Advanced Computer Networks

Network Routing

Jianping Pan
Summer 2007

Internet design

- Design principles
  - store-and-forward packet switching
  - end-to-end argument
- Endpoint transport-layer protocol control
  - connection management
  - flow, error and congestion control
- What's left for network layer?
  - addressing
  - routing

Internet addressing

- IPv4 address
  - 32-bit address space
  - class-based address allocation

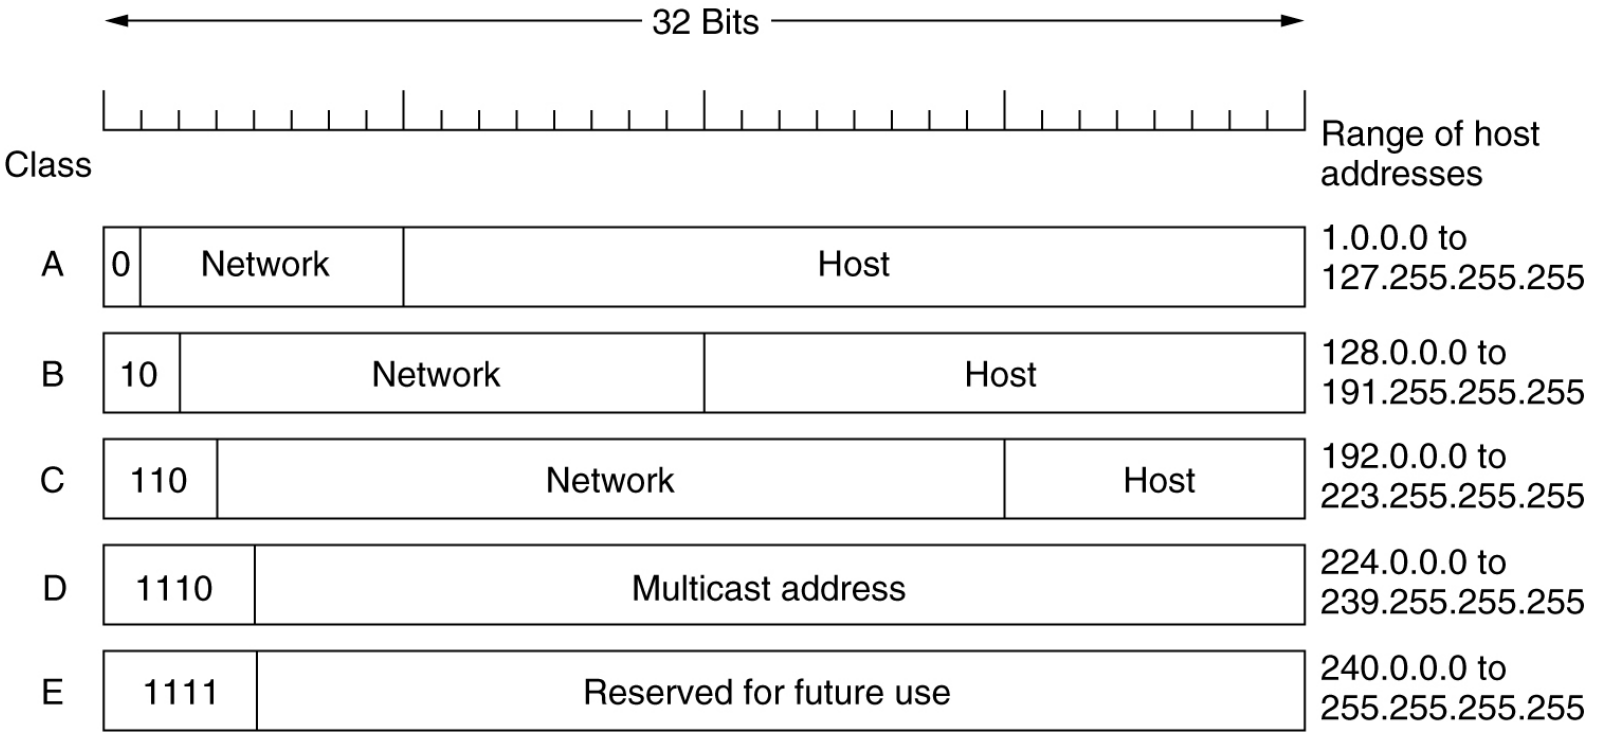

IPv6: 128-bit address space

Subnetwork

- E.g., UVicNet: 142.104.0.0/16
  - class B address
- Engineering Network (EngrNet)
  - 142.104.96.0/19
  - moving to 142.104.64.0/19
- Another example

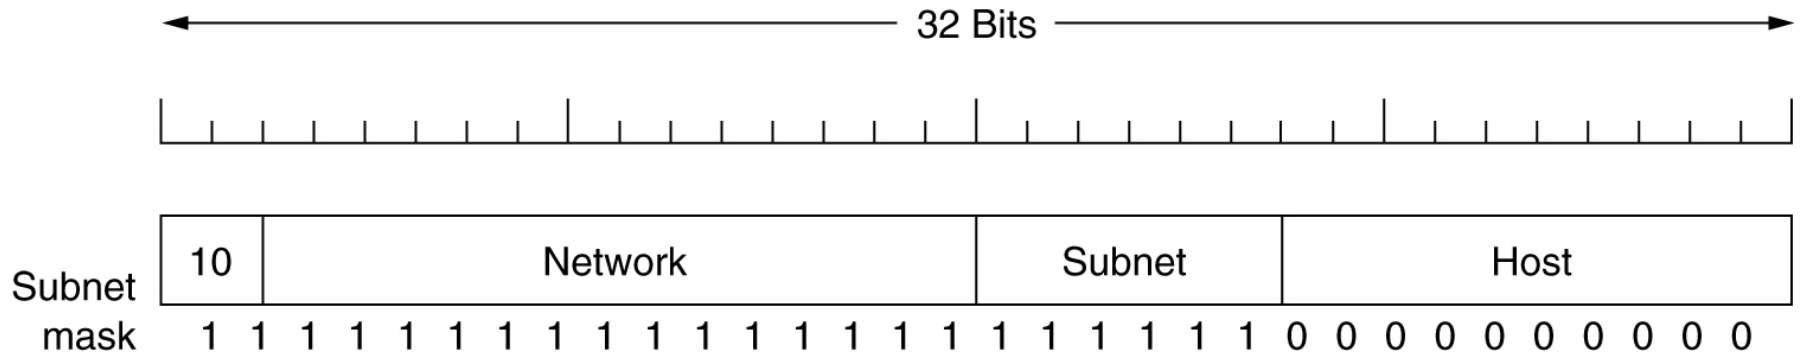

Classless Inter-Domain Routing

- CIDR addressing
  - no longer class A/B/C
    - a “class A” can be partitioned into “smaller” networks
    - a few “class C”s can be combined in a “bigger” network

University	First address	Last address	How many	Written as
Cambridge	194.24.0.0	194.24.7.255	2048	194.24.0.0/21
Edinburgh	194.24.8.0	194.24.11.255	1024	194.24.8.0/22
(Available)	194.24.12.0	194.24.15.255	1024	194.24.12/22
Oxford	194.24.16.0	194.24.31.255	4096	194.24.16.0/20

Hierarchical addressing

- Hierarchical address allocation
  - customer, provider, APNIC/ARIN/RIPE..., ICANN

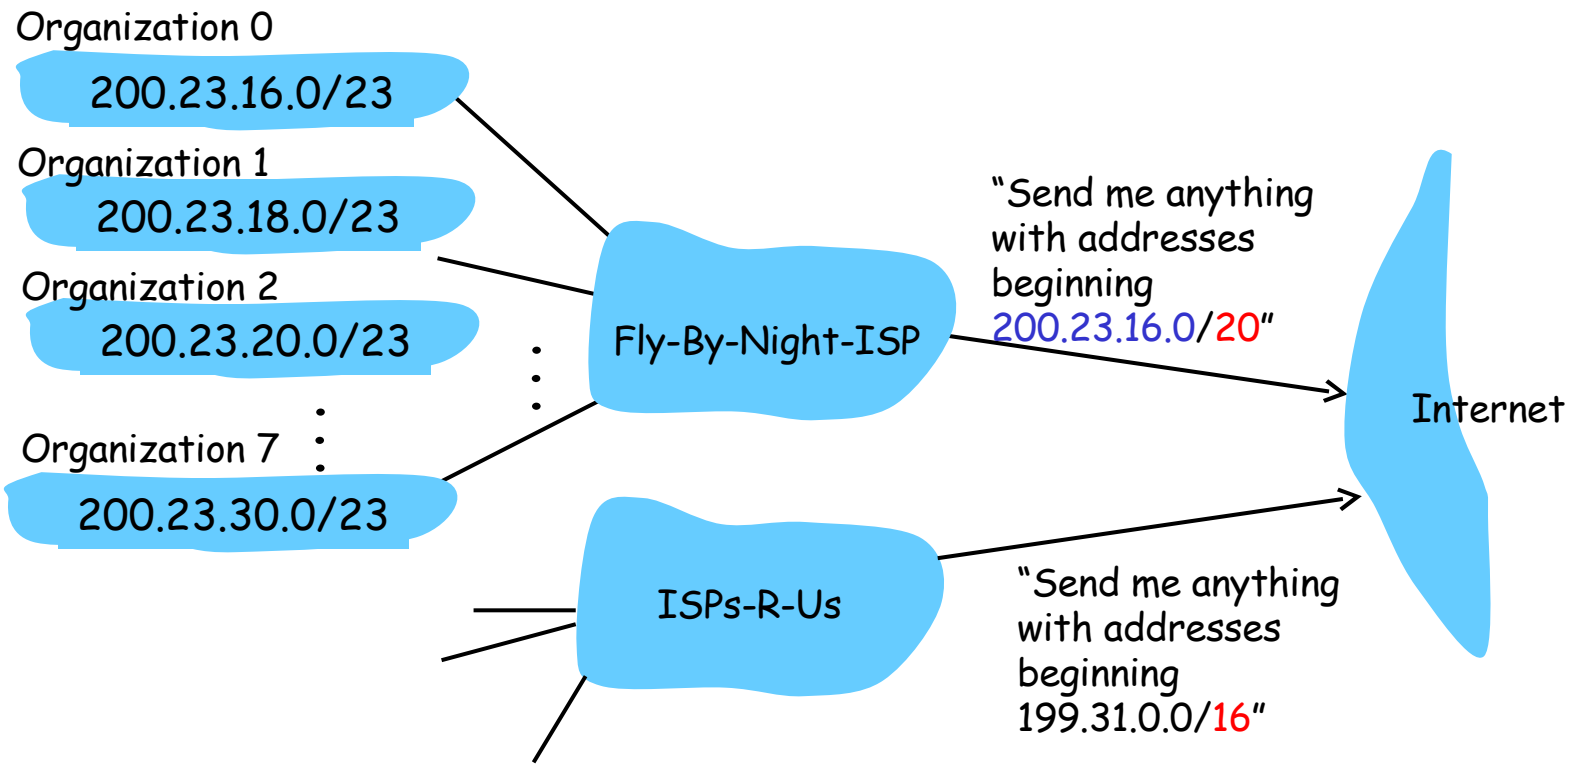

Address renumbering

- Hierarchical addressing
  - longest prefix match

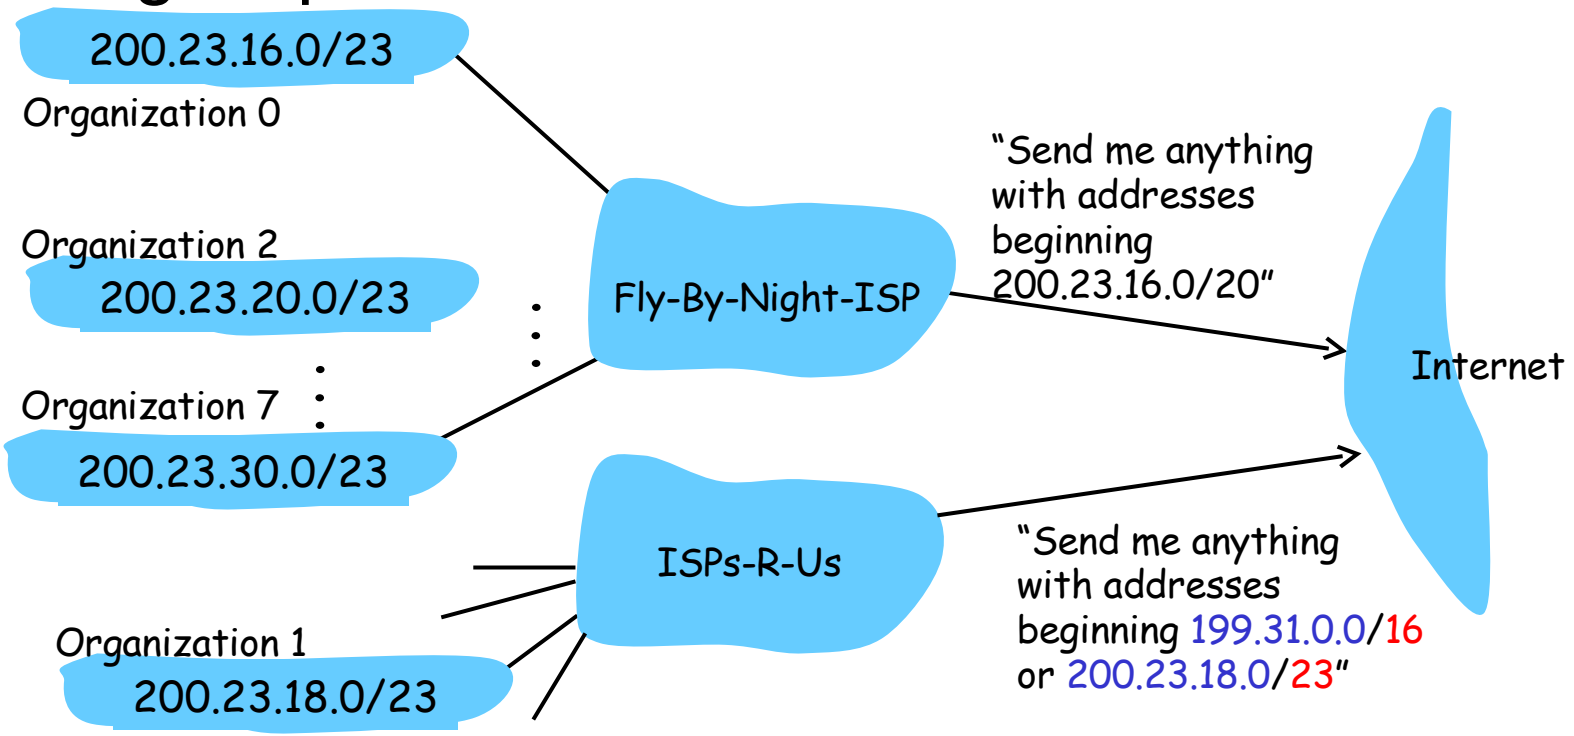

- Private addressing with NAT

Internet routing

- Internet routing
  - intra-domain vs inter-domain routing
  - autonomous system (AS)
- Intra-domain routing
  - RIP: distance vector
  - OSPF: link state
- Inter-domain routing
  - BGP: path vector

Distance vector routing

- Example

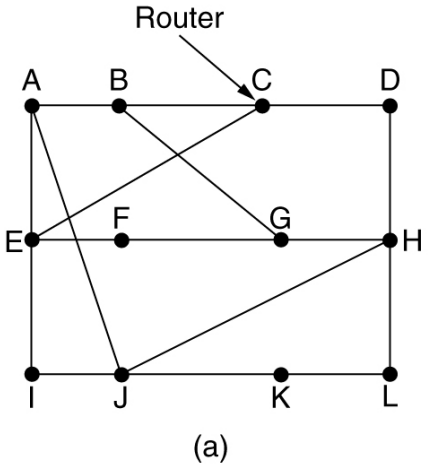

(a)

New estimated delay from J

To	A	I	H	K	Line
A	0	24	20	21	8 A
B	12	36	31	28	20 A
C	25	18	19	36	28 I
D	40	27	8	24	20 H
E	14	7	30	22	17 I
F	23	20	19	40	30 I
G	18	31	6	31	18 H
H	17	20	0	19	12 H
I	21	0	14	22	10 I
J	9	11	7	10	0 –
K	24	22	22	0	6 K
L	29	33	9	9	15 K
	JA delay is 8	JI delay is 10	JH delay is 12	JK delay is 6	New routing table for J

Vectors received from J's four neighbors

(b)

Q: problems with DV routing?

Count-to-infinite

- Example

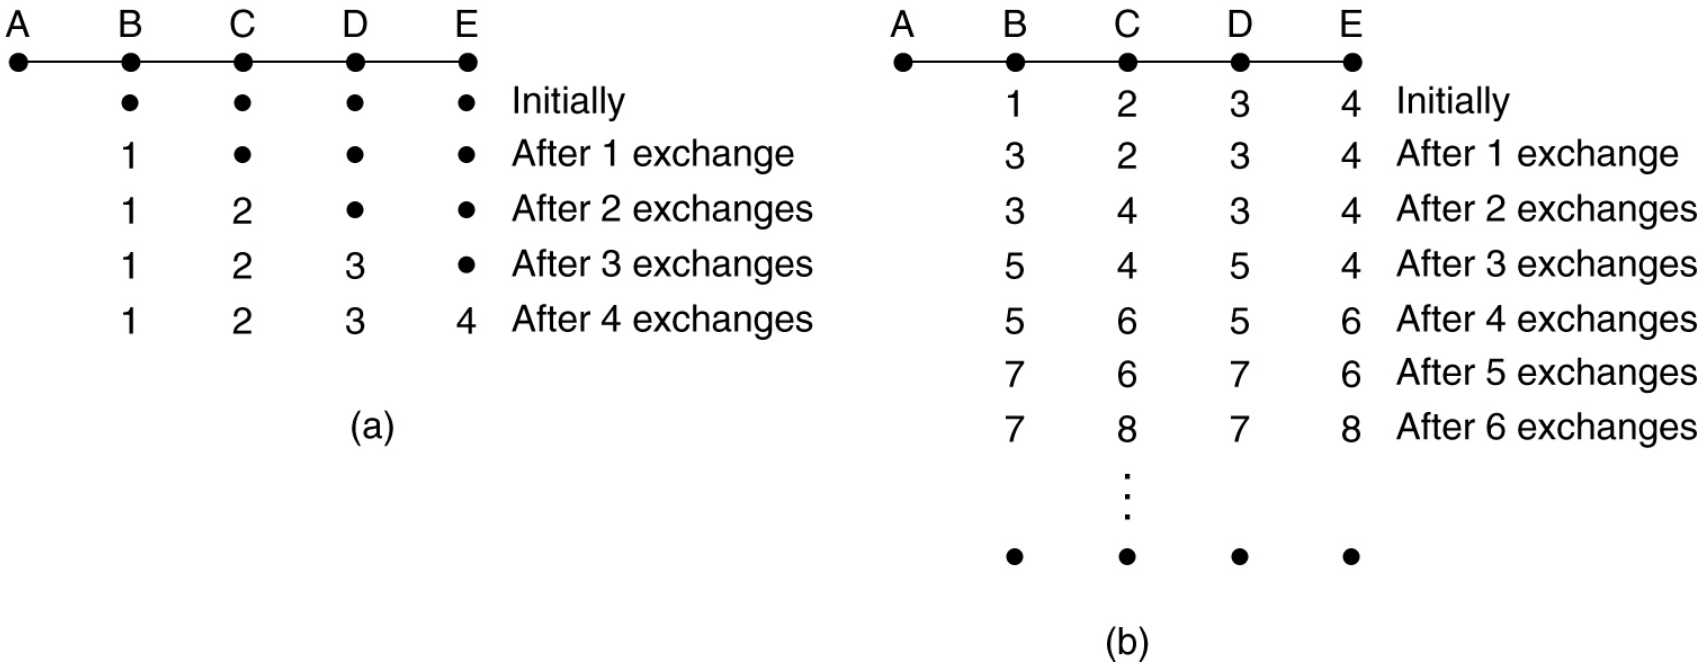

Routing Information Protocol

- RIP
  - distance vector
  - included in BSD Unix since 1982
  - max hops: 15
- Distance vector
  - exchanged between neighbors every 30s
    - up to 25 destinations within an AS
  - if no advertisement for 180s
    - neighbor is dead
    - invalidate routes going through the neighbor

Link state routing

- Example

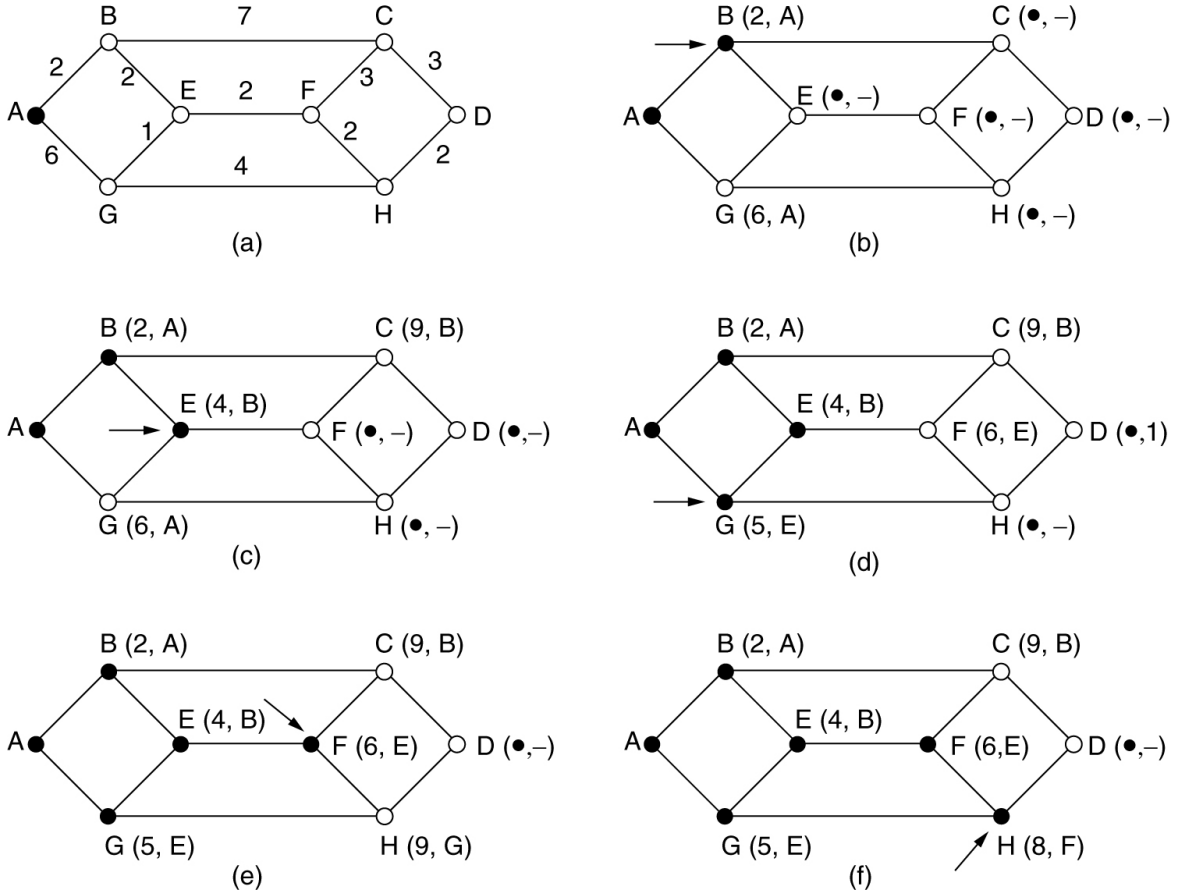

Open Shortest-Path First

- OSPF
  - using link state routing algorithm
  - link state dissemination
    - flooding, directly over IP
  - topology map at each node
  - Dijkstra's algorithm at each node
- Hierarchical OSPF
  - intra-domain areas: backbone and areas
    - flooding in an area
    - area border routers

Hierarchical routing

- Scalability: go hierarchical!

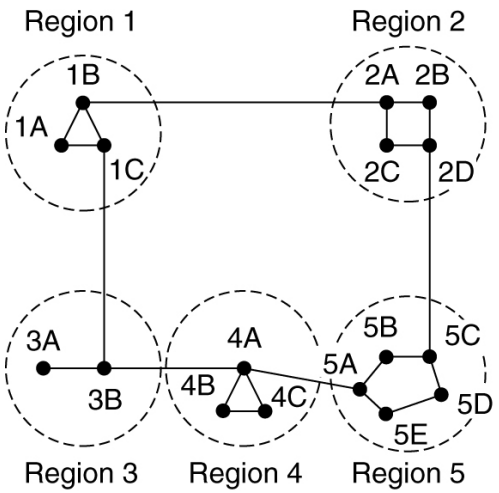

(a)

Full table for 1A

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

(b)

Hierarchical table for 1A

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

(c)

Border Gateway Protocol: basics

- BGPv4: based on distance vector
  - the de facto inter-domain routing standard
  - heavily policy-influenced, over TCP
- Reachability information
  - “you can reach X through me”
- AS path
  - “with a path of AS numbers”
  - AS: autonomous system (e.g., ISP domains)

BGP advertisement

- Prefix:
- AS Path:
- Next-hop:

1.2.3.4/16	1.2.3.4/16
AS3	AS3, AS1
3a	1b

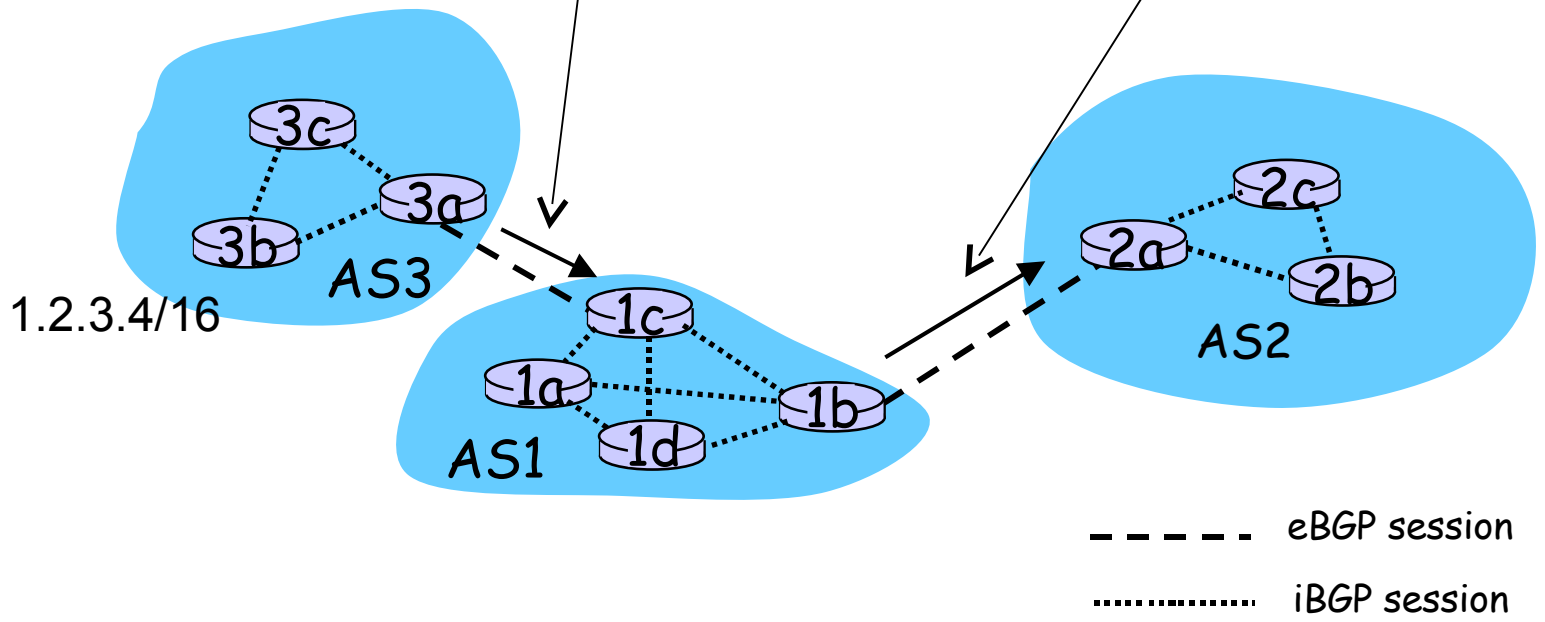

Student presentation

- Ching-Chang Chen: HN-SPF
  - [KZ90] A. Khanna and J. Zinky, "A Revised ARPANET Routing Metric," ACM SIGCOMM '89, pp. 45-56, September 1989.

Further discussion

- Internet routing

This lecture

- Internet addressing
  - class-based addressing
  - CIDR
- Internet routing
  - distance vector, link state, path vector
  - RIP, OSPF, BGP
- How to choose a “good” routing metric

Next lectures

- June 27
  - [LMJ97] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability". In Proceedings of ACM SIGCOMM'97, September 1997.
- July 4
  - [GR00] Lixin Gao and Jennifer Rexford, "Stable Internet Routing Without Global Coordination". In Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems. 2000.