CMSC 417

Computer Networks

Prof. Ashok K Agrawala

© 2018 Ashok Agrawala

October 9, 2018 (a)

Message, Segment, Packet, and Frame

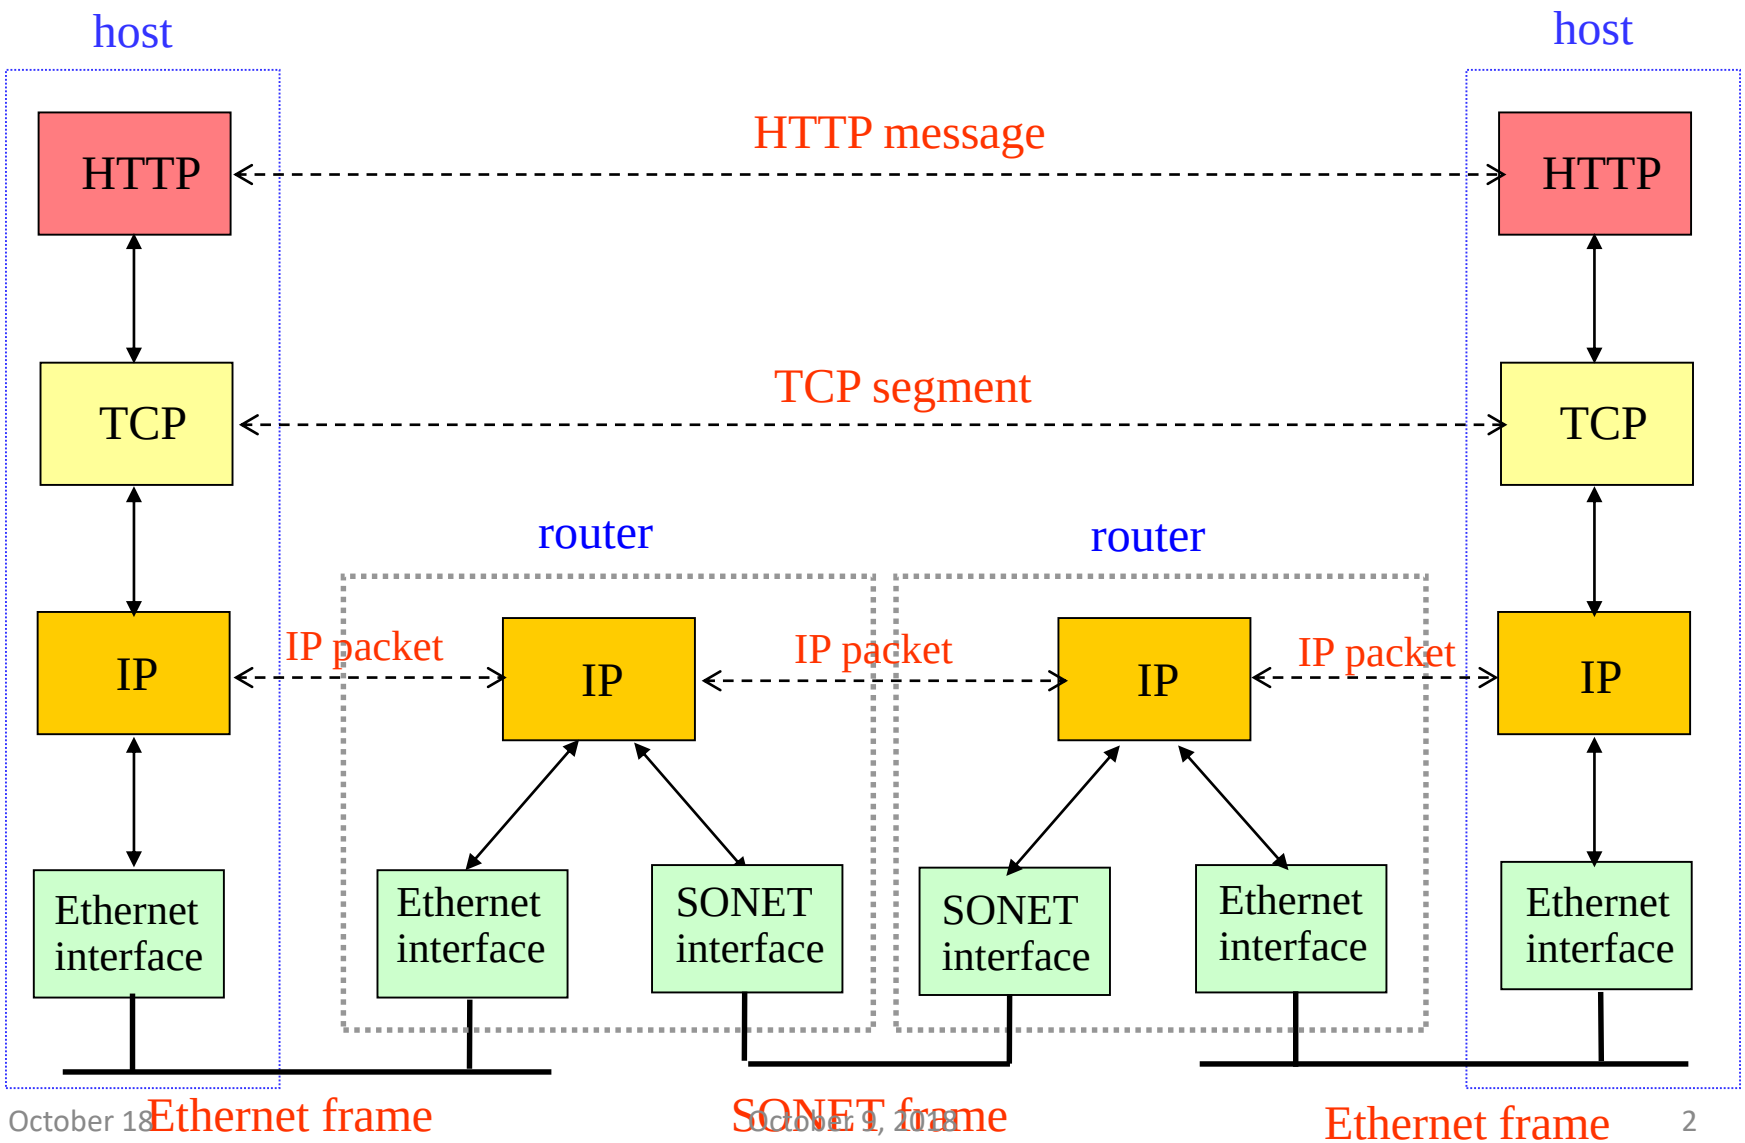

OSPF— Interior Routing Protocol (1)

OSPF computes routes for a single network (e.g., ISP)

- Models network as a graph of weighted edges

Network:

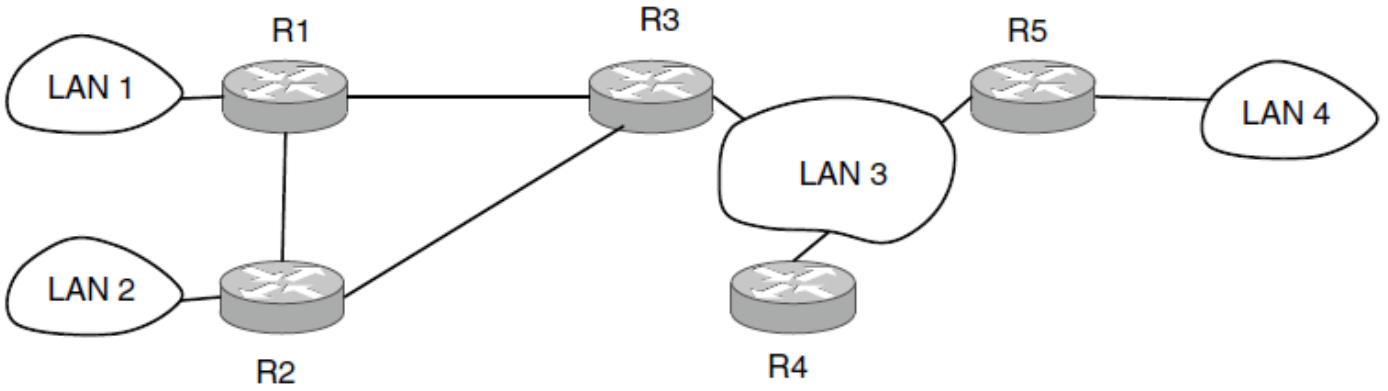

Graph:

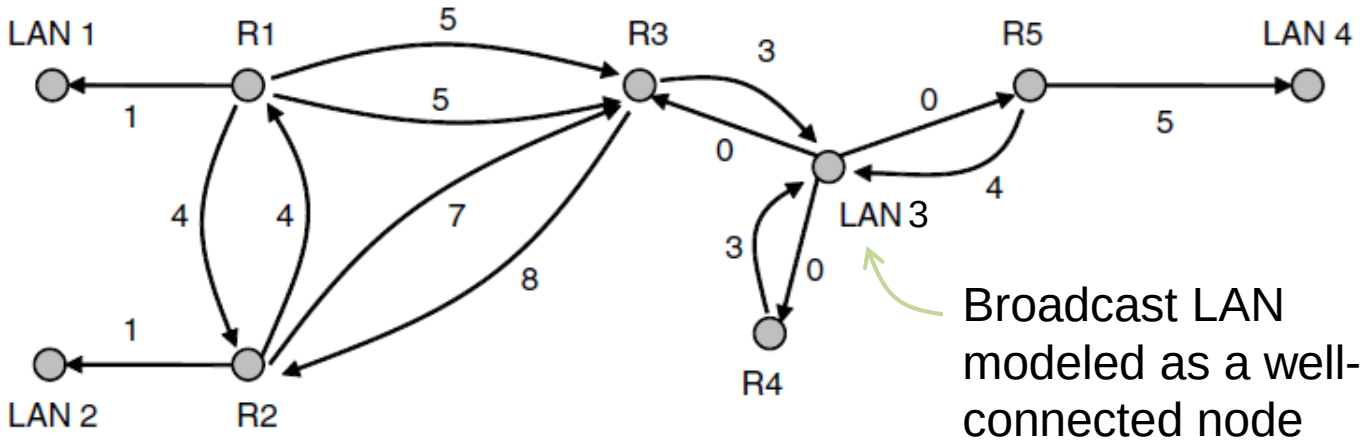

Broadcast LAN modeled as a well-connected node

OSPF— Interior Routing Protocol (2)

OSPF divides one large network (Autonomous System) into areas connected to a backbone area

- Helps to scale; summaries go over area borders

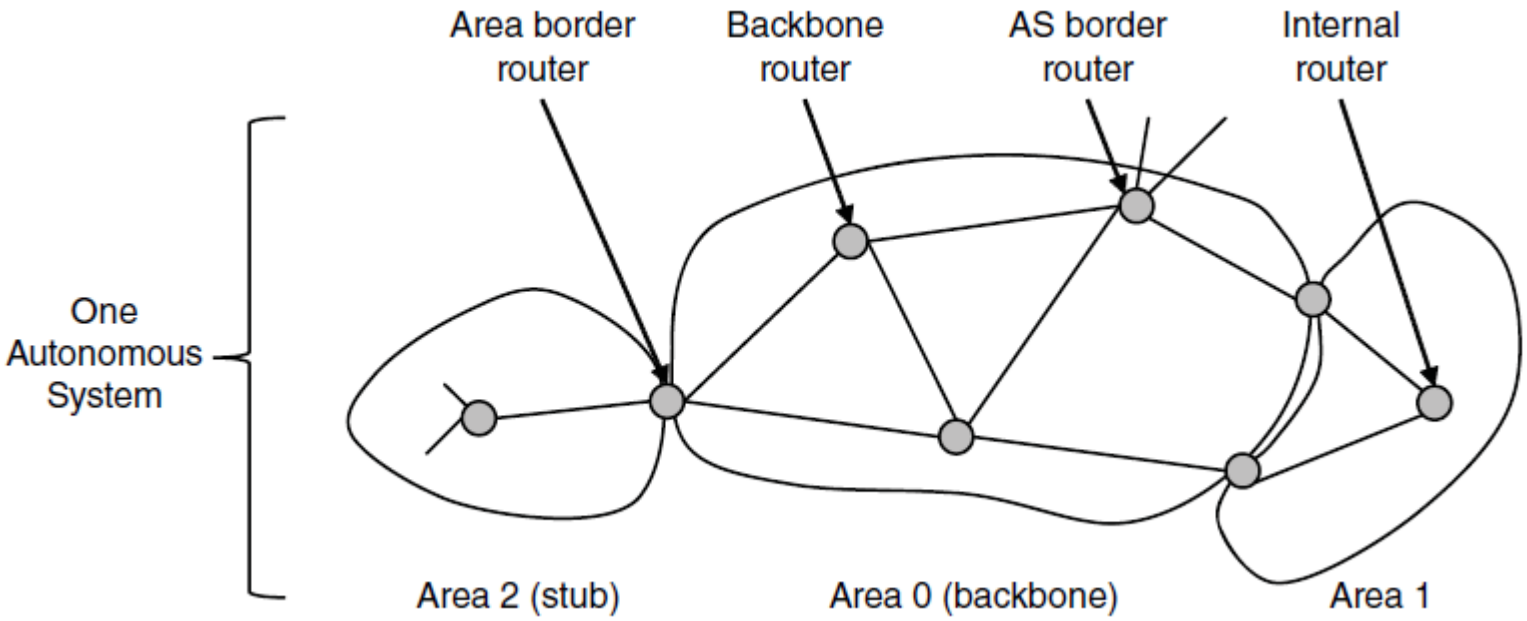

OSPF— Interior Routing Protocol (3)

OSPF (Open Shortest Path First) is link-state routing:

- Uses messages below to reliably flood topology
- Then runs Dijkstra to compute routes

Message type	Description
Hello	Used to discover who the neighbors are
Link state update	Provides the sender's costs to its neighbors
Link state ack	Acknowledges link state update
Database description	Announces which updates the sender has
Link state request	Requests information from the partner

Physical Layout

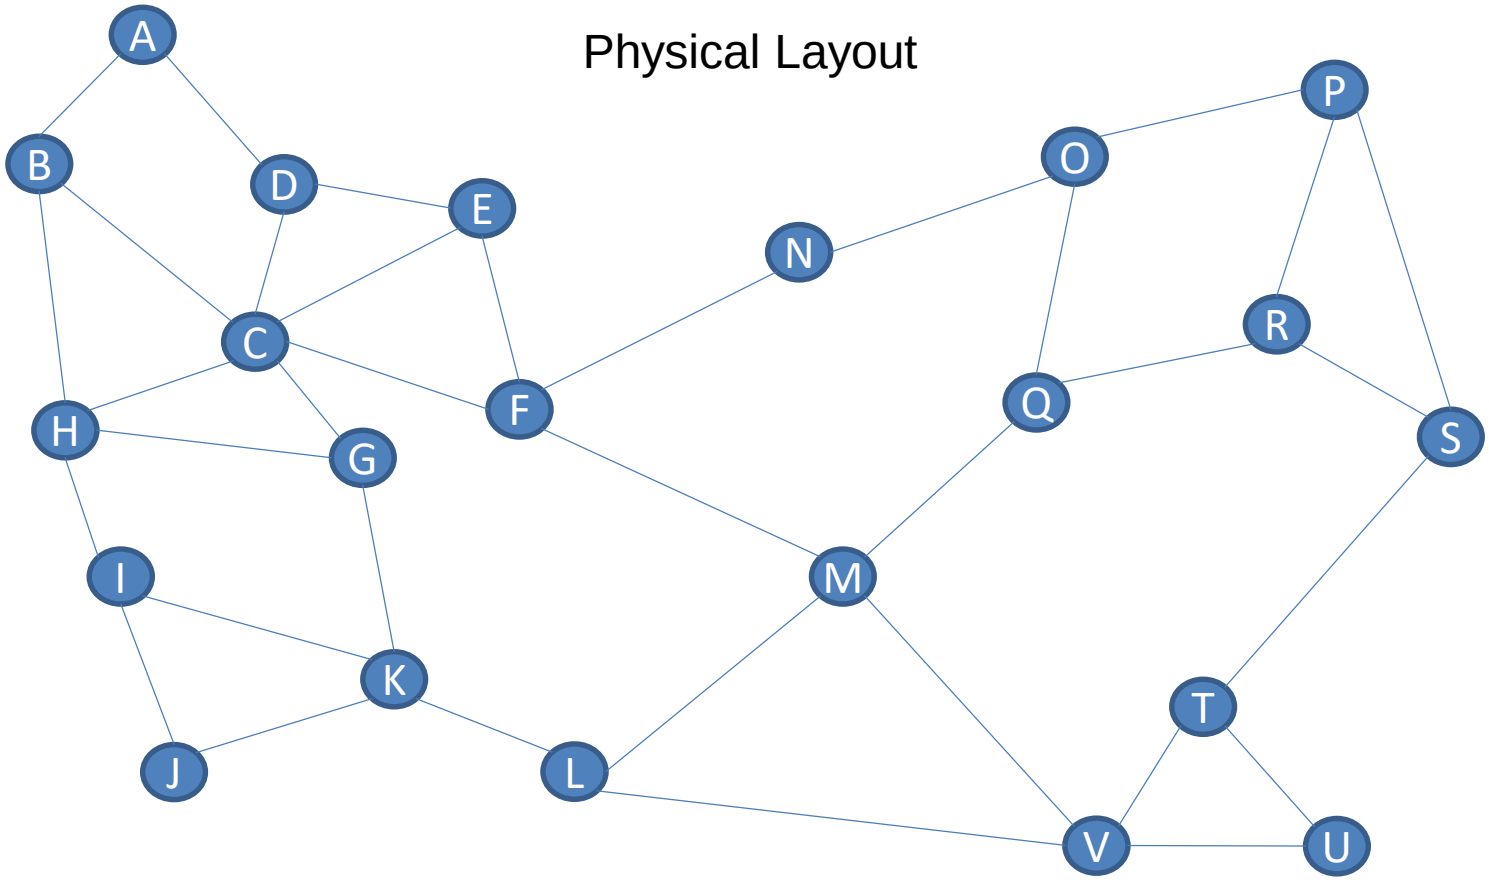

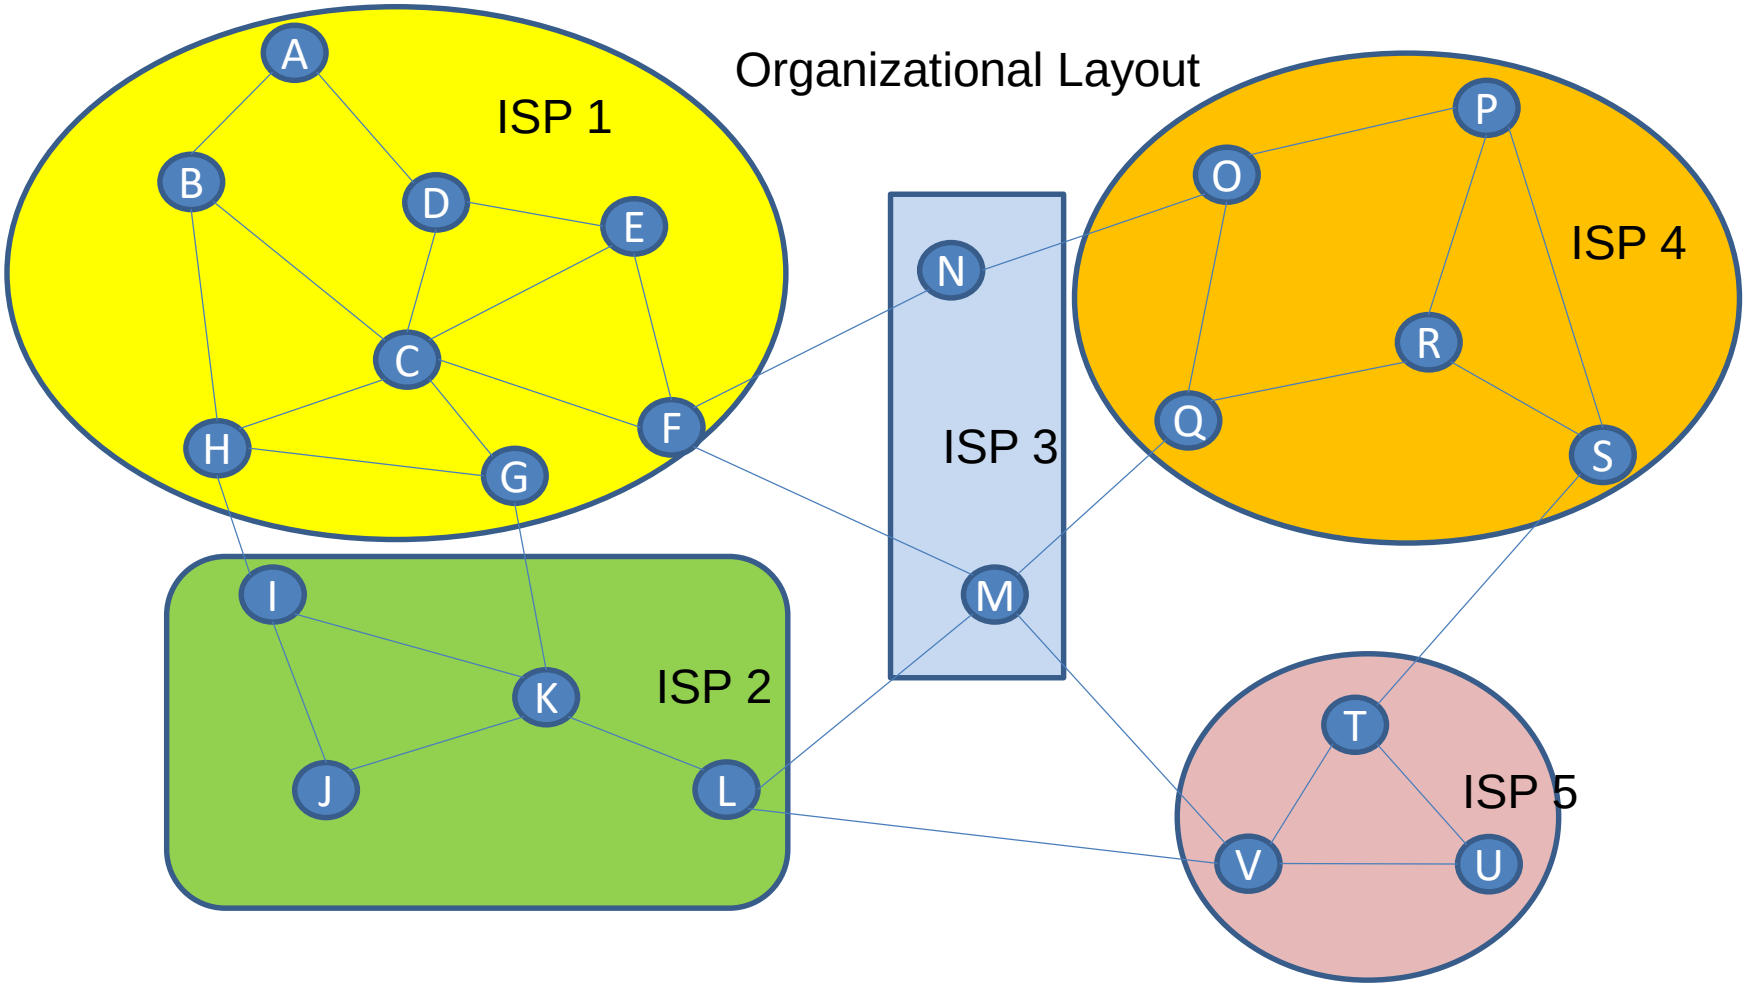
Organizational Layout
ISP 1
A
B
D
E
C
H
G
F
ISP 2
I
K
J
L
ISP 3
N
M
ISP 4
P
O
R
Q
S
ISP 5
T
V
U

What Is BGP?

- Border Gateway Protocol BGP-4
- The de-facto interdomain routing protocol
- BGP enables policy in routing:
  - Which information gets advertised and how
- BGP is a Distance Vector like protocol
- Within an AS, Internal Gateway Protocol (IGP or I-BGP)

How A BGP graph Looks Like

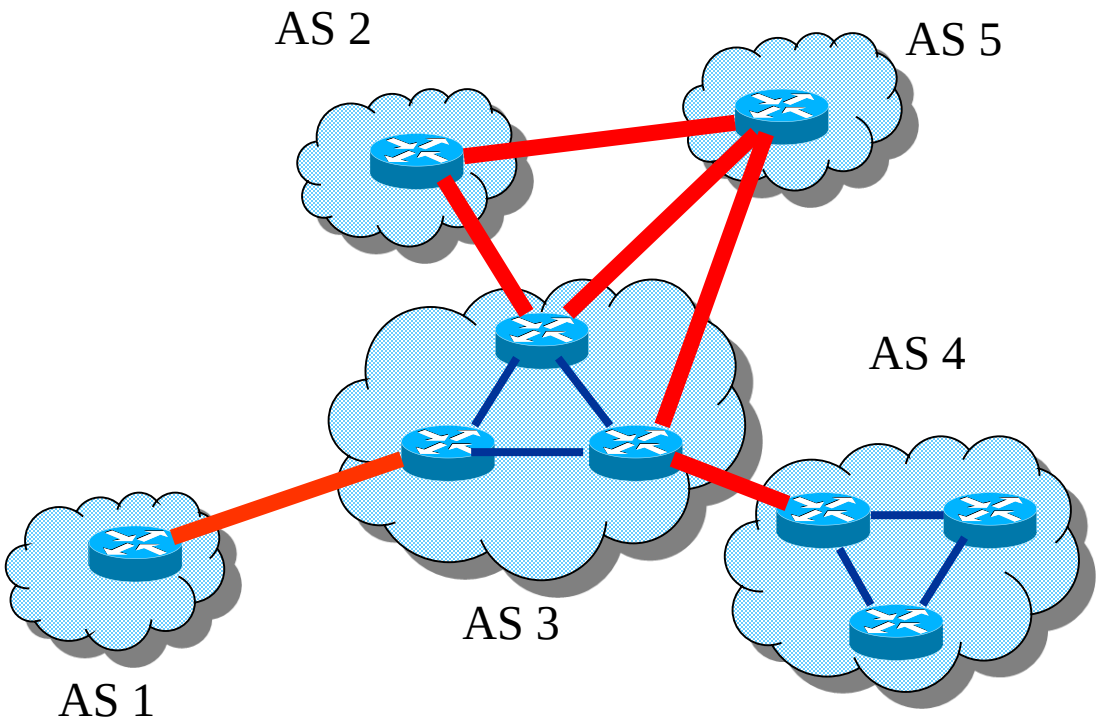

- Each AS has designated BGP routers
- BGP routers of an AS communicate internally with another protocol (IGP)

What is different with BGP?

- BGP goal: enable business relationships
- Opts for: flexibility, scalability
- Performance optimization is secondary

Some Basic Numbers

- ~20,000 Autonomous Systems approx.
  - Corporate Networks
  - ISP Internal Networks
  - National Service Providers
- Identified by ASN a 16 bit value
  - Assigned by IANA
- Superlinear growth

IP Addresses and Prefixes

- IP addresses have 32 bits: 4 octets of bits (IPv4)
- A prefix is a group of IP addresses
- 128.32.101.5 is an IP address (32 bits)
- 128.32.0.0/16 is a prefix of the 16 first bits:
  - 128.32.0.0 – 128.32.255.255 (2^16 addresses)
- 128.32.4.0/24 is a prefix of the 24 first bits - longer

Advertising Routing Information

- Each AS advertises what it can reach from each BGP router
- Policies I: filter what you advertise
- Policies II: filter from what you hear advertised
- Build up a BGP routing table
  - Remember which prefix you hear from which link

What Does a Routing Table Look Like?

Prefix	Origin AS	Path
128.32.0.0/16	123	14 56 123
	123	34 101 203 123
128.32.101.0/24	15	50 15 15

- Origin AS “owns” the address
- Routing tables can have peculiarities

Route Advertising

- Distance Vector style protocol
- Hear advertisements: IP prefix, AS-path
- Filter if desired (i.e. ignore)
- Append yourself: IP prefix, myAS+AS-path
- Forward to appropriate ASs

Basic AS relationships

- Customer – Provider
  - Customer pays Provider for service
  - The Customer is always right
- Peer to Peer: mutual cooperation
  - Ex. MCI and AT&T
- Sibling-Sibling
  - Ex. AT&T research and AT&T wireless

The Internet as a Directed Graph

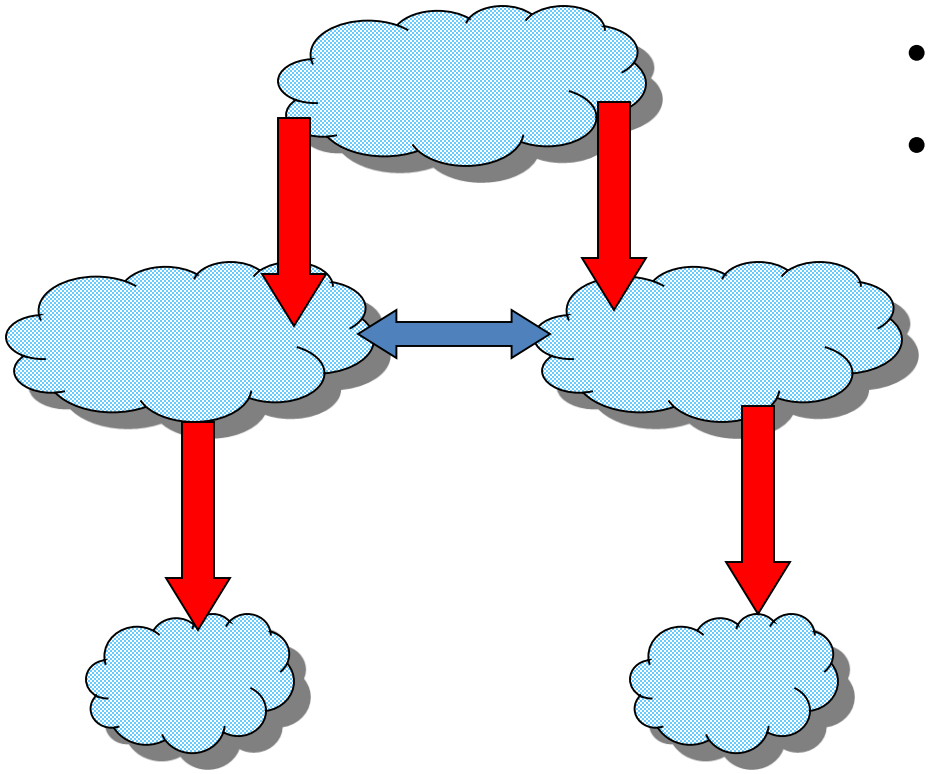

- Every edge is bidirectional
- Business relationships are represented

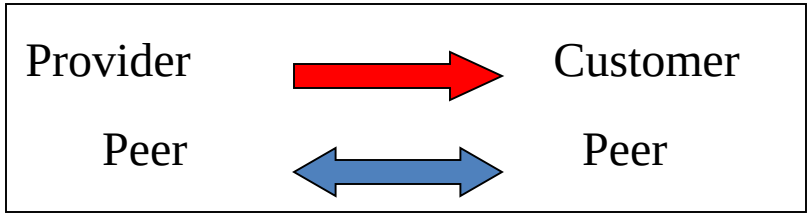

The Initial Idea

- Data flows between customers-providers
- Top level providers are peers
  - They exchange information to ensure connectivity
- What can possibly go wrong?

And then came the rain…

- Thousands of ASs
- Complicated relationships
- Multiple providers for one AS!!
  - Multihoming
- Traffic engineering
  - I want to use multiple paths and load balance

AS Relationships

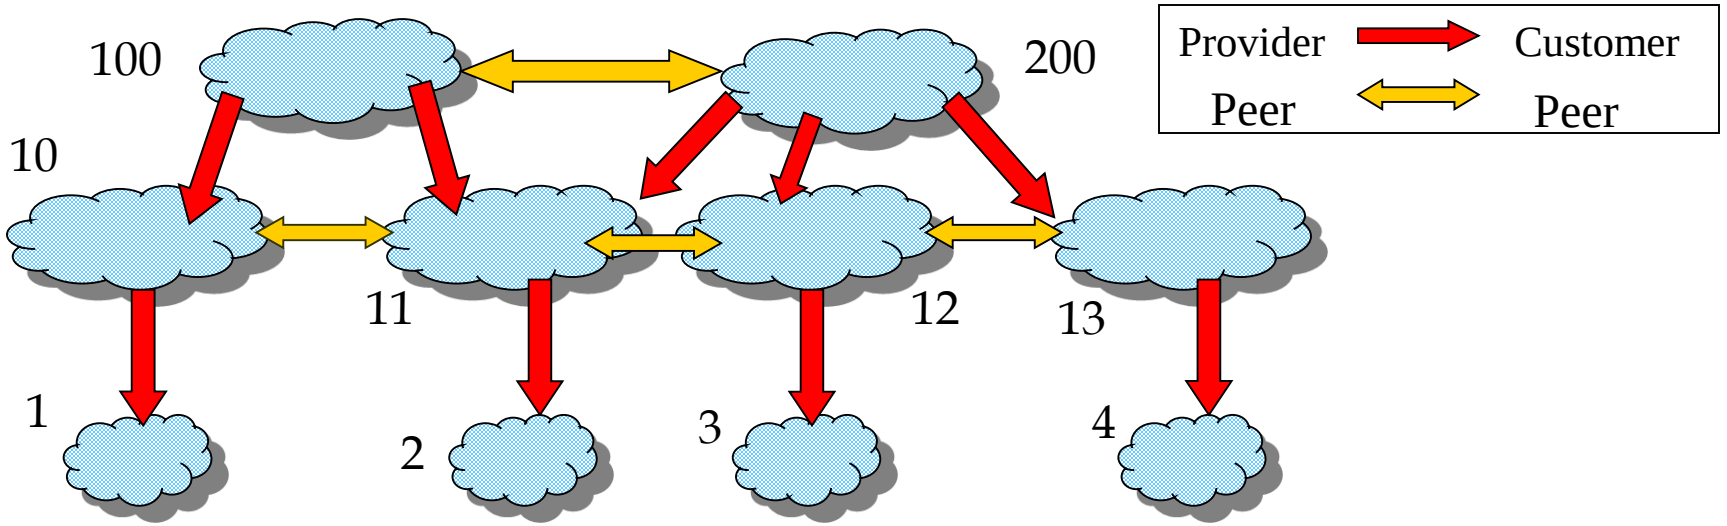

- Customer – Provider: customer pays and is always right
- Peer to Peer: Exchange traffic only between their customers
- Sibling-Sibling: Exchange traffic at will

The Rules of BGP Routing

- Transit traffic: traffic that does not go to my customers (or their customers)
- A provider carries any traffic to or from a customer
- Peers exchange traffic only if between their customers

How BGP Policy Restricts Routing

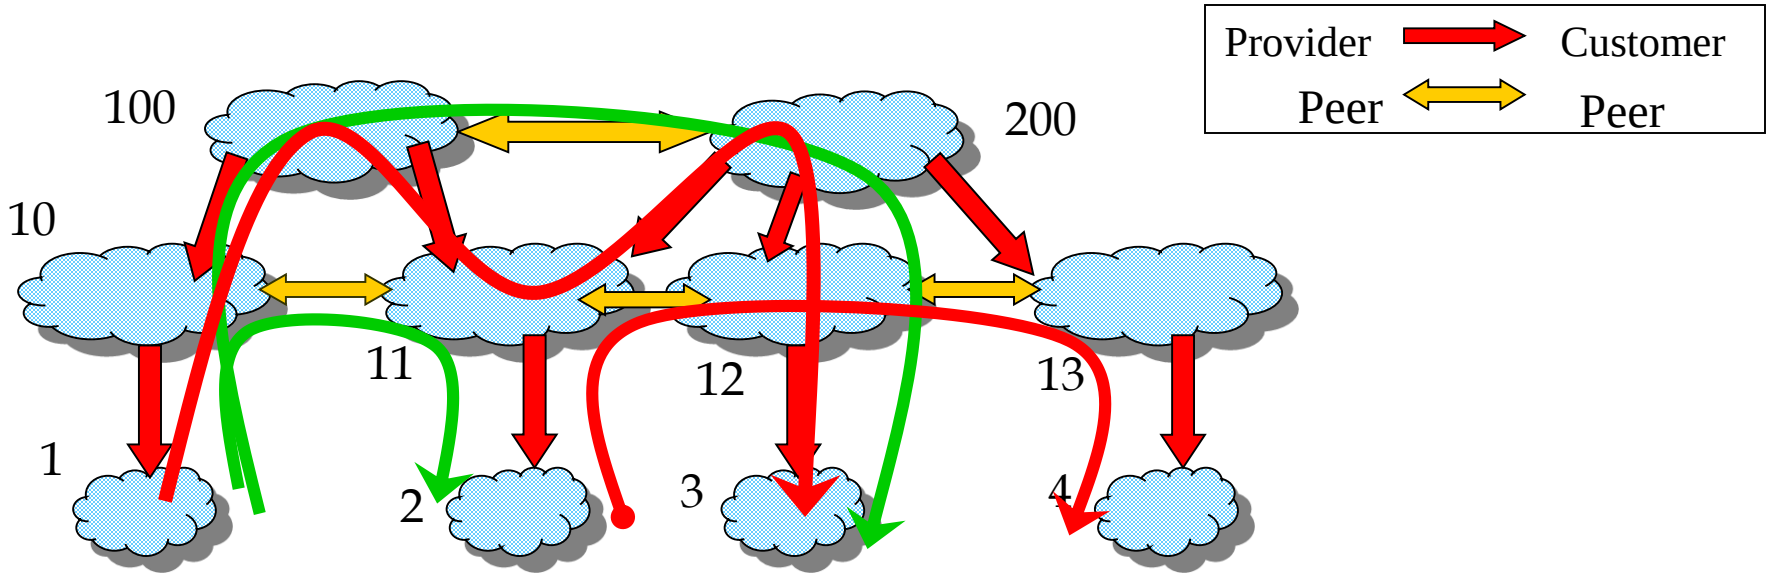

- Routing rules:
  - Provider accept everything
  - Peer only if it is for its customers
- Path Properties:
  - Up then down
  - No up-down-up, at most 1 peer-peer steps

Implementing BGP Rules

- What do you do with an advertisement:
- Through customer link
  - Advertise to all (customers, peers, providers)
- Through provider link
  - Advertise to customer only (and possibly siblings)
- Through peer link
  - Advertise to customer only (and possibly siblings)
- Through sibling link
  - Advertise to all

How Policies Affect Routing

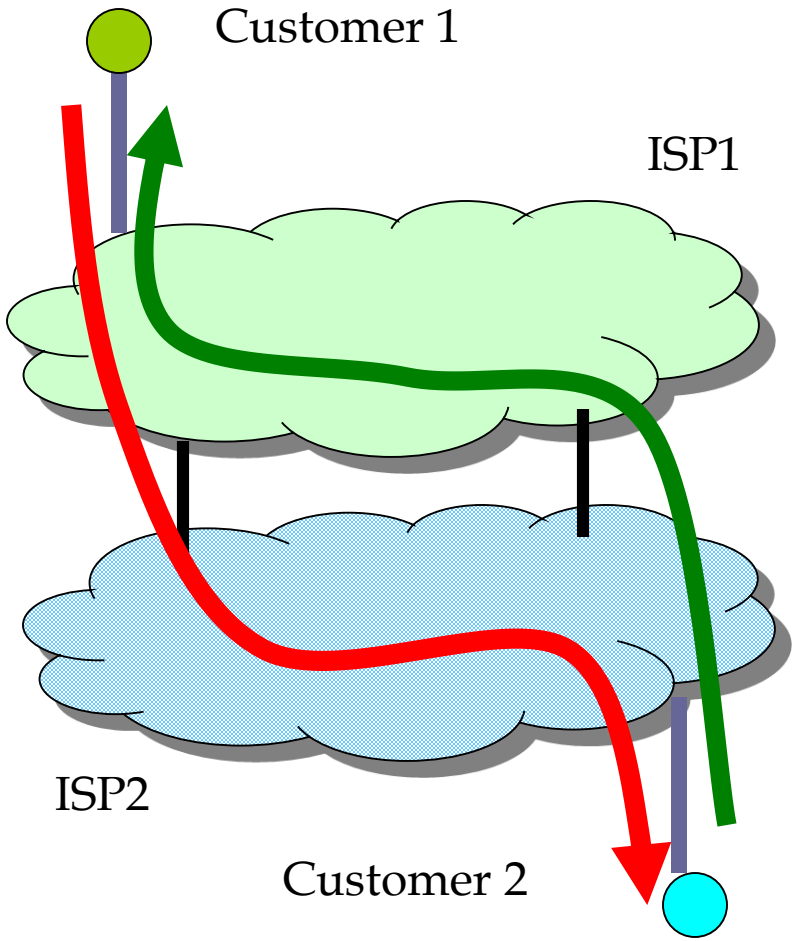

- A Provider will get rid of traffic as soon as possible,
- But a Provider will carry the traffic for its customer
- Did anyone say traffic is asymmetric?

BGP— Exterior Routing Protocol (1)

BGP (Border Gateway Protocol) computes routes across interconnected, autonomous networks

- Key role is to respect networks' policy constraints

Example policy constraints:

- No commercial traffic for educational network
- Never put Iraq on route starting at Pentagon
- Choose cheaper network
- Choose better performing network
- Don't go from Apple to Google to Apple

BGP— Exterior Routing Protocol (2)

Common policy distinction is transit vs. peering:

- Transit carries traffic for pay; peers for mutual benefit
- AS1 carries AS2↔AS4 (Transit) but not AS3 (Peer)

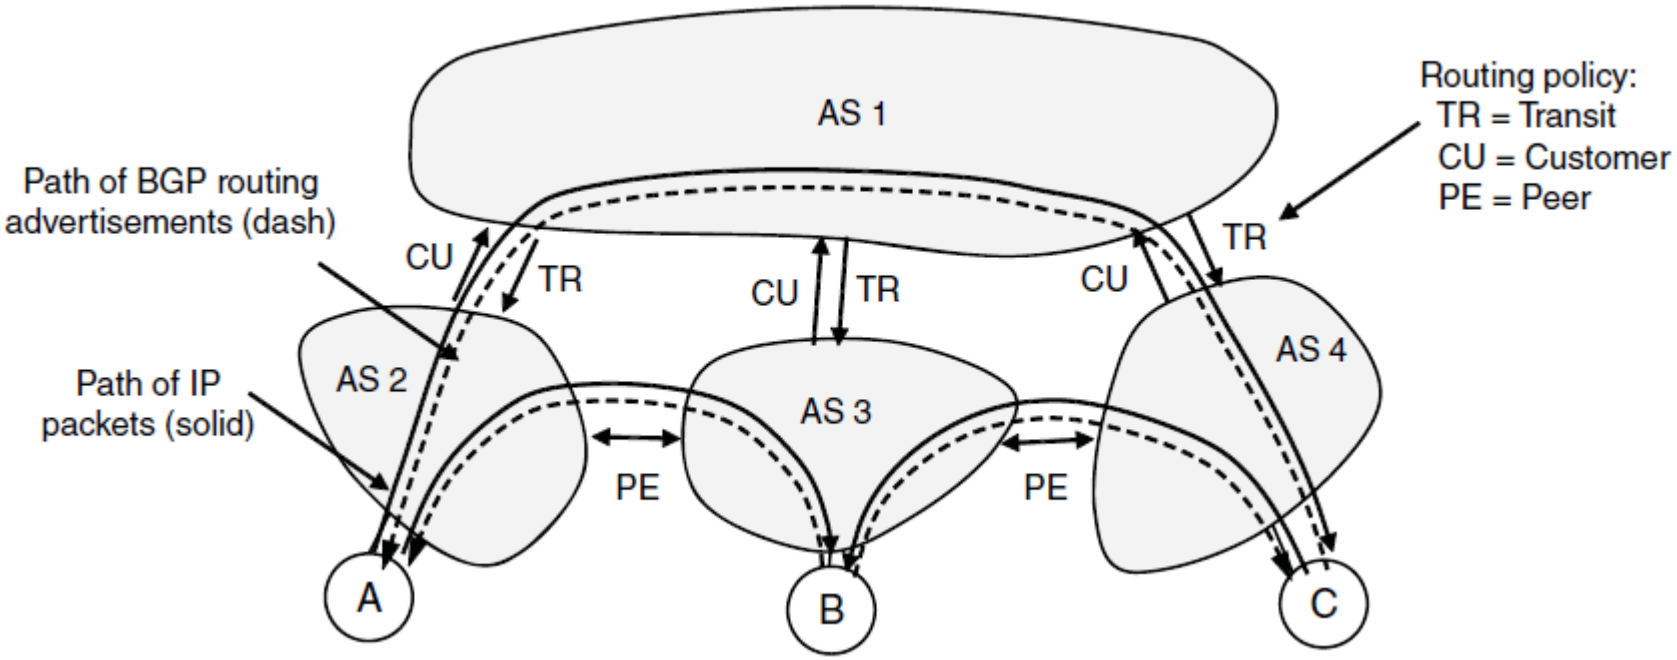

BGP— Exterior Routing Protocol (3)

- BGP propagates messages along policy-compliant routes
  - Message has prefix, AS path (to detect loops) and next-hop IP (to send over the local network)