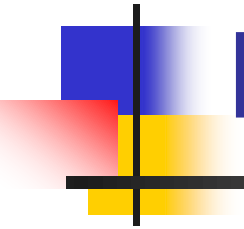

Routing

รศ.ดร. อนันต์ ผลเพิ่ม

Asso. Prof. Anan Phonphoem, Ph.D.

anan.p@ku.ac.th

http://www.cpe.ku.ac.th/~anan

Computer Engineering Department

Kasetsart University, Bangkok, Thailand

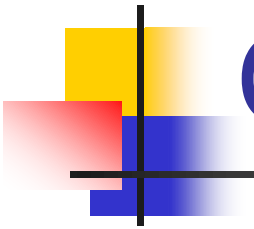

Outline

- Routing Principle
- Routing table calculation
- **Routing protocol**

Routing Protocol

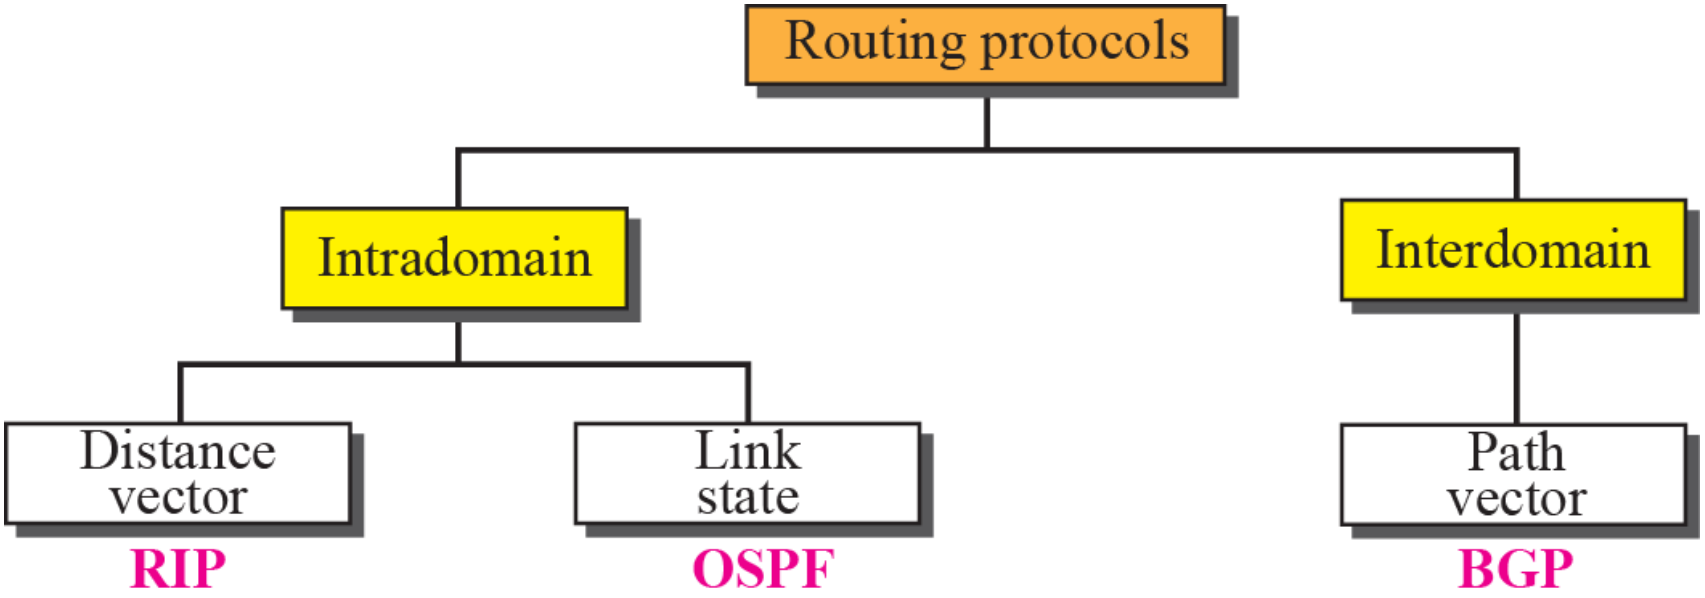

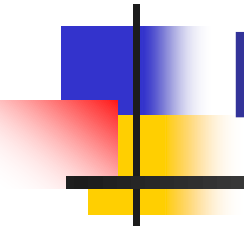

Distance-vector Routing

Distance-vector Routing

- Distributed route computation
- Router computes the routing table locally
- Each node sends the results to its neighbors periodically
- Each node keeps update its routing table
- **Routing by rumor**

Configure RIP

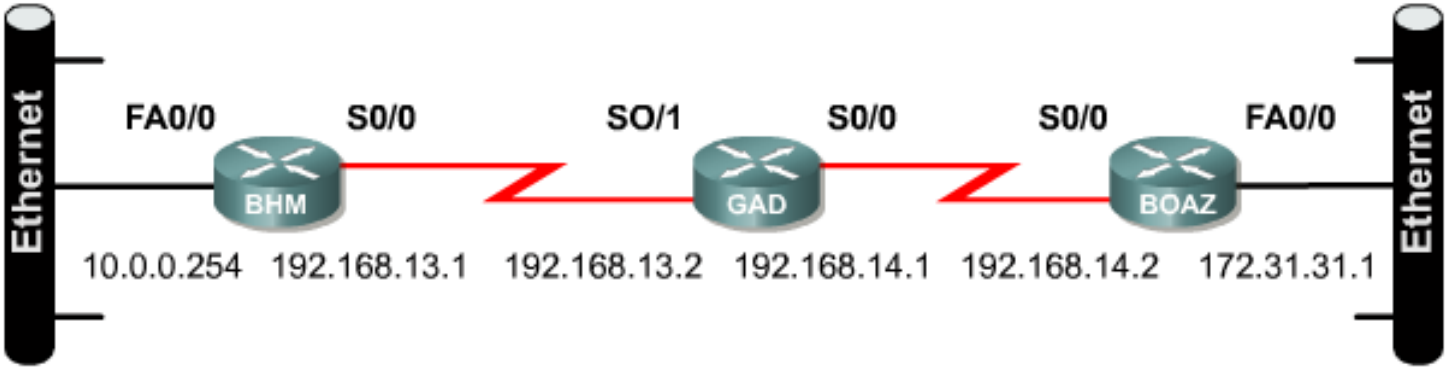


```
BHM(config)#router rip
BHM(config-router)#network 10.0.0.0
BHM(config-router)#network 192.168.13.0
```

```
GAD(config)#router rip
GAD(config-router)#network 192.168.14.0
GAD(config-router)#network 192.168.13.0
```

```
BOAZ(config)#router rip
BOAZ(config-router)#network 192.168.14.0
BOAZ(config-router)#network 172.31.0.0
```

Routing Table Example

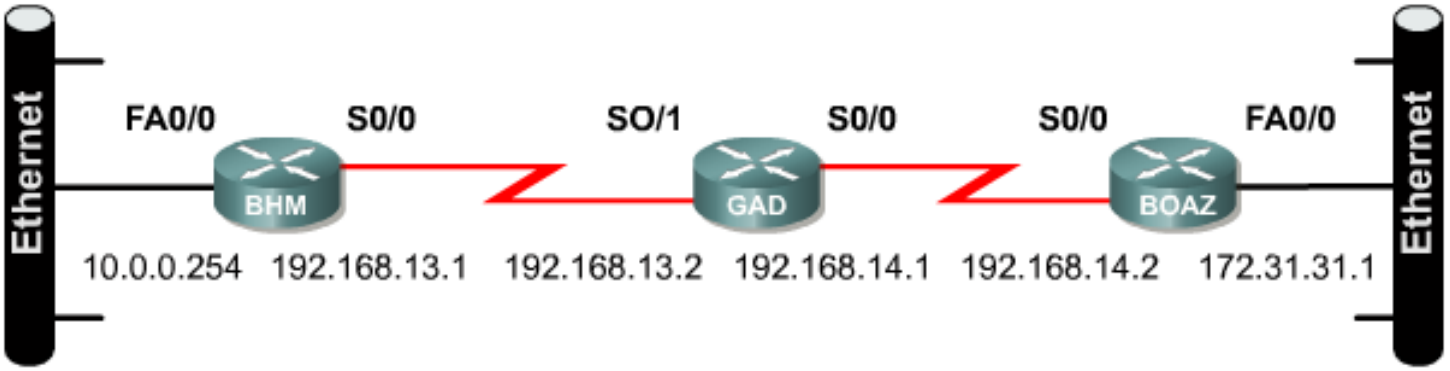

BOAZ#sh ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2, E1 - OSPF external type 1, E2 - OSPF external type 2, E – EGP, i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area, * - candidate default, U - per-user static route, o – ODR, P - periodic downloaded static route

Gateway of last resort is not set

R 10.0.0.0/8 [120/2] via 192.168.14.1, 00:00:16, Serial0/0/0

C 172.31.0.0/16 is directly connected, FastEthernet0/0

R 192.168.13.0/24 [120/1] via 192.168.14.1, 00:00:16, Serial0/0/0

C 192.168.14.0/24 is directly connected, Serial0/0/0

RIPv2

- RIPv2 supports
  - CIDR supernets
  - VLSM
  - Discontiguous networks (no auto-summary)
- Improves efficiency by
  - defining multicast address (224.0.0.9) → route update
  - some reserved fields are assigned (RIPv1 sets to zero)
  - Authentication
- Can coexist with RIP v1
  - RIPv2 use multicast must be configurable per interface
  - RIPv2 accept RIPv1 request/response

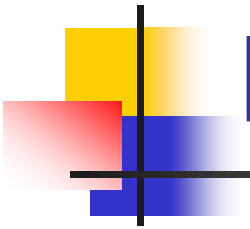

Distance Vector Summary

- Good
  - Only need communicate with neighbors (so little bandwidth is wasted on protocol overhead)
  - Relatively little processing of info
- Bad
  - Count to infinity problem
  - Slow convergence (the real issue)
- Despite this, RIP is popular
  - Because included in original BSD implementation

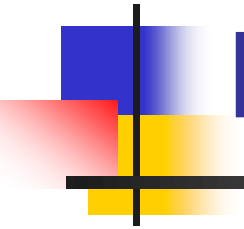

Link State Routing

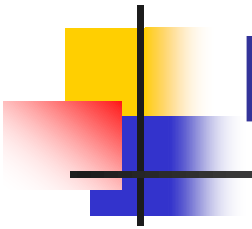

Link State Routing

- Each router is responsible for
  - meeting its neighbors
  - learning their names
- Each router constructs a packet
  - called a Link State Packet (LSP)
  - or Link State Advertisement (LSA)
  - containing a list of names and cost assigned to each neighbor

Router Discovery

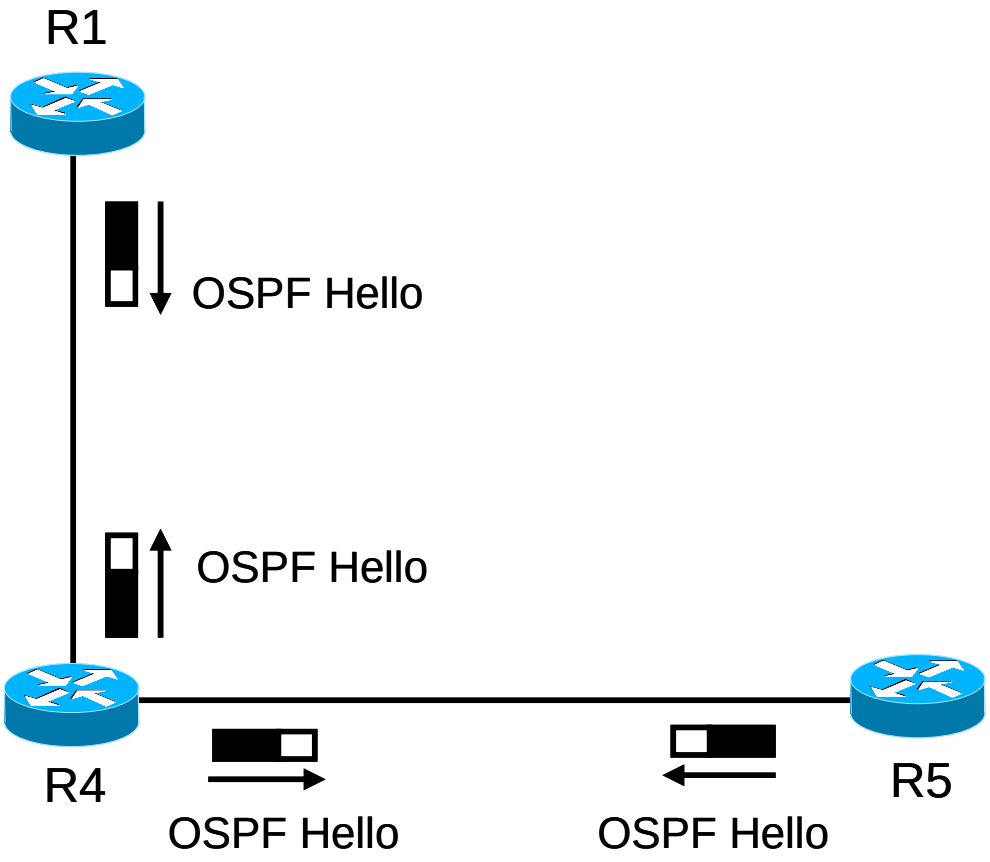

Link State Routing

- Each router LSP is sent to all other routers (Reliable Flooding)
- Each router now has a topological map of the network (Link-State Database)
- Apply the Shortest Path First (SPF) algorithm to compute routes to each destination
  - SPF also know as Dijkstra algorithm
  - Link-State algorithms advertise the state of its local network links not distances

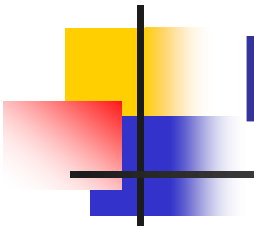

Link State Packet (LSP)

- ID of the node that created the LSP
- Cost of the link to each directly connected neighbor
- sequence number (SEQNO)
- time-to-live (TTL) for this packet

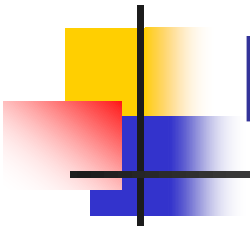

Reliable flooding

- store most recent LSP from each node
- forward LSP to all nodes but one that sent it
- generate new LSP periodically
  - increment SEQNO
- start SEQNO at 0 when reboot
- decrement TTL of each stored LSP
  - discard when TTL=0

Flooding

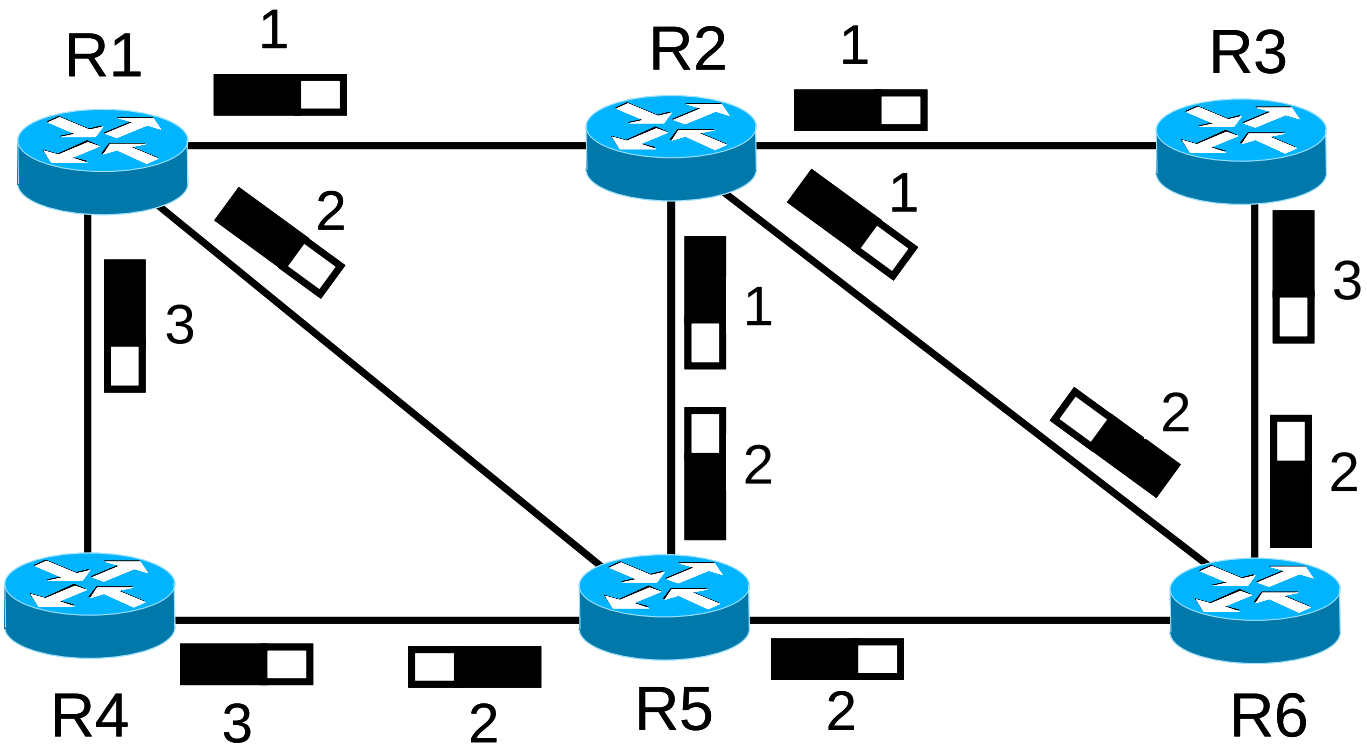

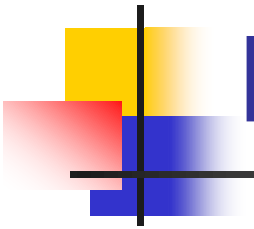

Route Calculation

- Dijkstra's Algorithm
  - Finding the shortest path from a source to other nodes in network
- Weight represents distance between two nodes
- Sum of weights along the path is the total distance
- Choose the lowest total distance

Link State Routing

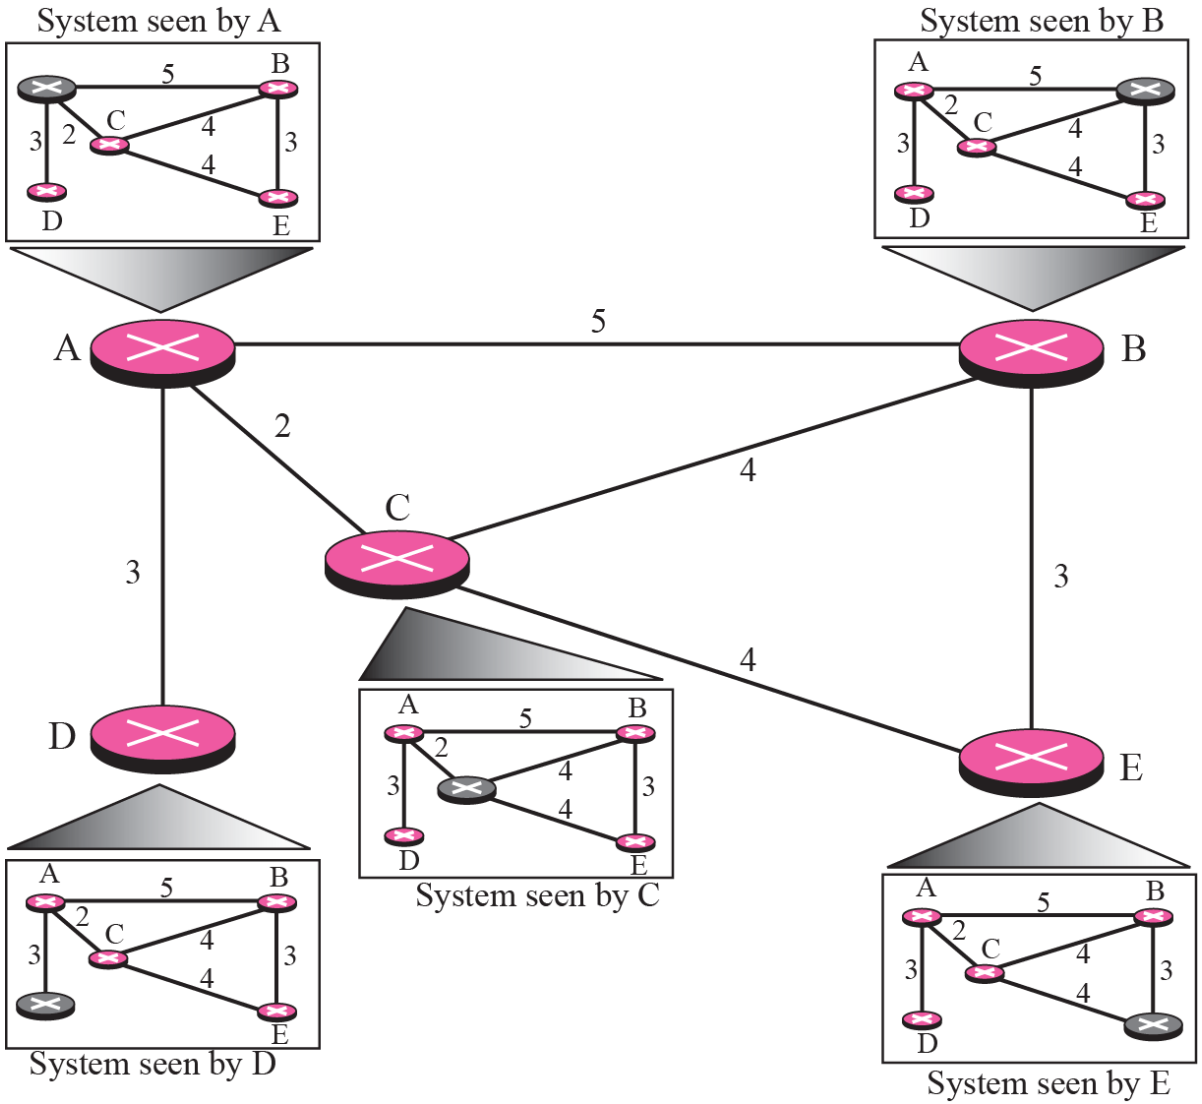

Dijkstra Algorithm

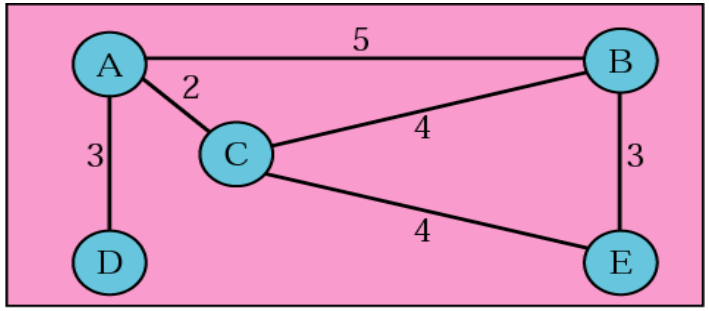

Topology

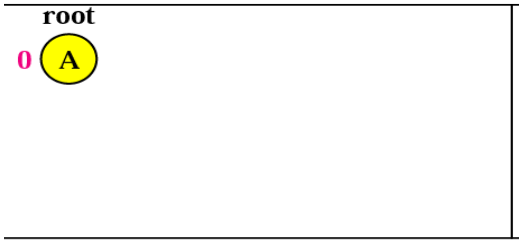

1. Set root to A and move A to tentative list

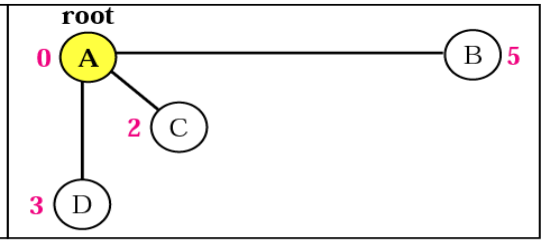

2. Move A to permanent list and add B, C, and D to tentative list

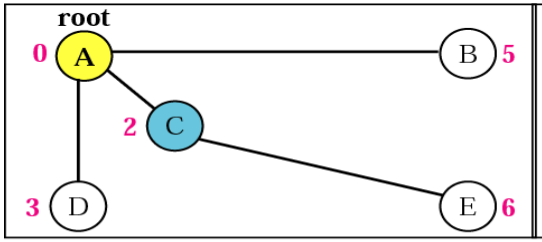

3. Move C to permanent and add E to tentative list

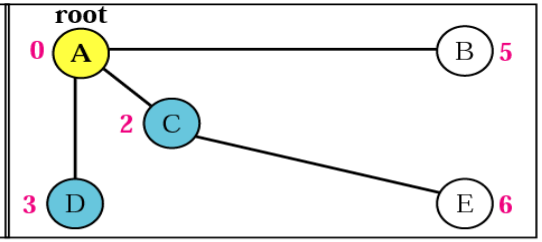

4. Move D to permanent list.

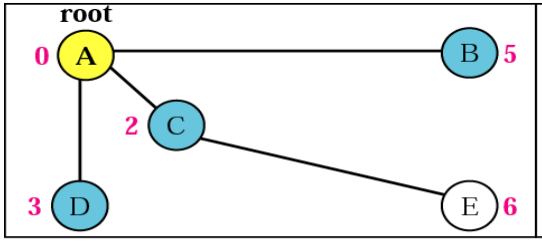

5. Move B to permanent list

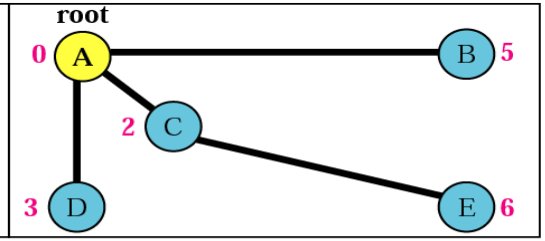

6. Move E to permanent list (tentative list is empty)

Routing Table

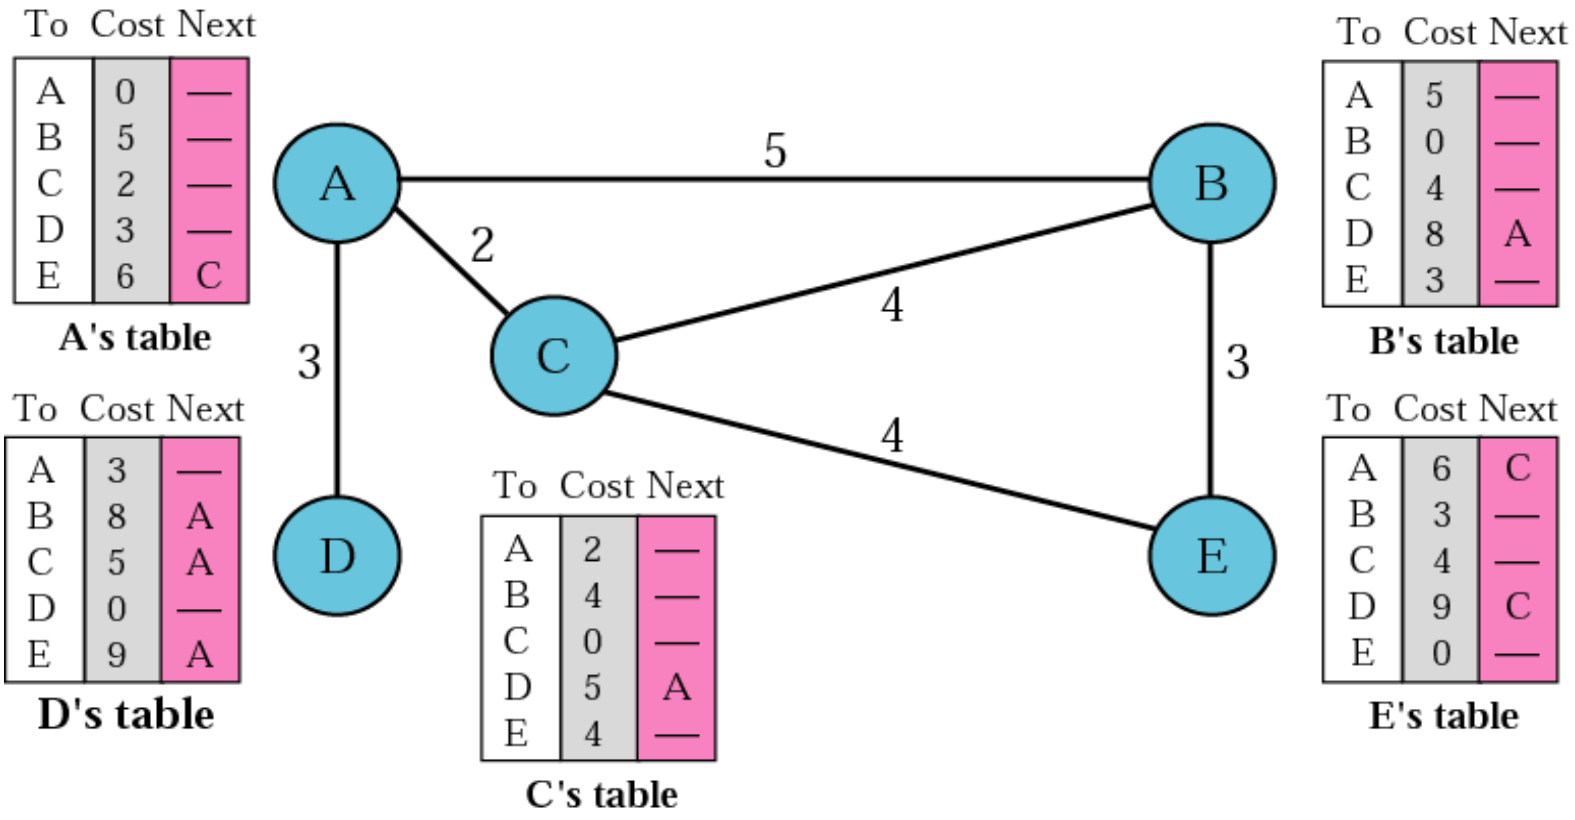

Dijkstra's Algorithm In action (I)

Japan: http://www-b2.is.tokushima-u.ac.jp/~ikeda/suuri/dijkstra/Dijkstra.shtml

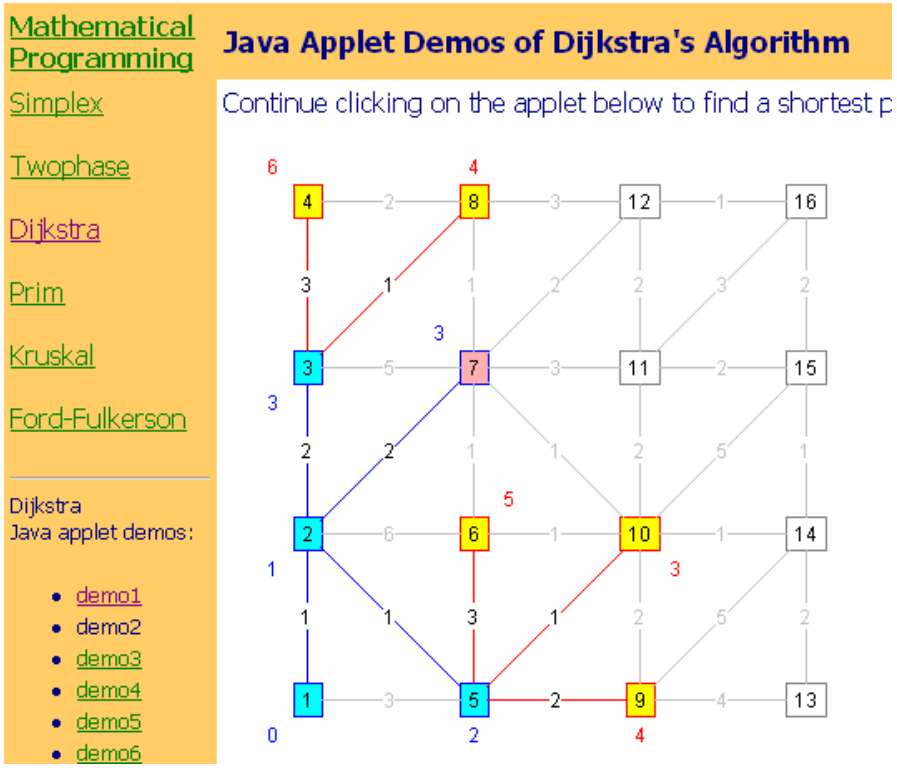

Dijkstra's Algorithm In action (II)

by Carla Laffra of Pace University

http://www.dgp.toronto.edu/people/JamesStewart/270/9798s/Laffra/DijkstraApplet.html

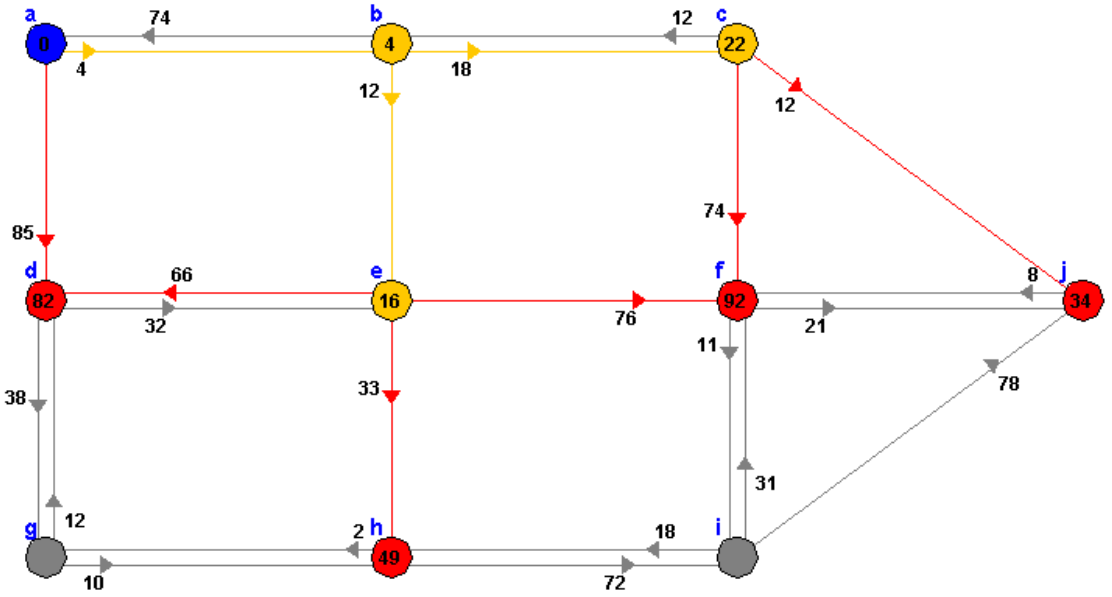

Routing Table

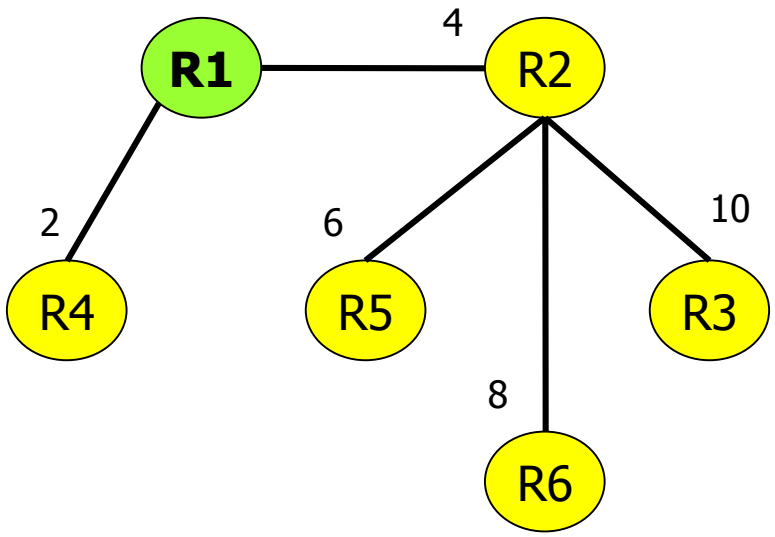

Final **R1** Routing Table

Destination	Next Hop	Cost
R2	-	4
R3	R2	10
R4	-	2
R5	R2	6

Network Change

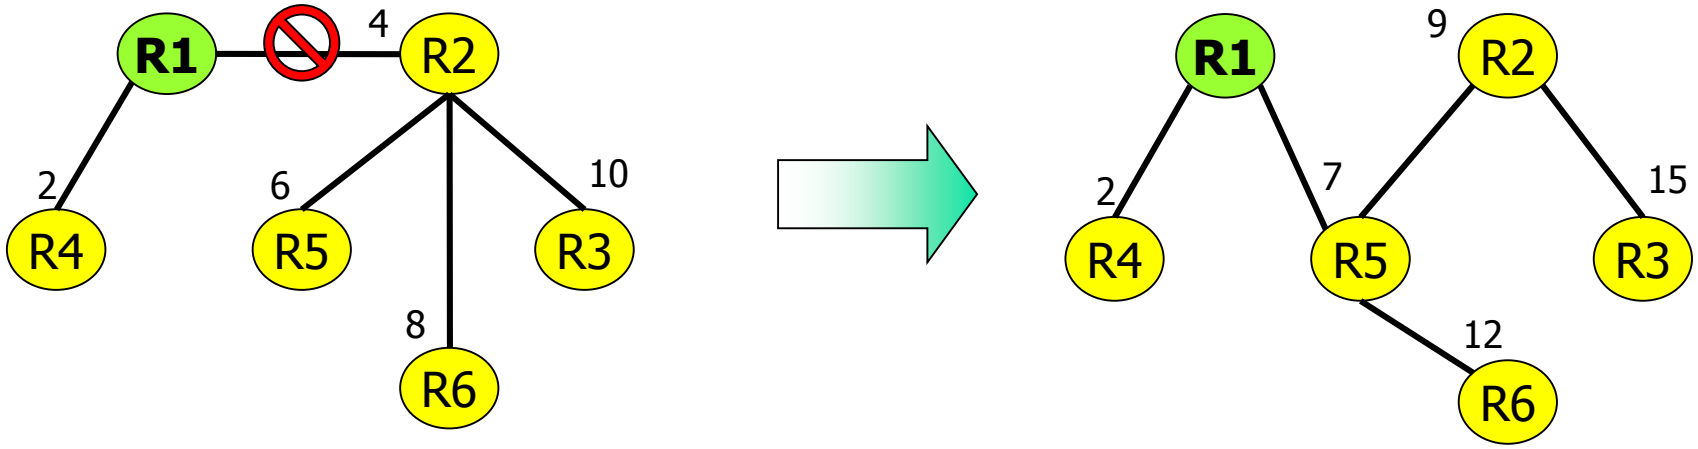

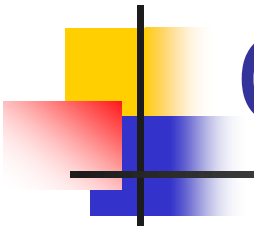

OSPF Routing Protocol

- Link State Protocol
  - Link is connection between two routers
  - Routing table stores > just its hop count: cost, reliability, etc.
  - Allows OSPF routers to optimize routing based on these variables

Autonomous System

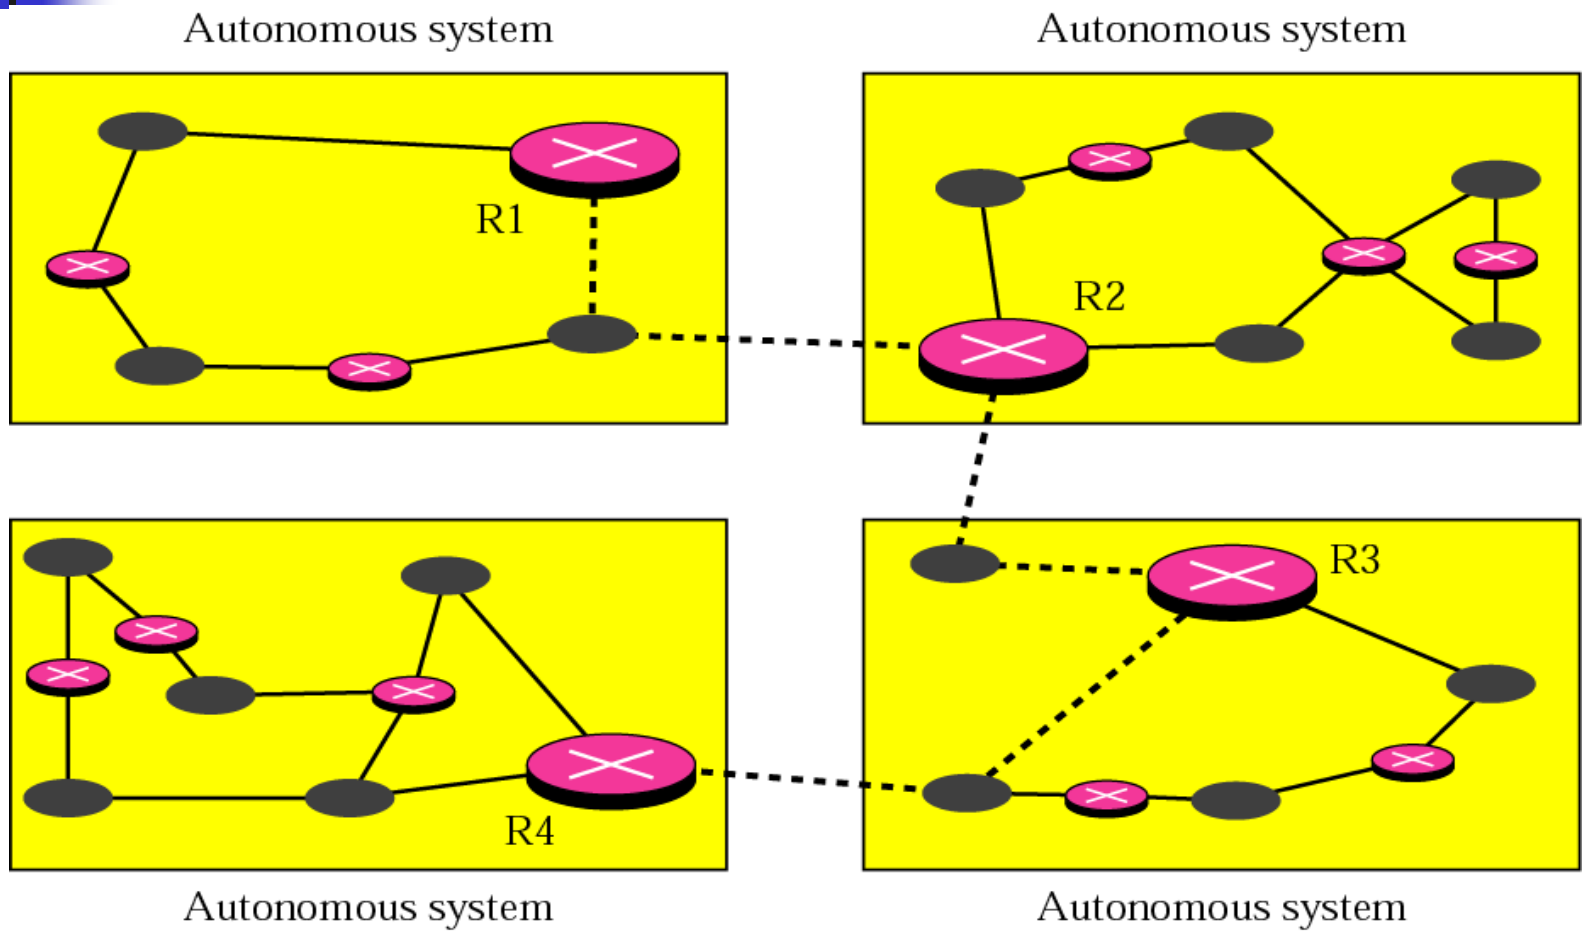

OSPF-Autonomous System

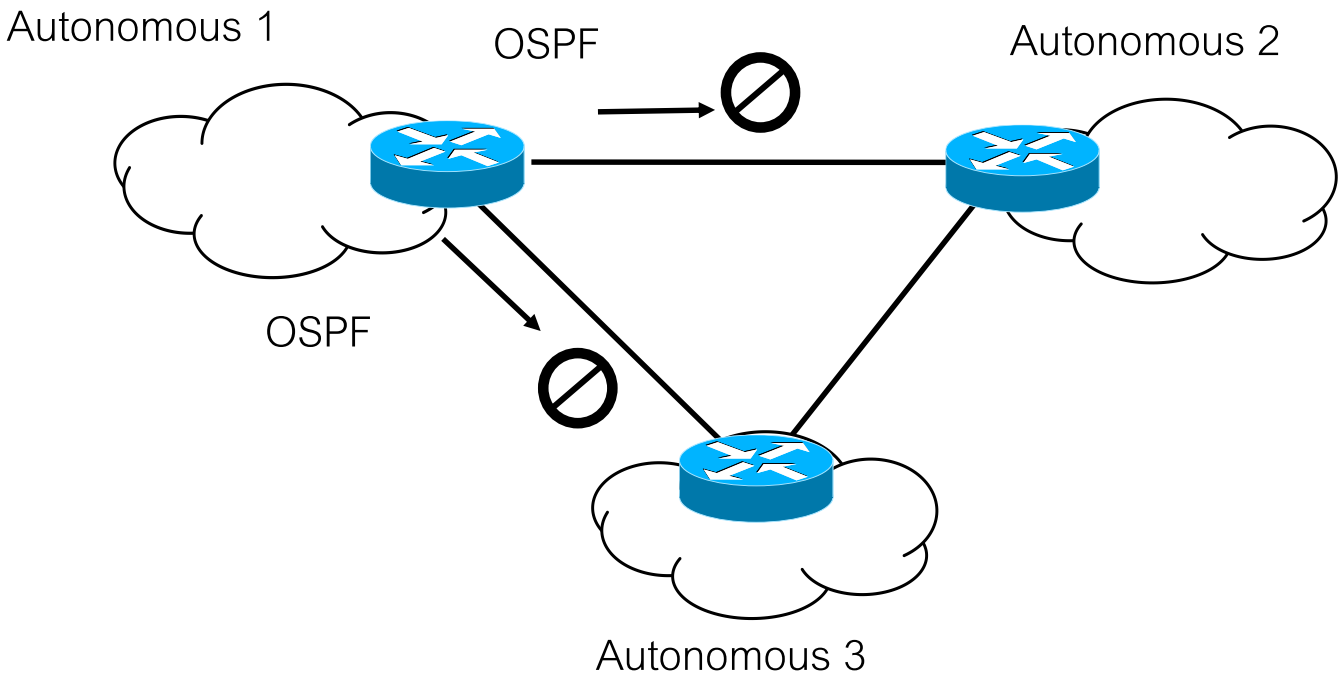

OSPF Area

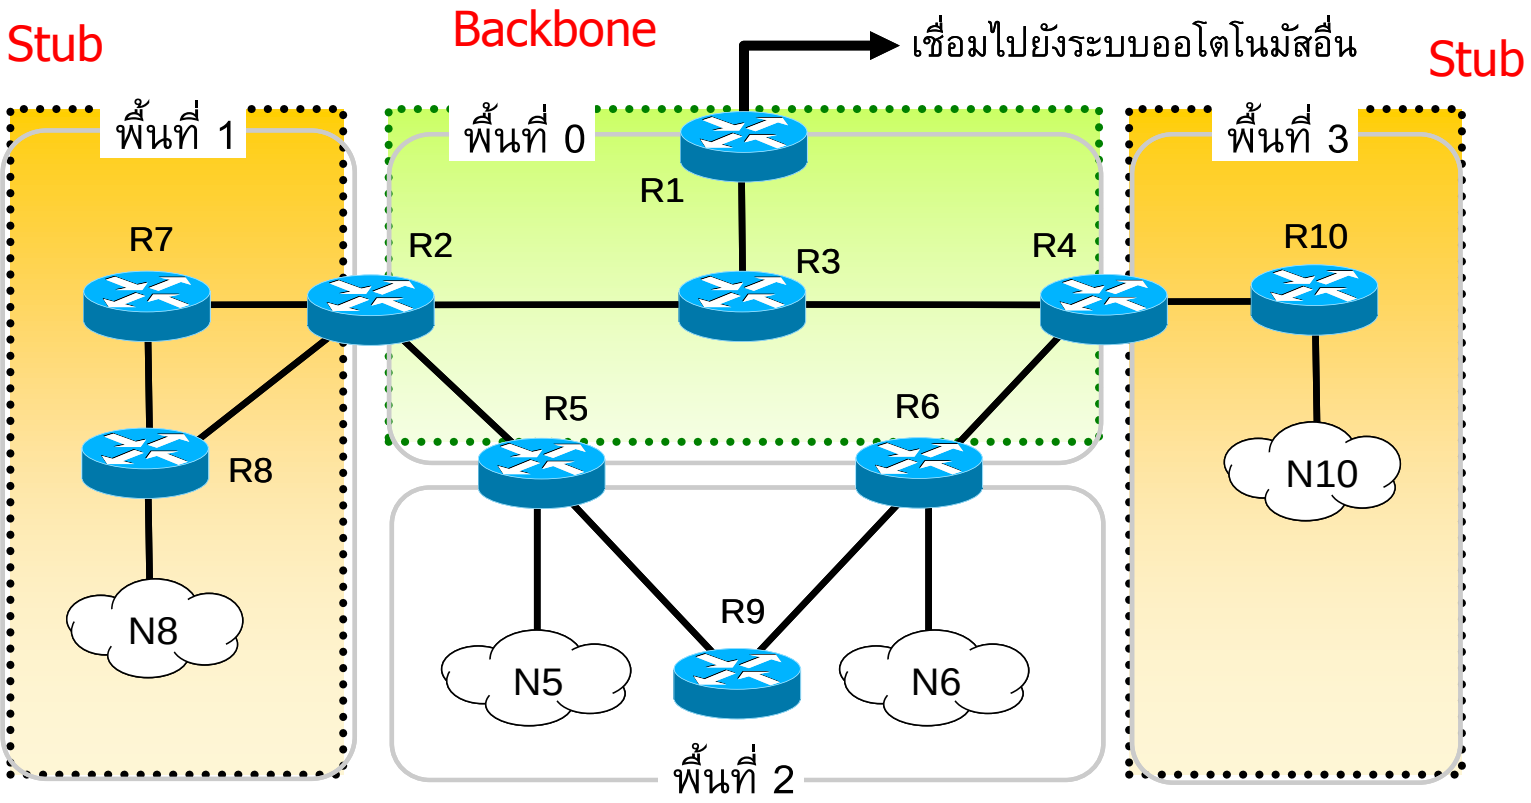
Stub
Backbone
เชื่อมไปยังระบบออโตโนมัสอื่น
Stub
พื้นที่ 1
พื้นที่ 0
พื้นที่ 3
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
N5
N6
N8
N10
พื้นที่ 2

Route between Area

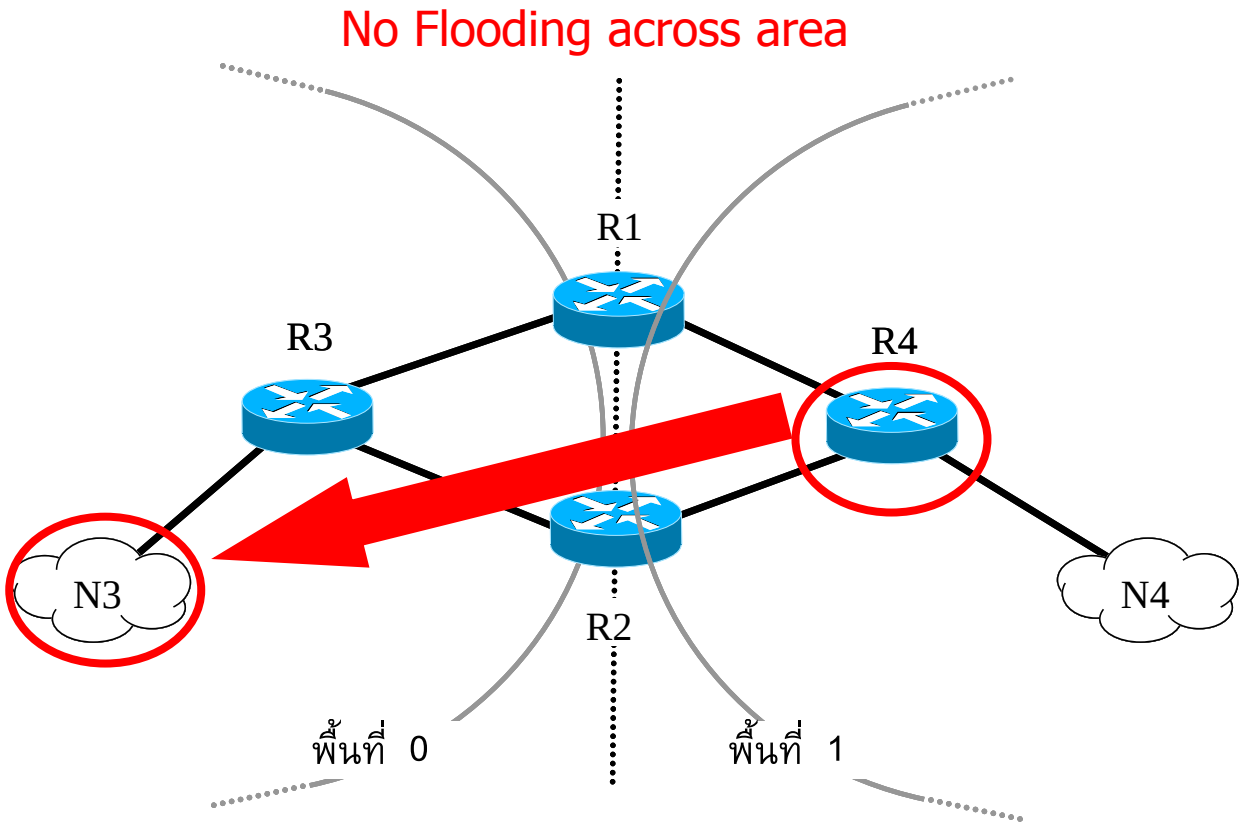

OSPF Routers

- Network is Divided into Areas
  - Each area has a designated router

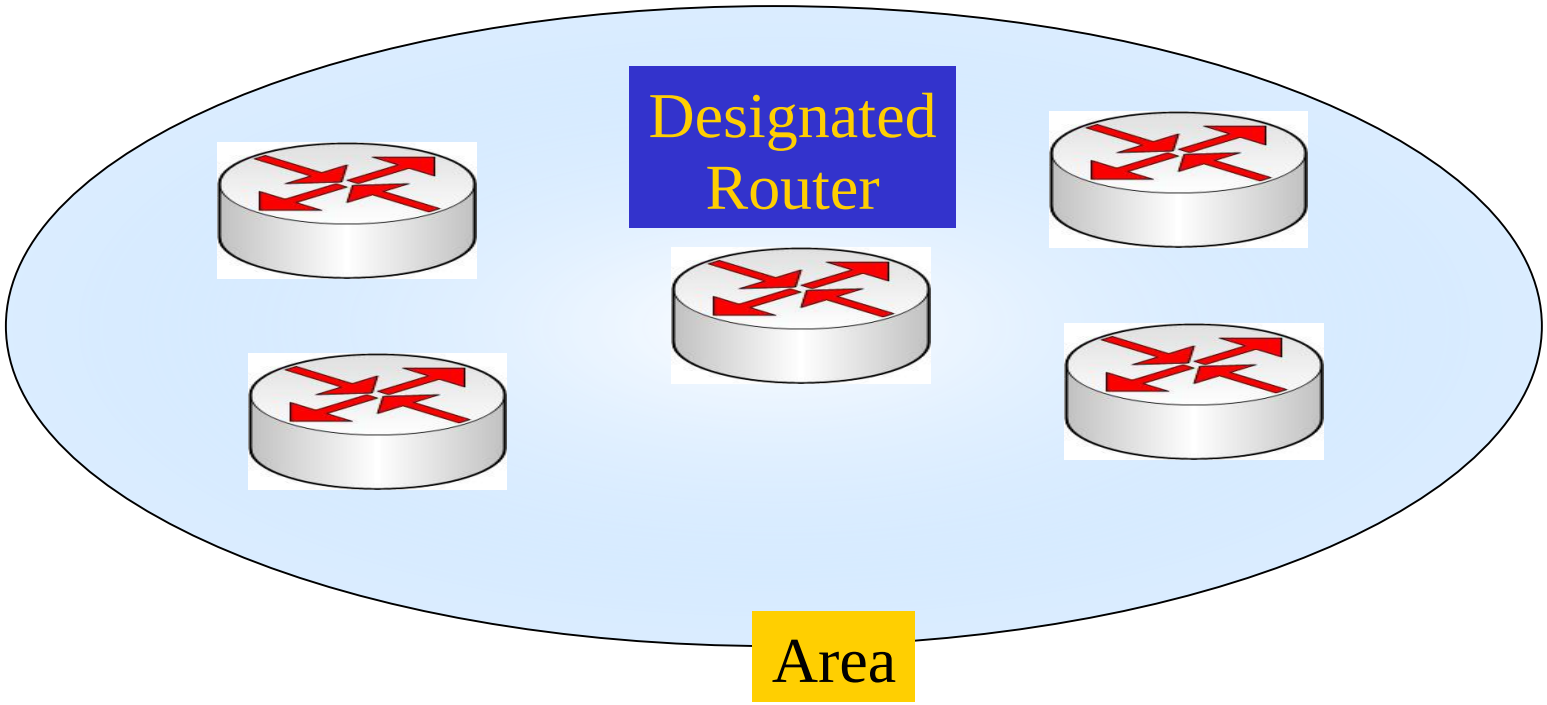

OSPF Routers

- When a router senses a link state change
  - Sends this information to the designated router

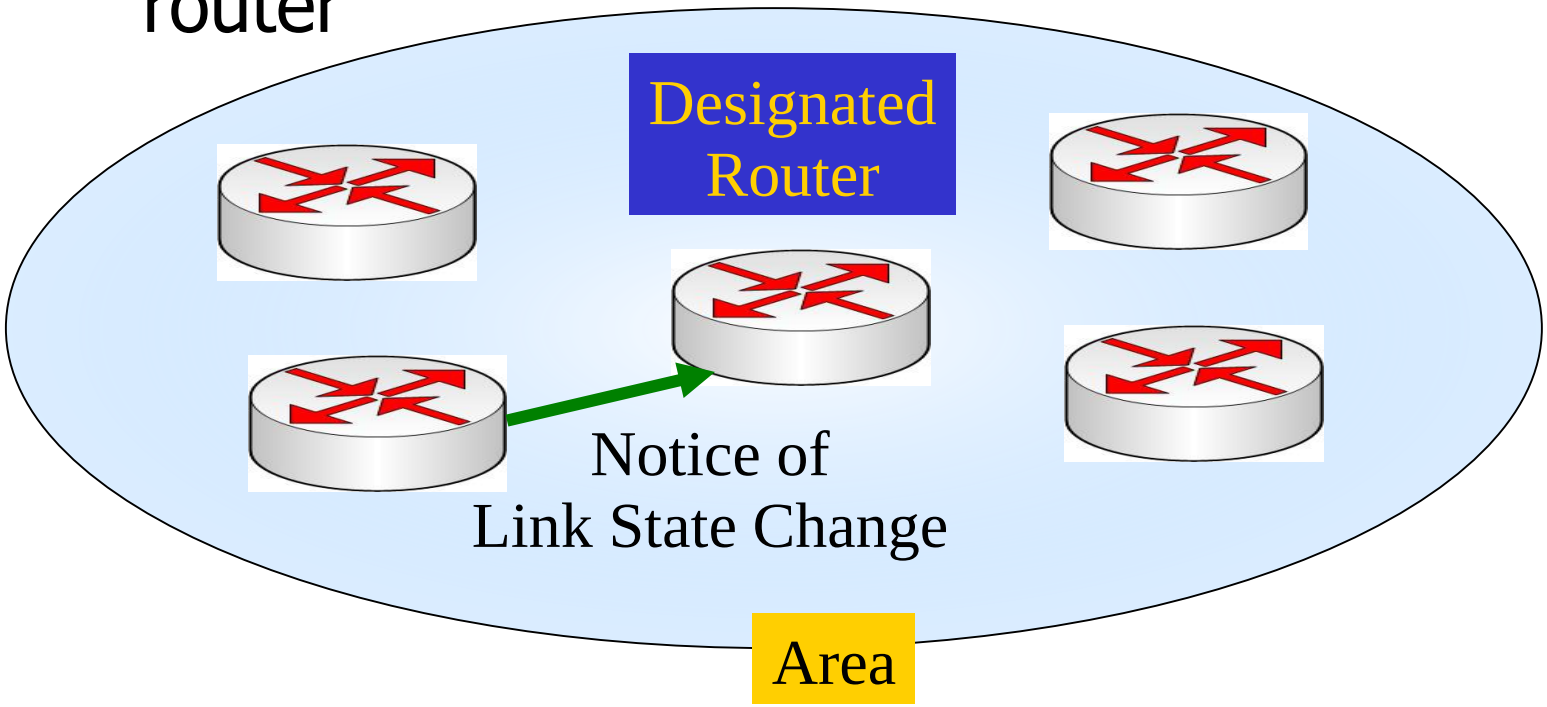

OSPF Routers

- Designed Router Notifies all Routers
  - Within its area

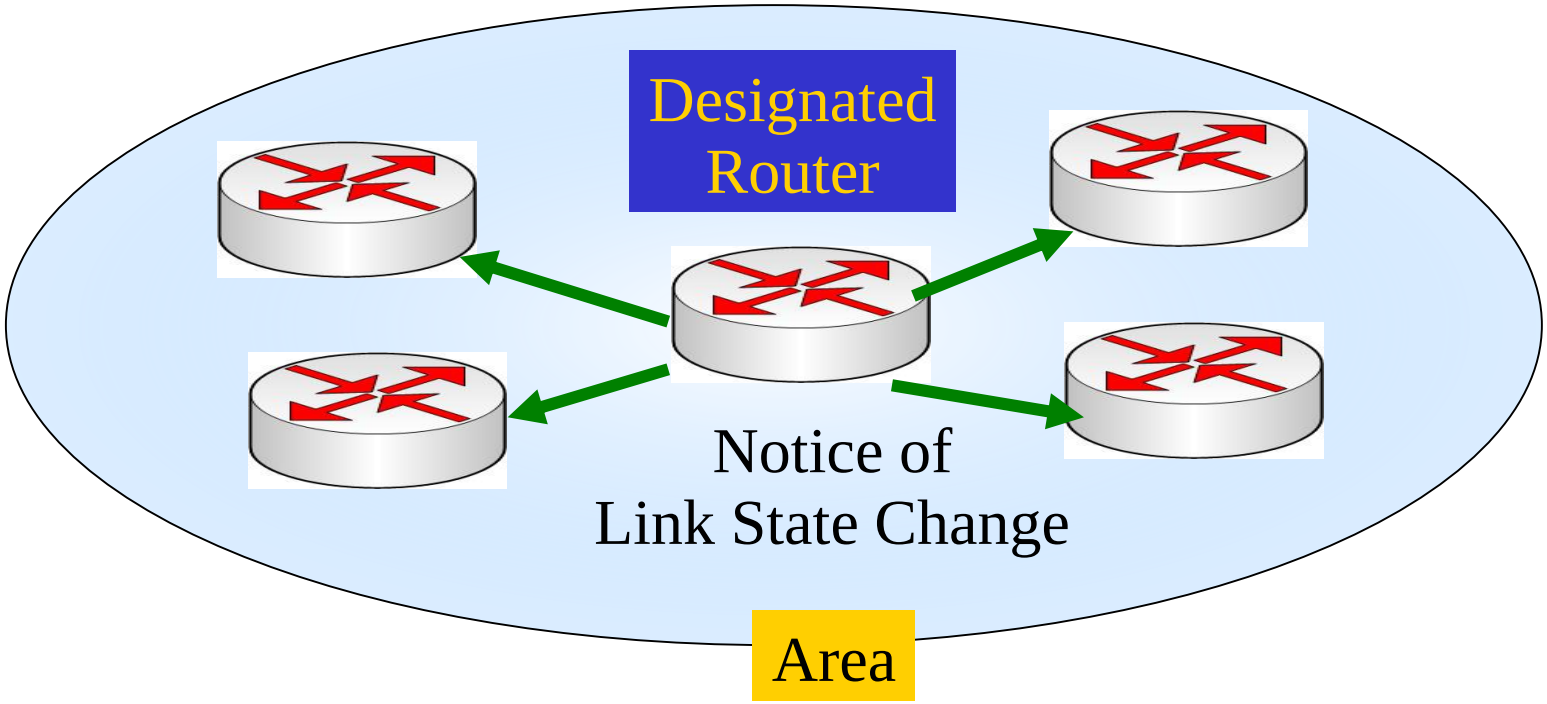

OSPF Routers

- Efficient
  - Only routers are informed (not hosts)
  - Usually only updates are transmitted, not whole tables

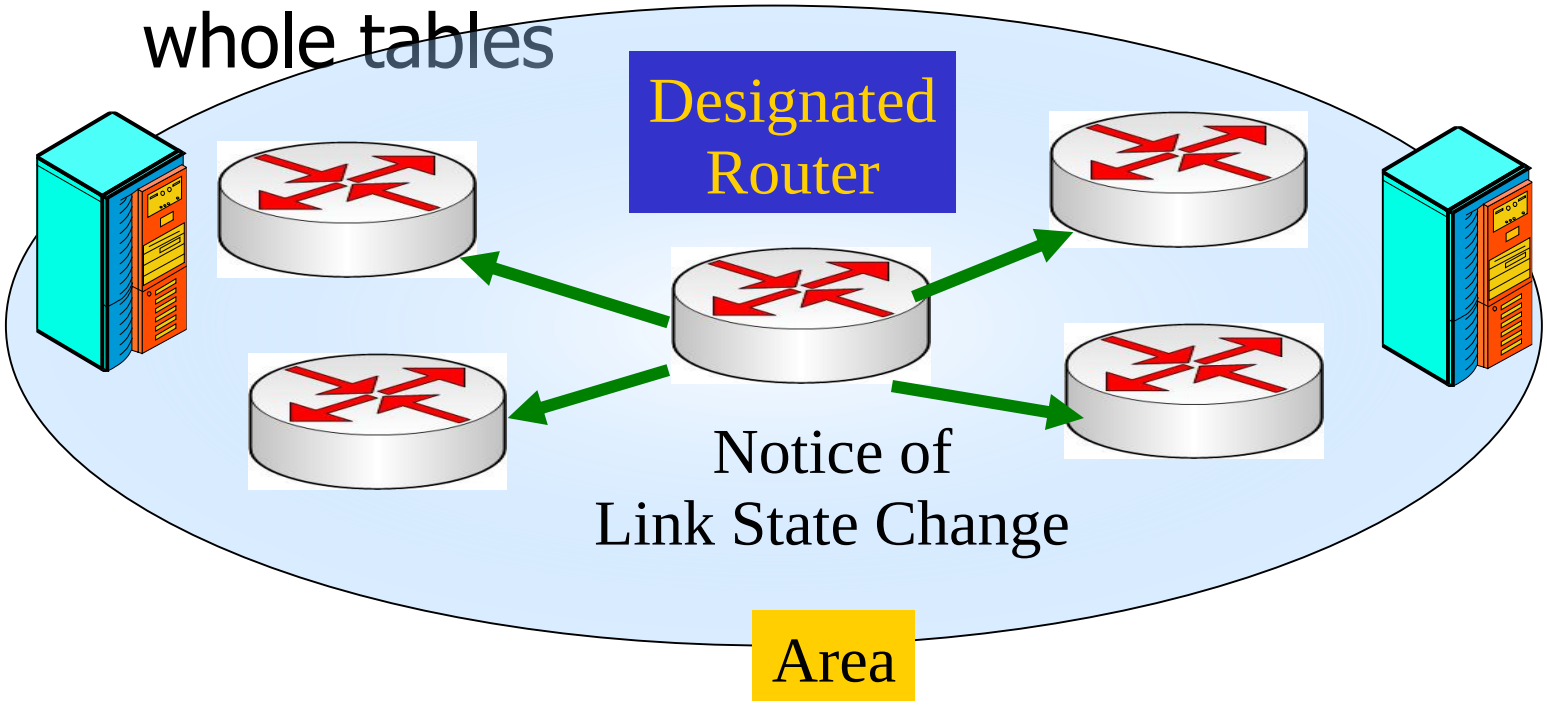

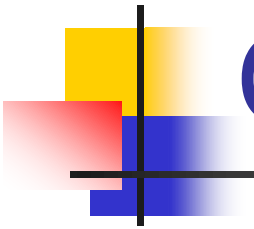

OSPF

- Fast Convergence
  - When a failure occurs, a router transmits the notice to the designated router
  - Designated router sends the information back out to other routers immediately

OSPF

- Carried in data field of IP packet
  - Encapsulation Protocol value is 89
- IP is unreliable, so OSPF messages do not always get through
- A single lost OSPF message does little or no harm

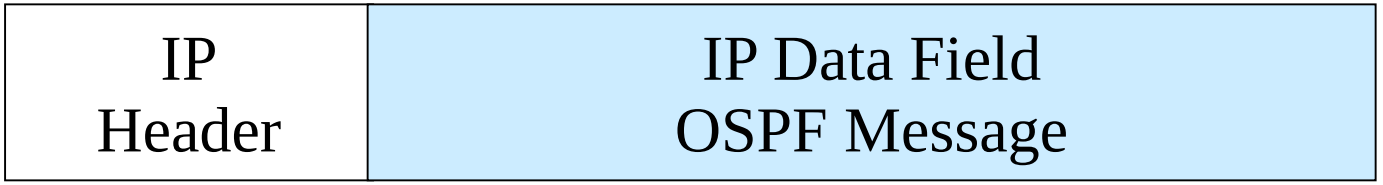

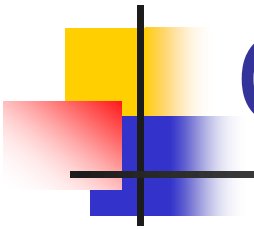

OSPF

- Typical link-state but with enhancements
  - Authentication of routing messages
  - Additional hierarchy (to help with scalability)
  - Load balancing

Configure OSPF

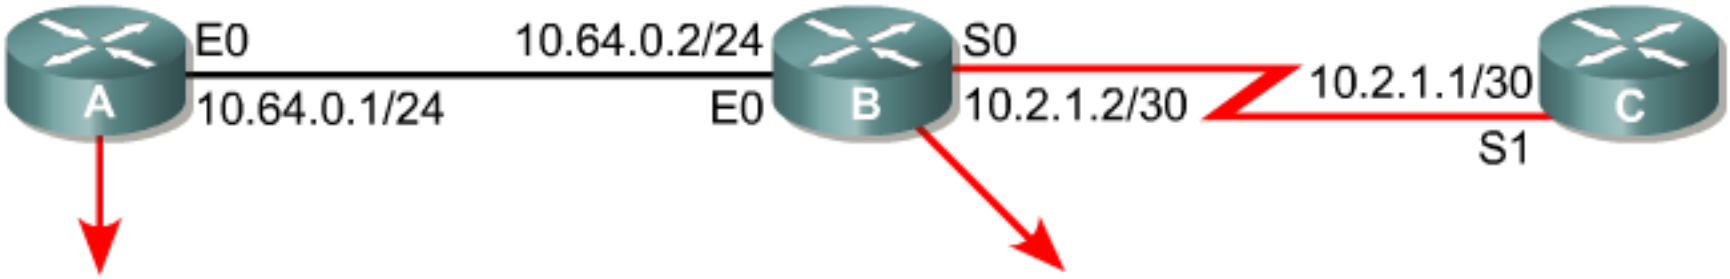


```
<Output Omitted>
interface Ethernet0
ip address 10.64.0.1 255.255.255.0
!
<Output Omitted>
router ospf 1
network 10.64.0.0 0.0.0.255 area 0
```

```
<Output Omitted>
interface Ethernet0
ip address 10.64.0.2 255.255.255.0
!
interface Serial0
ip address 10.2.1.2 255.255.255.252
<Output Omitted>
router ospf 1
network 10.2.1.0 0.0.0.3 area 0
network 10.64.0.0 0.0.0.255 area 0
```

Routing table @A

A#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O 10.2.1.0/30 [110/65] via 10.64.0.2, 00:05:36, FastEthernet0/0
C 10.64.0.0/24 is directly connected, FastEthernet0/0

Routing table @B

B#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.2.1.0/30 is directly connected, Serial0/0/0
C 10.64.0.0/24 is directly connected, FastEthernet0/0

Routing table @C

C#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.2.1.0/30 is directly connected, Serial0/0/1
O 10.64.0.0/24 [110/65] via 10.2.1.2, 00:00:01, Serial0/0/1

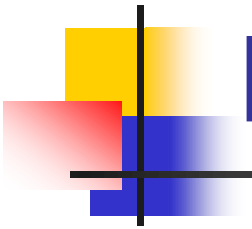

Link-State Summary

- Good
  - Converges relatively quickly
- Bad
  - Lots of information stored at each node because LSP for each node in network must be stored at each node (scalability problem)
  - Flooding of LSPs uses bandwidth
  - Potential security issue (if false LSP propagates)

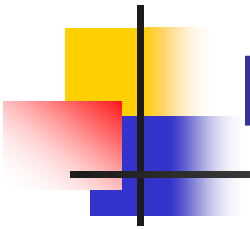

Distance Vector VS. Link-state

- Key philosophical difference
  - Distance vector talks only to directly connected neighbors and tells them what is has learned
  - Link-state talks to everybody, but only tells them what it knows

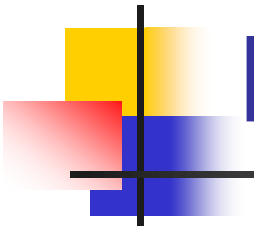

References

- "Gregory Kesden" lecture of 20-770 Communications and Networking
- "Nina Taft", The Basics of BGP Routing and its Performance in Today's Internet, Sprint.
- "Anonymous"
  - lecture of Addressing and Domain Name System, CS640
  - Telecom App2b
- Cisco CCNA Material
- "Jennifer Rexford", Internet Routing (COS 598A)

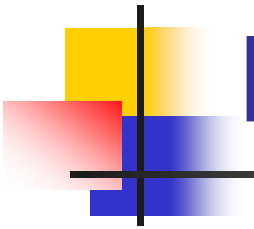

References

- Slide from “Agilent Technologies”
- Slide from “Anonymous” CS 332, Spring 2002
- Panko’s Business Data Networking and Telecommunications, 5th edition
- Cisco CCNA Course Material
- สถาปัตยกรรมและโปรโตคอลทีซีพี/ไอพี สุรศักดิ์ สงวนพงษ์