NATs and all that } Last time
- Shortest path routing
- Distance Vector routing
- Flooding
- Link-state routing
- Equal-cost multi-path
- Inter-domain routing (BGP) } This time

Computer Networks

9

## Shortest Path Routing (§5.2.1-5.2.2)

- Defining "best" paths with link costs
  - These are shortest path routes




Computer Networks

10

## What are "Best" paths anyhow?

- Many possibilities:
  - Latency, avoid circuitous paths
  - Bandwidth, avoid slow links
  - Money, avoid expensive links
  - Hops, to reduce switching
- But only consider topology
  - Ignore workload, e.g., hotspots



Computer Networks

11

## Shortest Paths

We'll approximate "best" by a cost function that captures the factors

- Often call lowest "shortest"

- Assign each link a cost (distance)
- Define best path between each pair of nodes as the path that has the lowest total cost (or is shortest)
- Pick randomly to any break ties

Computer Networks

12

### Shortest Paths (2)

- Find the shortest path A → E
- All links are bidirectional, with equal costs in each direction
  - Can extend model to unequal costs if needed

Computer Networks 13

### Shortest Paths (3)

- ABCE is a shortest path
- $\text{dist}(ABCE) = 4 + 2 + 1 = 7$
- This is less than:
  - $\text{dist}(ABE) = 8$
  - $\text{dist}(ABFE) = 9$
  - $\text{dist}(AE) = 10$
  - $\text{dist}(ABCDE) = 10$

Computer Networks 14

### Shortest Paths (4)

- Optimality property:
  - Subpaths of shortest paths are also shortest paths
- ABCE is a shortest path
  - So are ABC, AB, BCE, BC, CE

Computer Networks 15

### Sink Trees

- Sink tree for a destination is the union of all shortest paths towards the destination
  - Similarly source tree
- Find the sink tree for E

Computer Networks 16

### Sink Trees (2)

- Implications:
  - Only need to use destination to follow shortest paths
  - Each node only need to send to the next hop
- Forwarding table at a node
  - Lists next hop for each destination
  - Routing table may know more

Computer Networks 17


### Computing Shortest Paths with Dijkstra (§5.2.2)

- How to compute shortest path given the network topology
  - With Dijkstra's algorithm

Computer Networks 18

### Edsger W. Dijkstra (1930-2002)

- Famous computer scientist
  - Programming languages
  - Distributed algorithms
  - Program verification
- Dijkstra's algorithm, 1959
  - Single-source shortest paths, given network with non-negative link costs



By Hamilton Richards, CC BY-SA 3.0, via Wikimedia Commons

Computer Networks 19

### Dijkstra's Algorithm

Algorithm:

- Mark all nodes tentative, set distances from source to 0 (zero), and  $\infty$  (infinity) for all other nodes
- While tentative nodes remain:
  - Extract N, a node with lowest distance
  - Add link to N to the shortest path tree
  - Relax the distances of neighbors of N by lowering any better distance estimates

Computer Networks 20

### Dijkstra's Algorithm (2)

- Initialization

We'll compute shortest paths from A

Computer Networks 21

### Dijkstra's Algorithm (3)

- Relax around A

Computer Networks 22

### Dijkstra's Algorithm (4)

- Relax around B

Computer Networks 23

### Dijkstra's Algorithm (5)

- Relax around C

Computer Networks 24

### Dijkstra's Algorithm (6)

- Relax around G (say) Didn't fall ...

Computer Networks 25

### Dijkstra's Algorithm (7)

- Relax around F (say) Relax has no effect

Computer Networks 26

### Dijkstra's Algorithm (8)

- Relax around E

Computer Networks 27

### Dijkstra's Algorithm (9)

- Relax around D

Computer Networks 28

### Dijkstra's Algorithm (10)

- Finally, H ... done

Computer Networks 29


### Dijkstra Comments

- Finds shortest paths in order of increasing distance from source
  - Leverages optimality property
- Runtime depends on efficiency of extracting min-cost node
  - Superlinear in network size (grows fast)
- Gives complete source/sink tree
  - More than needed for forwarding!
  - But requires complete topology

Computer Networks 30

### Distance Vector Routing (§5.2.4)

- How to compute shortest paths in a distributed network
  - The Distance Vector (DV) approach



Computer Networks

31

### Distance Vector Routing

- Simple, early routing approach
  - Used in ARPANET, and RIP (Routing Information Protocol)
- One of two main approaches to routing
  - Distributed version of Bellman-Ford
  - Works, but very slow convergence after some failures
- Link-state algorithms are now typically used in practice
  - More involved, better behavior

Computer Networks

32

### Distance Vector Setting

Each node computes its forwarding table in a distributed setting:

1. Nodes know only the cost to their neighbors; not the topology
2. Nodes can talk only to their neighbors using messages
3. All nodes run the same algorithm concurrently
4. Nodes and links may fail, messages may be lost

Computer Networks

33

### Distance Vector Algorithm

Each node maintains a vector of distances (and next hops) to all destinations


1. Initialize vector with 0 (zero) cost to self,  $\infty$  (infinity) to other destinations
2. Periodically send vector to neighbors
3. Update vector for each destination by selecting the shortest distance heard, after adding cost of neighbor link
  - Use the best neighbor for forwarding

Computer Networks

34

### Distance Vector Example

- Consider a simple network. Each node runs on its own
  - E.g., node A can only talk to nodes B and D




Computer Networks

35

### DV Example (2)

- First exchange, A hears from B, D and finds 1-hop routes
  - A always learns  $\min(B+3, D+7)$



Computer Networks

36

### DV Example (3)

- First exchange for all nodes to find best 1-hop routes
  - E.g., B learns  $\min(A+3, C+6, D+3)$

To	A says	B says	C says	D says
A	0	$\infty$	$\infty$	$\infty$
B	$\infty$	0	$\infty$	$\infty$
C	$\infty$	$\infty$	0	$\infty$
D	$\infty$	$\infty$	$\infty$	0

A learns	B learns	C learns	D learns
Cost Next	Cost Next	Cost Next	Cost Next
0 -- 3 A	$\infty$ -- 7 A	$\infty$ -- 6 C	$\infty$ -- 2 D
3 B 0 --	0 -- 6 B	0 -- 2 C	0 --
$\infty$ -- 6 C	$\infty$ -- 3 D	0 -- 2 C	0 --
7 D 3 D 2	0 --	0 --	0 --

= learned better route

Computer Networks 37

### DV Example (4)

- Second exchange for all nodes to find best 2-hop routes

To	A says	B says	C says	D says
A	0	3	$\infty$	7
B	3	0	6	3
C	$\infty$	6	0	2
D	7	3	2	0

A learns	B learns	C learns	D learns
Cost Next	Cost Next	Cost Next	Cost Next
0 -- 3 A	9 B 6 B	9 B 6 B	6 B
3 B 0 --	5 D 3 B	5 D 3 B	3 B
9 D 5 D	0 -- 2 C	0 -- 2 C	2 C
6 B 3 D	2 D 0 --	2 D 0 --	0 --

= learned better route

Computer Networks 38

### DV Example (5)

- Third exchange for all nodes to find best 3-hop routes

To	A says	B says	C says	D says
A	0	3	9	6
B	3	0	5	3
C	9	5	0	2
D	6	3	2	0

A learns	B learns	C learns	D learns
Cost Next	Cost Next	Cost Next	Cost Next
0 -- 3 A	8 D 6 B	8 D 6 B	6 B
3 B 0 --	5 D 3 B	5 D 3 B	3 B
8 B 5 D	0 -- 2 C	0 -- 2 C	2 C
6 B 3 D	2 D 0 --	2 D 0 --	0 --

= learned better route

Computer Networks 39

### DV Example (5)

- Fourth and subsequent exchanges; converged

To	A says	B says	C says	D says
A	0	3	8	6
B	3	0	5	3
C	8	5	0	2
D	6	3	2	0

A learns	B learns	C learns	D learns
Cost Next	Cost Next	Cost Next	Cost Next
0 -- 3 A	8 D 6 B	8 D 6 B	6 B
3 B 0 --	5 D 3 B	5 D 3 B	3 B
8 B 5 D	0 -- 2 C	0 -- 2 C	2 C
6 B 3 D	2 D 0 --	2 D 0 --	0 --

= learned better route

Computer Networks 40

## Distance Vector Dynamics

- Adding routes:
  - News travels one hop per exchange
- Removing routes
  - When a node fails, no more exchanges, other nodes forget
- But partitions (unreachable nodes in divided network) are a problem
  - “Count to infinity” scenario

Computer Networks 41

## DV Dynamics (2)

- Good news travels quickly, bad news slowly (inferred)

A	B	C	D	E
•	•	•	•	•
1	•	•	•	•
1	2	•	•	•
1	2	3	•	•
1	2	3	4	•

A	B	C	D	E
•	•	•	•	•
1	2	3	4	•
3	2	3	4	•
3	4	3	4	•
5	4	5	4	•
5	6	5	6	•
7	6	7	6	•
7	8	7	8	•
•	•	•	•	•

Desired convergence

“Count to infinity” scenario

Computer Networks 42

### DV Dynamics (3)

- Various heuristics to address
  - e.g., “Split horizon, poison reverse” (Don’t send route back to where you learned it from.)
- But none are very effective
  - Link state now favored in practice in intra-domain (LAN) settings
  - Except when very resource-limited

Computer Networks

43

### RIP (Routing Information Protocol)


- DV protocol with hop count as metric
  - Infinity is 16 hops; limits network size
  - Includes split horizon, poison reverse
- Routers send vectors every 30 secs
  - Runs on top of UDP
  - Timeout in 180 secs to detect failures
- RIPv1 specified in RFC1058 (1988)

Computer Networks

44

### Flooding (§5.2.3)

- How to broadcast a message to all nodes in the network with flooding
  - Simple mechanism, but inefficient



Computer Networks

45

### Flooding


- Rule used at each node:
  - Sends an incoming message on to all other neighbors
  - Remember the message so that it is only sent once over each link (called duplicate suppression)
- Inefficient because one node may receive multiple copies of message

Computer Networks

46

### Flooding (2)

- Consider a flood from A; first reaches B via AB, E via AE




Computer Networks

47

### Flooding (3)

- Next B floods BC, BE, BF, BG, and E floods EB, EC, ED, EF



Computer Networks

48



### Flooding (4)

- C floods CD, CH; D floods DC; F floods FG; G floods GF

F gets another copy

Computer Networks 49

### Flooding (5)

- H has no-one to flood ... and we're done

Each link carries the message, and in at least one direction

Computer Networks 50

### Flooding Details

- Remember message (to stop flood) using source and sequence number
  - Used for duplicate suppression, so same message is only sent once to neighbors
  - So subsequent message (with higher sequence number) will again be flooded
- To make flooding reliable, use ARQ
  - So receiver acknowledges, and sender resends if needed

Computer Networks 51

### Link State Routing (§5.2.5, 5.6.6)

- How to compute shortest paths in a distributed network
  - The Link-State (LS) approach

Computer Networks 52

### Link-State Routing

- One of two approaches to routing
  - Trades more computation than distance vector for better dynamics
- Widely used in practice
  - Used in Internet/ARPANET from 1979
  - Modern networks use OSPF and IS-IS for intra-domain routing

Computer Networks 53

### Link-State Setting

Each node computes their forwarding table in the same distributed setting as distance vector:

1. Node knows only the cost to its neighbors; not the topology
2. Node can talk only to its neighbors using messages
3. Nodes run the same algorithm concurrently
4. Nodes/links may fail, messages may be lost

Computer Networks 54

## Link-State Algorithm

Proceeds in two phases:


1. Nodes **flood** topology in the form of link state packets
  - Each node learns full topology
2. Each node computes its own forwarding table
  - By running Dijkstra (or equivalent)

## Phase 1: Topology Dissemination

- Each node floods **link state packet** (LSP) that describes their portion of the topology

Node E's LSP flooded to A, B, C, D, and F

Seq. #
A
B
C
D
F




## Phase 2: Route Computation

- Each node has full topology
  - By combining all LSPs
- Each node simply runs Dijkstra
  - Some replicated computation, but finds required routes directly
  - Compile forwarding table from sink/source tree
  - That's it folks!

## Forwarding Table

Source Tree for E (from Dijkstra)



E's Forwarding Table

To	Next
A	C
B	C
C	C
D	D
E	--
F	F
G	F
H	C

## Handling Changes


- On change, flood updated LSPs, and re-compute routes
  - E.g., nodes adjacent to failed link or node initiate

B's LSP

Seq. #
A
C
E
F
G

F's LSP

Seq. #
B
E
G



## Handling Changes (2)

- Link failure
  - Both nodes notice, send updated LSPs
  - Link is removed from topology
- Node failure
  - All neighbors notice a link has failed
  - Failed node can't update its own LSP
  - But it is OK: all links to node removed

### Handling Changes (3)

- Addition of a link or node
  - Add LSP of new node to topology
  - Old LSPs are updated with new link
- Additions are the easy case ...

Computer Networks 61

### Link-State Complications

- Things that can go wrong:
  - Seq. number reaches max, or is corrupted
  - Node crashes and loses seq. number
  - Network partitions then heals
- Strategy:
  - Include age on LSPs and forget old information that is not refreshed
- Much of the complexity is due to handling corner cases (as usual!)

Computer Networks 62

### DV/LS Comparison

Goal	Distance Vector	Link-State
Correctness	Distributed Bellman-Ford	Replicated Dijkstra
Efficient paths	Approx. with shortest paths	Approx. with shortest paths
Fair paths	Approx. with shortest paths	Approx. with shortest paths
Fast convergence	Slow – many exchanges	Fast – flood and compute
Scalability	Excellent – storage/compute	Moderate – storage/compute

Computer Networks 63

### IS-IS and OSPF Protocols

- Widely used in large enterprise and ISP networks
  - IS-IS = Intermediate System to Intermediate System
  - OSPF = Open Shortest Path First
- Link-state protocol with many added features
  - E.g., “Areas” for scalability

Computer Networks 64

### Equal-Cost Multi-Path Routing (§5.2.1, 5.6.6)

- More on shortest path routes
  - Allow multiple shortest paths

Computer Networks 65

### Multipath Routing

- Allow multiple routing paths from node to destination be used at once
  - Topology has them for redundancy
  - Using them can improve performance
- Questions:
  - How do we find multiple paths?
  - How do we send traffic along them?

Computer Networks 66

### Equal-Cost Multipath Routes

- One form of multipath routing
  - Extends shortest path model by keeping set if there are ties
- Consider A → E
  - ABE = 4 + 4 = 8
  - ABCE = 4 + 2 + 2 = 8
  - ABCDE = 4 + 2 + 1 + 1 = 8
  - Use them all!

67

### Source "Trees"

- With ECMP, source/sink "tree" is a directed acyclic graph (DAG)
  - Each node has set of next hops
  - Still a compact representation

68

### Source "Trees" (2)

- Find the source "tree" for E
  - Procedure is Dijkstra, simply remember set of next hops
  - Compile forwarding table similarly, may have set of next hops
- Straightforward to extend DV too
  - Just remember set of neighbors

69

### Source "Trees" (3)

Source Tree for E

E's Forwarding Table

Node	Next hops
A	B, C, D
B	B, C, D
C	C, D
D	D
E	--
F	F
G	F
H	C, D

70

### Forwarding with ECMP

- Could randomly pick a next hop for each packet based on destination
  - Balances load, but adds jitter
- Instead, try to send packets from a given source/destination pair on the same path
  - Source/destination pair is called a flow
  - Map flow identifier to single next hop
  - No jitter within flow, but less balanced

71

### Forwarding with ECMP (2)

Multipath routes from F/E to C/H

E's Forwarding Choices

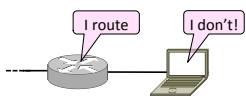
Flow	Possible next hops	Example choice
F → H	C, D	D
F → C	C, D	D
E → H	C, D	C
E → C	C, D	C

Use both paths to get to one destination

72

### Combining Hosts and Routers

- How routing protocols work with IP
  - The Host/Router distinction



Computer Networks 73

### Recap

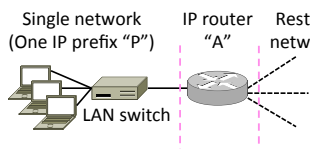
- In the Internet:
  - Hosts on same network have IP addresses in the same IP prefix
  - Hosts just send off-network traffic to the nearest router to handle
  - Routers discover the routes to use
  - Routers use longest prefix matching to send packets to the right next hop

Computer Networks 74

### Host/Router Combination

- Hosts attach to routers as IP prefixes
  - Router needs table to reach all hosts

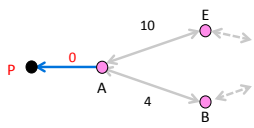
Single network (One IP prefix "P")      IP router "A"      Rest of network



Computer Networks 75

### Network Topology for Routing

- Group hosts under IP prefix connected directly to router
  - One entry for all hosts



Computer Networks 76

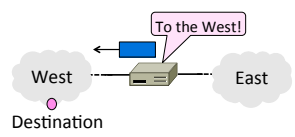
### Network Topology for Routing (2)

- Routing now works as before!
  - Routers advertise IP prefixes for hosts
  - Router addresses are "/32" prefixes
  - Lets all routers find a path to hosts
  - Hosts find by sending to their router

Computer Networks 77

### Hierarchical Routing (§5.2.6)

- How to scale routing with hierarchy in the form of regions
  - Route to regions, not individual nodes



Computer Networks 78

### Internet Growth

Internet Domain Survey Host Count

- At least a billion Internet hosts and growing ...

Source: Internet Systems Consortium ([www.isc.org](http://www.isc.org))

Computer Networks 79

### Internet Routing Growth

- Internet growth translates into routing table growth (even using prefixes) ...

Source: By Mro (Own work), CC-BY-SA-3.0, via Wikimedia Commons

Computer Networks 80

### Impact of Routing Growth

- Forwarding tables grow
  - Larger router memories, may increase lookup time
- Routing messages grow
  - Need to keep all nodes informed of larger topology
- Routing computation grows
  - Shortest path calculations grow faster than the size of the network

Computer Networks 81

### Techniques to Scale Routing

- IP prefixes
  - Route to blocks of hosts

} Last time
- Network hierarchy
  - Route to network regions

} This time
- IP prefix aggregation
  - Combine, and split, prefixes

Computer Networks 82

### Hierarchical Routing

- Introduce a larger routing unit
  - IP prefix (hosts) ← from one host
  - Region, e.g., ISP network
- Route first to the region, then to the IP prefix within the region
  - Hide details within a region from outside of the region

Computer Networks 83

### Hierarchical Routing (2)

Dest.	Line	Hops
1A	—	—
1B	1B	1
1C	1C	1
2A	1B	2
2B	1B	3
2C	1B	3
2D	1B	4
3A	1C	3
3B	1C	2
4A	1C	3
4B	1C	4
4C	1C	4
5A	1C	4
5B	1C	5
5C	1B	5
5D	1C	6
5E	1C	5

Dest.	Line	Hops
1A	—	—
1B	1B	1
1C	1C	1
2	1B	2
3	1C	2
4	1C	3
5	1C	4

Computer Networks 84

### Hierarchical Routing (3)

Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2A	1B	2
2B	1B	3
2C	1B	3
2D	1B	4
3A	1C	3
3B	1C	2
4A	1C	3
4B	1C	4
4C	1C	4
5A	1C	4
5B	1C	5
5C	1B	5
5D	1C	6
5E	1C	5

Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2	1B	2
3	1C	2
4	1C	3
5	1C	4

Computer Networks 85

### Hierarchical Routing (4)

- Penalty is longer paths

Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2A	1B	2
2B	1B	3
2C	1B	3
2D	1B	4
3A	1C	3
3B	1C	2
4A	1C	3
4B	1C	4
4C	1C	4
5A	1C	4
5B	1C	5
5C	1B	5
5D	1C	6
5E	1C	5

Dest.	Line	Hops
1A	--	--
1B	1B	1
1C	1C	1
2	1B	2
3	1C	2
4	1C	3
5	1C	4

Computer Networks 86

### Observations

- Outside a region, nodes have one route to all hosts within the region
  - This gives savings in table size, messages and computation
- However, each node may have a different route to an outside region
  - Routing decisions are still made by individual nodes; there is no single decision made by a region

Computer Networks 87

### IP Prefix Aggregation and Subnets (§5.6.2)

- How to help scale routing by adjusting the size of IP prefixes
  - Split (subnets) and join (aggregation)

Computer Networks 88

### Recall

- IP addresses are allocated in blocks called IP prefixes, e.g., 18.31.0.0/16
  - Hosts on one network in same prefix
- A "/N" prefix has the first N bits fixed and contains  $2^{32-N}$  addresses
  - E.g., "/24"
  - E.g., "/16"

Computer Networks 89

### Key Flexibility

- Routers keep track of prefix lengths
  - Use it for longest prefix matching

Routers can change prefix lengths without affecting hosts

- More specific IP prefix
  - Longer prefix, fewer IP addresses
- Less specific IP prefix
  - Shorter prefix, more IP addresses

Computer Networks 90

### Prefixes and Hierarchy

- IP prefixes already help to scale routing, but we can go further
  - Can use a less specific prefix to name a region made up of several prefixes

Computer Networks 91

### Subnets and Aggregation

Two use cases for adjusting the size of IP prefixes; both reduce routing table size

- Subnets**
  - Internally split one less specific prefix into multiple more specific prefixes
- Aggregation**
  - Externally join multiple more specific prefixes into one large prefix

Computer Networks 92

### Subnets

- Internally split up one IP prefix

Computer Networks 93

### Aggregation

- Externally join multiple separate IP prefixes

Computer Networks 94

### Routing with Multiple Parties (§5.6.7)

- Routing when there are multiple parties, each with their own goals
  - Like Internet routing across ISPs ...

Computer Networks 95

### Structure of the Internet

- Networks (ISPs, CDNs, etc.) group hosts as IP prefixes
- Networks are richly interconnected, often using IXPs

Computer Networks 96



### Internet-wide Routing Issues

- Two problems beyond routing within an individual network
- 1. Scaling to very large networks
  - Techniques of IP prefixes, hierarchy, prefix aggregation
- 2. Incorporating policy decisions
  - Letting different parties choose their routes to suit their own needs

Yikes!

Computer Networks 97

### Effects of Independent Parties

- Each party selects routes to suit its own interests
  - e.g., shortest path in ISP
- What path will be chosen for A2→B1 and B1→A2?
  - What is the best path?

Computer Networks 98

### Effects of Independent Parties (2)

- Selected paths are longer than overall shortest path
  - And asymmetric too!
- This is a consequence of independent goals and decisions, not hierarchy

Computer Networks 99

### Routing Policies

- Capture the goals of different parties – could be anything
  - E.g., Internet2 only carries non-commercial traffic
- Common policies we'll look at:
  - ISPs give TRANSIT service to customers
  - ISPs give PEER service to each other

Computer Networks 100

### Routing Policies – Transit

- One party (customer) gets TRANSIT service from another party (ISP)
  - ISP accepts traffic from customer to deliver to the rest of Internet
  - ISP accepts traffic from the rest of the Internet to delivery to customer
  - Customer pays ISP for the privilege

Computer Networks 101


### Routing Policies – Peer

- Both party (ISPs in example) get PEER service from each other
  - Each ISP accepts traffic from the other ISP only for their customers
  - ISPs do not carry traffic to the rest of the Internet for each other
  - ISPs don't pay each other

Computer Networks 102

### Border Gateway Protocol (§5.6.7)

- How to route with multiple parties, each with their own routing policies
  - BGP computes Internet-wide routes




Computer Networks

103

### Recall

- Internet is made up of independently run networks
- Each network has its own route preferences (policies)




Computer Networks

104

### BGP (Border Gateway Protocol)

- BGP is the protocol that computes interdomain routes in the Internet
  - Path vector, a kind of distance vector



Computer Networks

105


### BGP (2)

- Different parties like ISPs are called AS (Autonomous Systems)
- Border routers of ASes announce BGP routes to each other
- Route announcements contain an IP prefix, path vector, next hop
  - Path vector is list of ASes on the way to the prefix; list is to find loops
- Route announcements move in the opposite direction to traffic

Computer Networks

106

### BGP (3)



Computer Networks

107

### BGP (4)

Policy is implemented in two ways:

1. Border routers of ISP announce paths only to other parties who may use those paths
  - Filter out paths others can't use
2. Border routers of ISP select the best path of the ones they hear in any, non-shortest way

Computer Networks

108

### BGP Example

- AS2 buys TRANSIT service from AS1 and PEER service from AS3

Routing policy:  
TR = Transit  
CU = Customer  
PE = Peer

Computer Networks 109

### BGP Example (2)

- CUSTOMER (other side of TRANSIT): AS2 says [A, (AS2)] to AS1

Routing policy:  
TR = Transit  
CU = Customer  
PE = Peer

Computer Networks 110

### BGP Example (3)

- TRANSIT: AS1 says [B, (AS1, AS3)], [C, (AS1, AS4)] to AS2

Routing policy:  
TR = Transit  
CU = Customer  
PE = Peer

Computer Networks 111

### BGP Example (4)

- PEER: AS2 says [A, (AS2)] to AS3, AS3 says [B, (AS3)] to AS2

Routing policy:  
TR = Transit  
CU = Customer  
PE = Peer

Computer Networks 112

### BGP Example (5)

- AS2 hears one route to C, and two routes to B (chooses AS3!)

Routing policy:  
TR = Transit  
CU = Customer  
PE = Peer

Computer Networks 113

### Closing Thoughts

- Much more beyond basics to explore!
- Policy is a substantial factor
  - Can we be sure independent decisions will yield sensible overall routes?
- Other important factors:
  - Convergence effects
  - How well it scales
  - Integration with routing within ISPs
  - And more ...

Computer Networks 114