Dynamic routing

- Routing Algorithm (Dijkstra / Bellman-Ford) – idealization
  - All routers are identical
  - Network is flat.
    - Not true in Practice

- Hierarchical routing
  - Internet = network of networks
  - Each network admin may want to control routing in its own routing network.
  - Hierarchical routing solves
    - Scale problems.
    - Administrative autonomy.

Routing Protocols RIP, OSPF, BGP

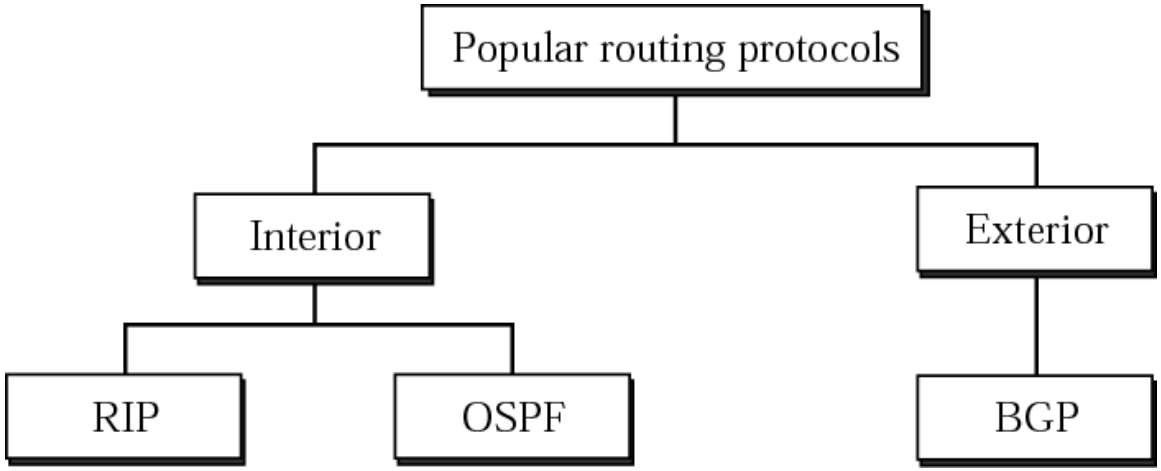

- Dynamic protocols
  - Sharing neighborhood information
- Use different metrics.
- RIP (one hop count, how many networks a packet crosses), Networks are treated equally
- BGP (depend on the policy, set by administrator)
- OSPF (TOS, minimize delay, maximize throughput)

Routing Information Protocol RFC 1058

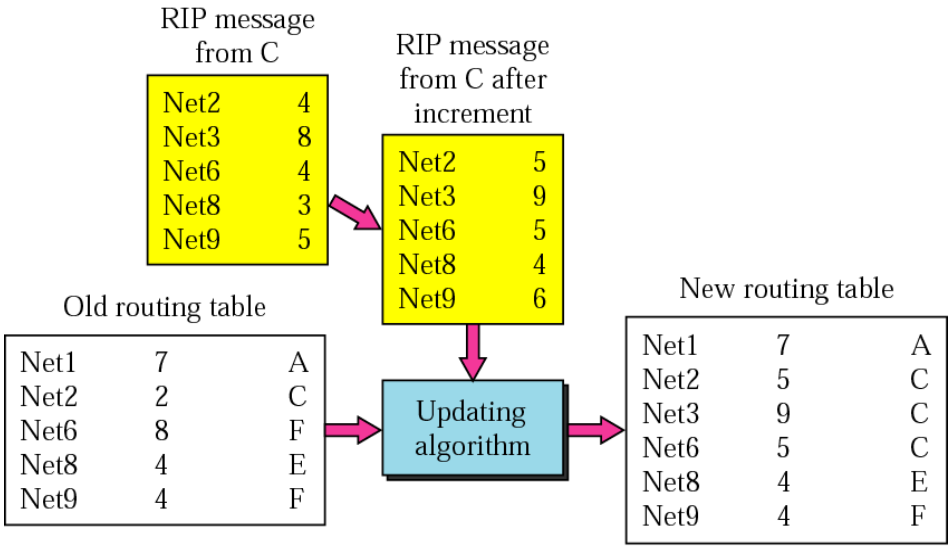

Net1: No news, do not change
Net2: Same next hop, replace
Net3: A new router, add
Net6: Different next hop, new hop count smaller, replace
Net8: Different next hop, new hop count the same, do not change
Net9: Different next hop, new hop count larger, do not change

- http://www.faqs.org/rfcs/rfc1058.html
- Distance vector algorithm (Bellman-Ford)
  - Sharing knowledge about the entire AS
  - Shares only with neighbors
  - Shares at regular intervals (different from DVA)

- Receive a RIP message (a response)
- Add one hop for each advertised dest
- Repeat
  - If (dest not in routing table)
    - Add the advertised info to the table
  - Else
    - If (next-hop is the same)
      - Replace with the advertised one
    - Else
      - If (advertised hop count < one in the table)
      - Replace entry in the routing table
- Return

Initializing and Updating Routing Table

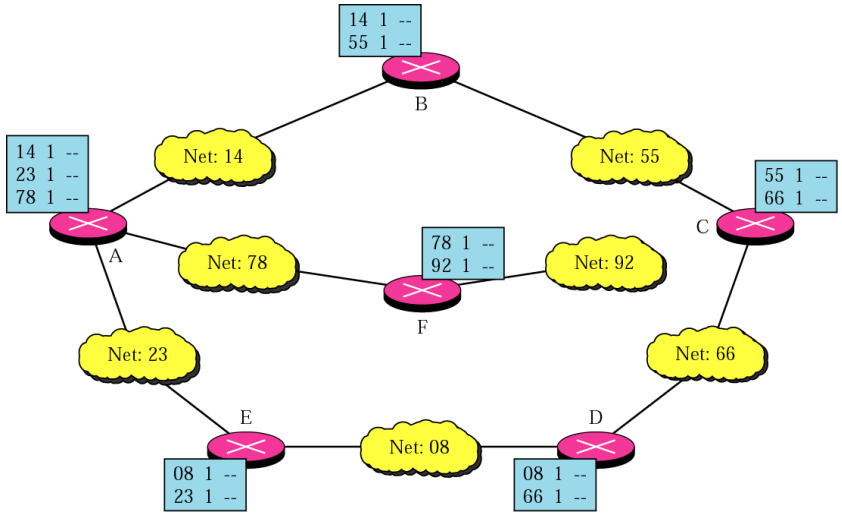

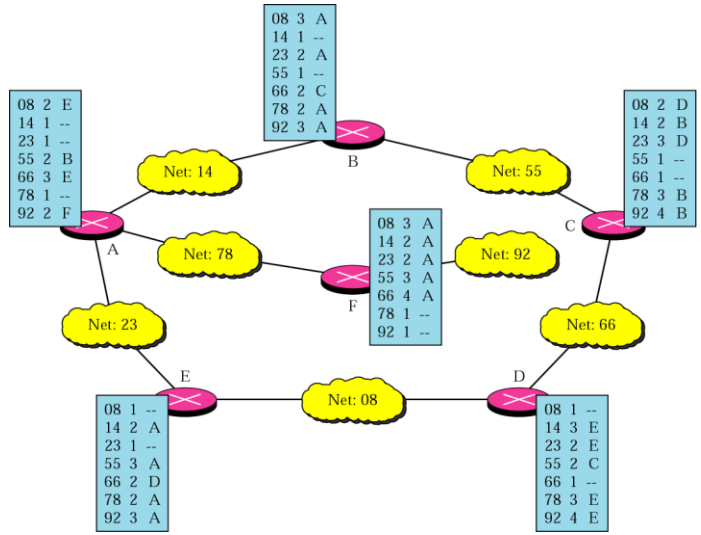

- Destination Metric Gateway
- Dest D G

- Rule to implement RIP based in DVA
- If Information is not from G
  - Update to reflect minimum cost
- If Information is from G
  - Always update

RIP message format

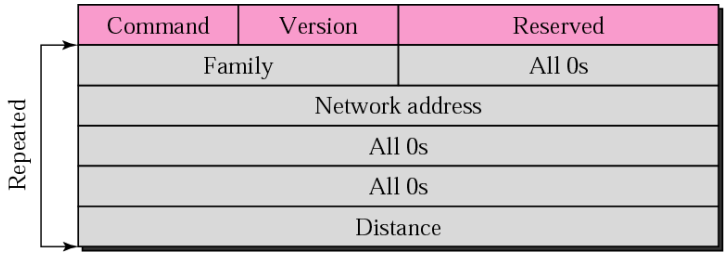

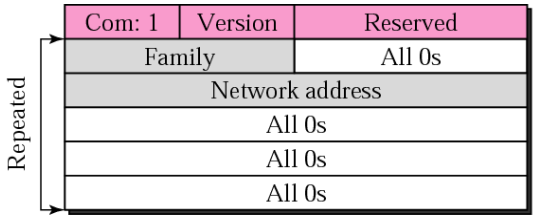

a. Request for some

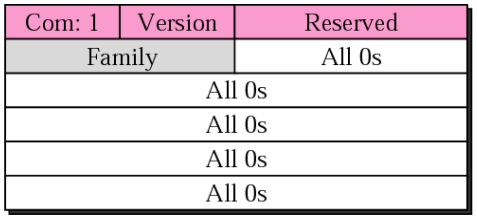

b. Request for all

- RIP messages uses UDP datagrams on port 520
  - Implemented in Unix systems by the ‘routed’ daemon.
- Size of datagram limited to 512 bytes (allow advertisement of 25 routes.

- Command: 8 bits
  - Request (1), reply (2)
- Version: 1 or 2
- Family: of protocol used for TCP/IP it is 2
- Network address : 32 bytes
- Distance: hop count from the advertising router to the destination network

- Response: solicited or unsolicited.

RIP message example

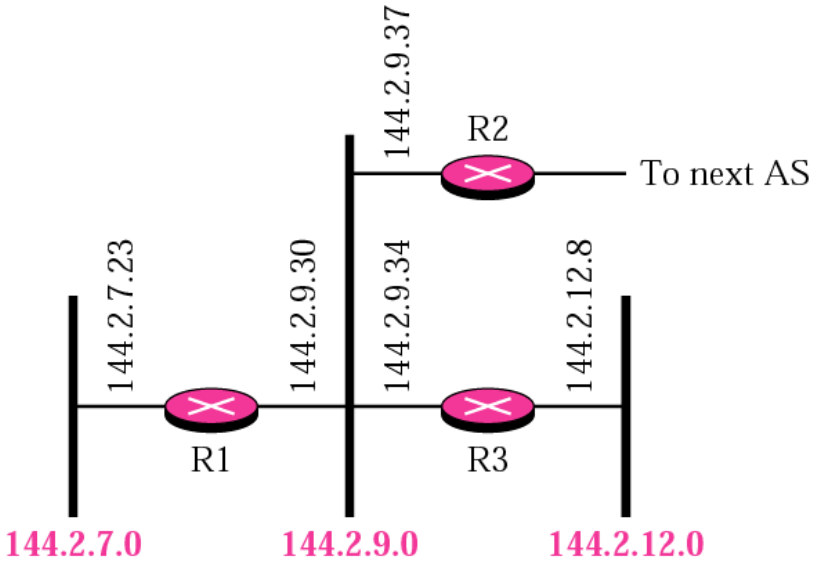

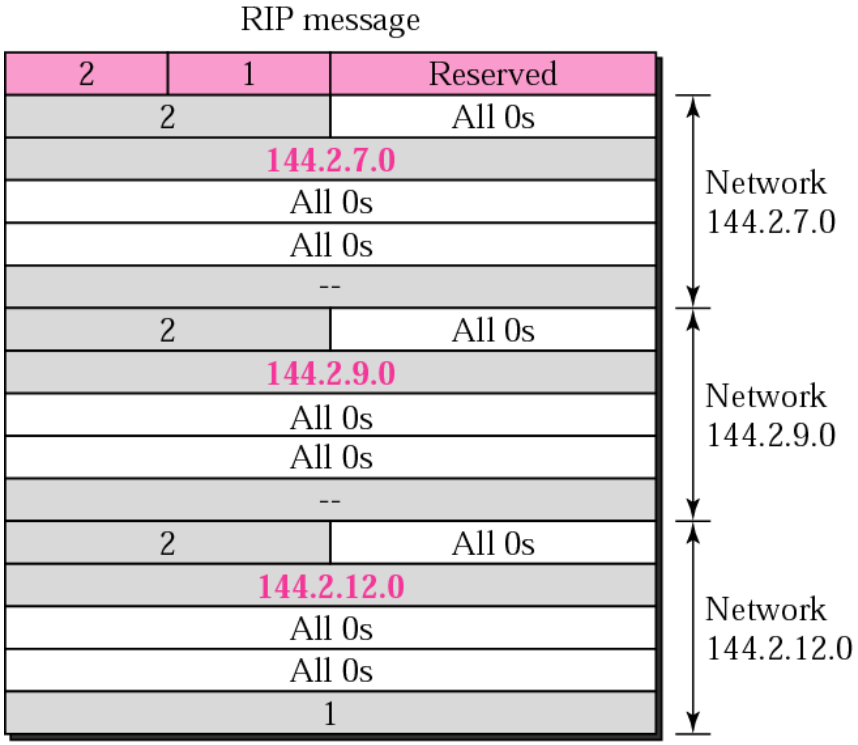

RIP timers

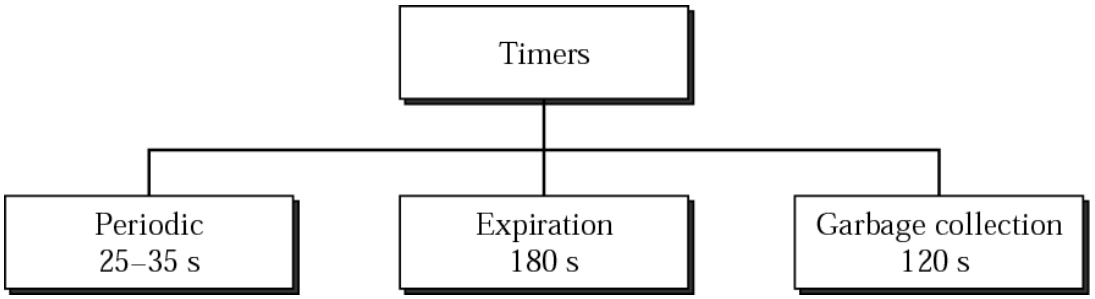

- Periodic timer: control advertising of regular update messages (25-35 sec)
- Expiration timer: governs the validity of a route (180 sec)
  - Every time an update (on a 30 sec average) is received the timer is reset
  - If no update received within this timer the metric is set to 16
- Garbage timer: 120 sec
  - A route can be advertised with a 16 metric for 120 sec before it get purged
    - Allow neighbors to have knowledge of the invalidity of a route

Slow Convergence Problem

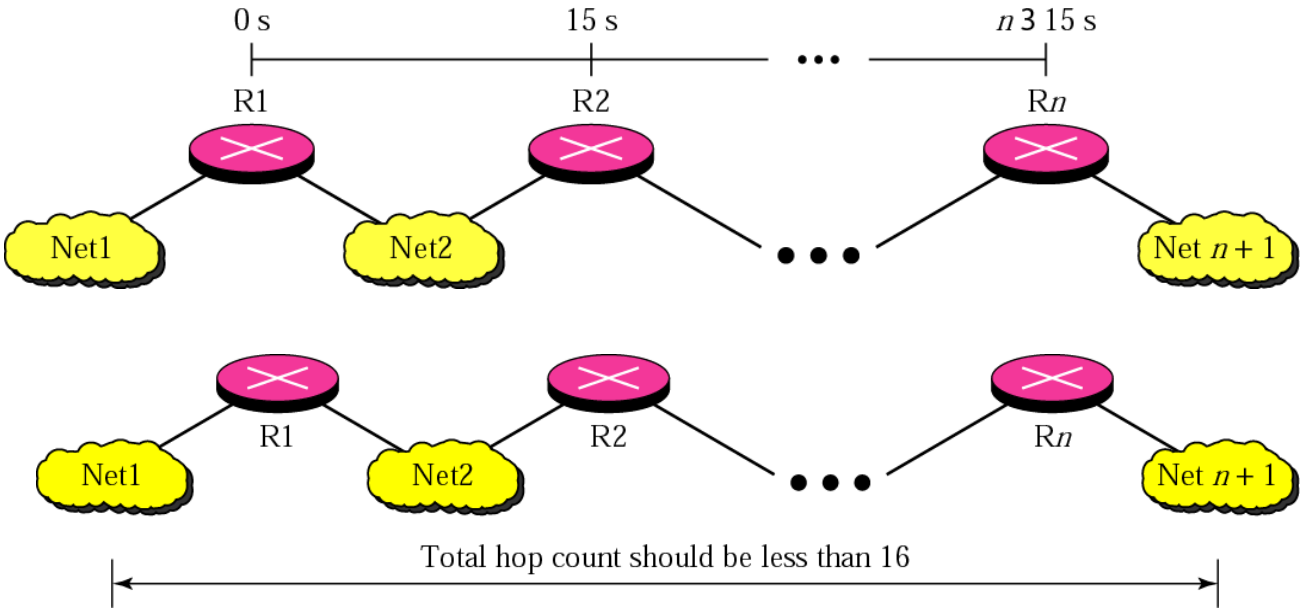

Solved by limiting number of hops to 15, 16 means
That the destination is unreachable

RIP Instability

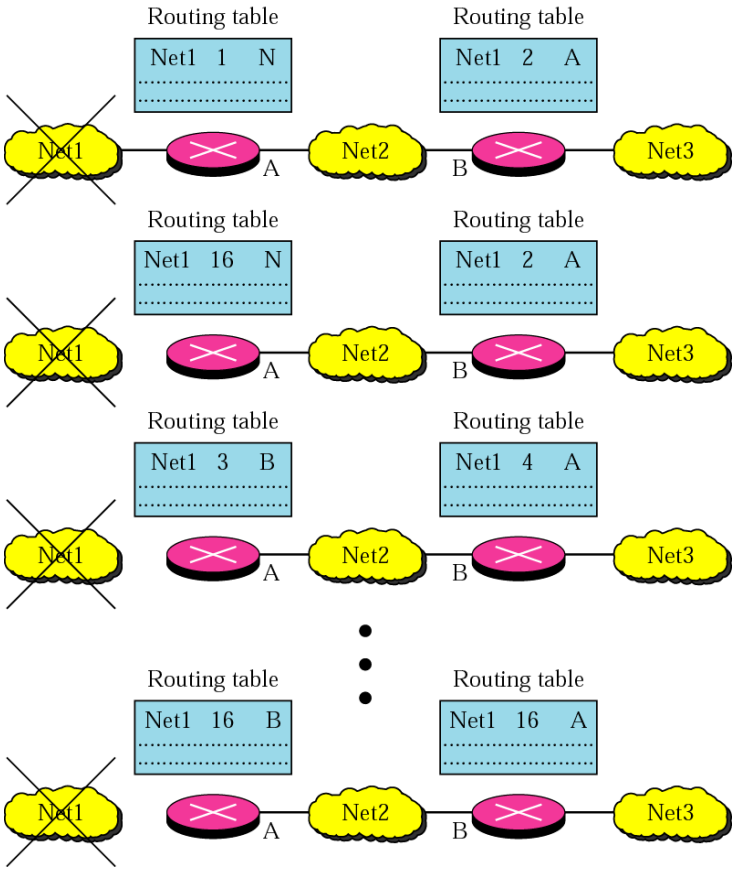

Triggered Updates

Split Horizons

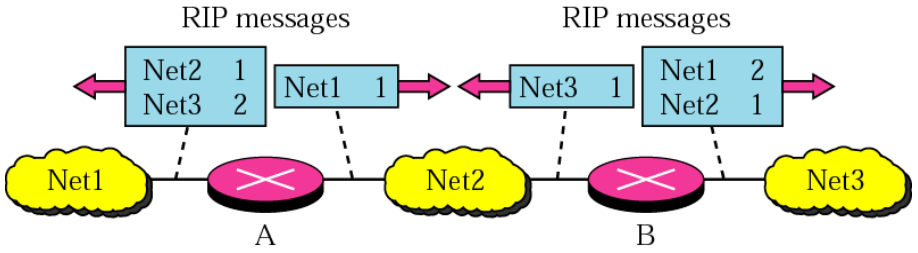

Poison reverse

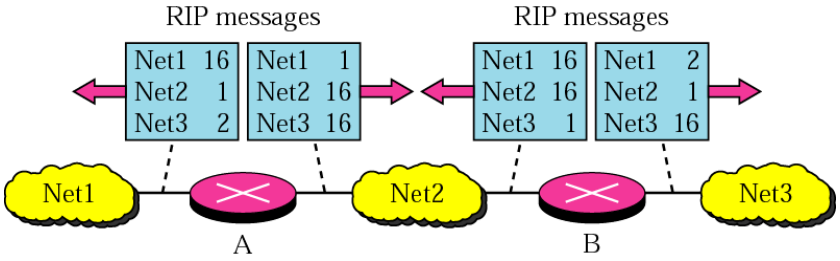

Border Gateway Protocol (RFC 1771)

- Based on the path vector routing.
- Distance-vector protocol not preferred for inter-AS routing (exterior routing protocol)
  - Assumes all routers have a common distance metrics to judge route preferences.
    - If routers have different meanings of a metric, it may not be possible to create stable, loop free routes.
  - A given AS may have different priorities from another AS.
  - Gives no information about the ASs that will be visited.
- Link-state routing protocol
  - Different metrics.
  - Flooding is not realistic.
- Path vector routing
  - No metrics,
  - Information about which networks can be reached by a given router and ASs to be crossed.
- Differs from DVA
  - Path vector approach does not include a distance or cost estimate
  - Lists all of the ASs visited to reach destination network.

BGP (continued)

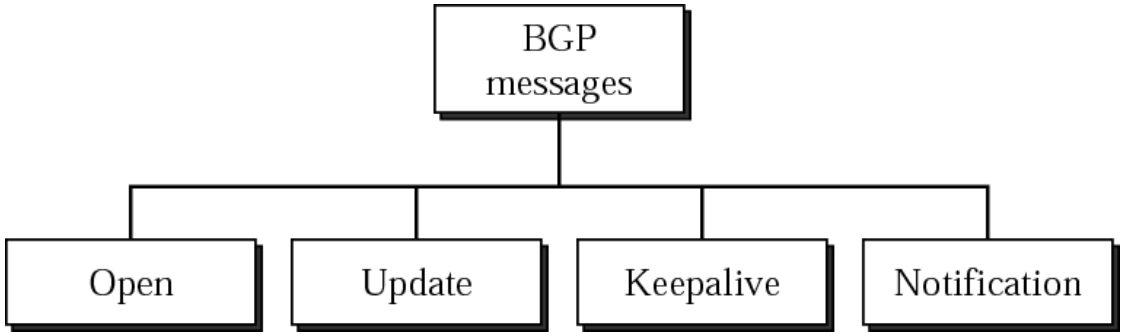

- Messages are sent over TCP connections on port 179.
- Functional procedures
  - Neighbor acquisition (open message, acceptance through Keepalive message)
  - Neighbor reachability (periodic Keepalive messages)
  - Network reachability (broadcast an update message)
    - Each routers maintains a database of networks that can be reached
    - + preferred route to this network.
- RFC does not address
    - How a router knows the address of another router.
    - Up to network admin.

BGP (cont.)

Example of Network Reachability

Network	Next router	Path
N1	R1	AS14,AS23,AS67
N2	R5	AS22,AS67,AS5,AS89
N3	R6	AS67,AS89,AS9,AS34
N4	R12	AS62,AS2,AS9

- Loop Prevention in BGP:
  - Checks the Path before updating its database. (If its AS is in the path ignore the message)
- Policy Routing:
  - If a path consist of an AS against the policy of the current AS, message discarded.

Example of Message adverstisements

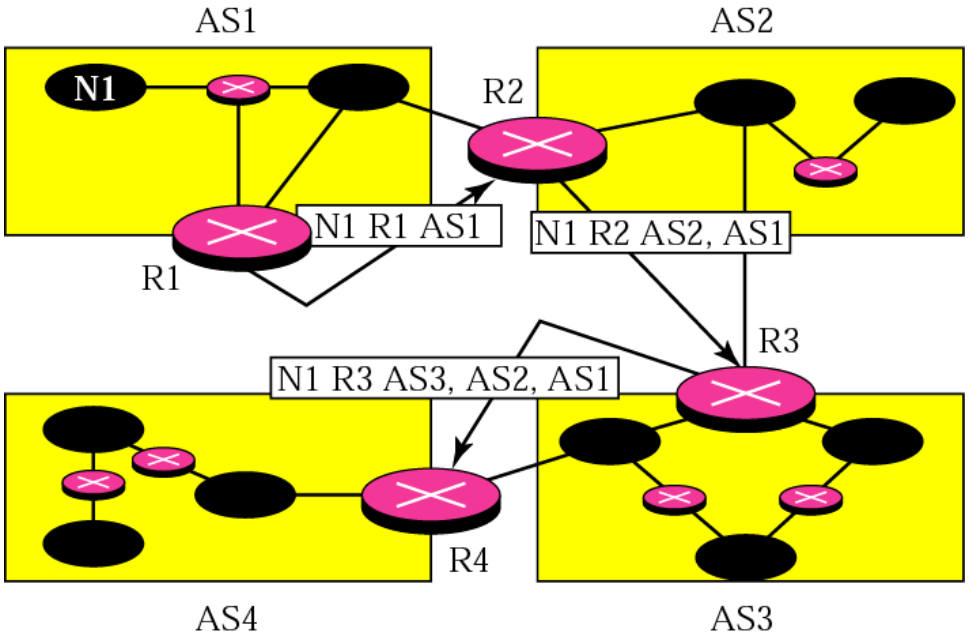

BGP message format (Open, Keepalive, Update, Notification

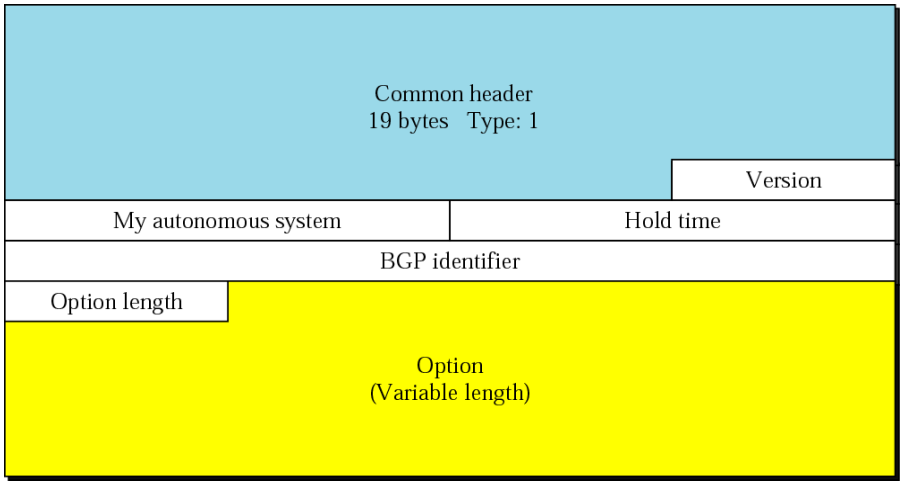

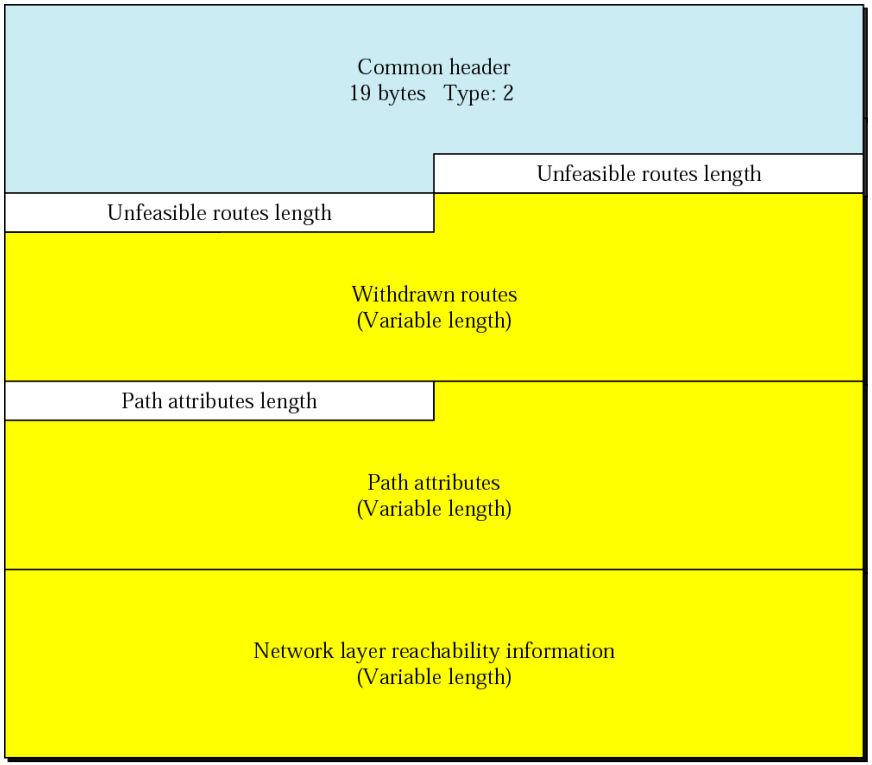

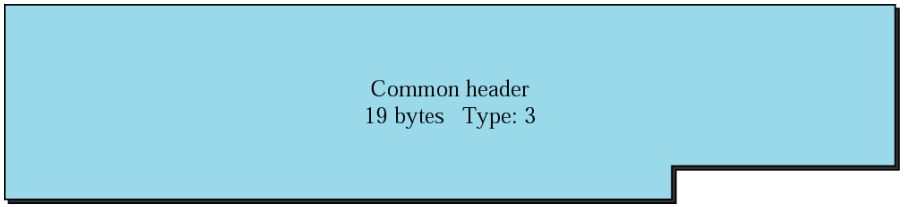

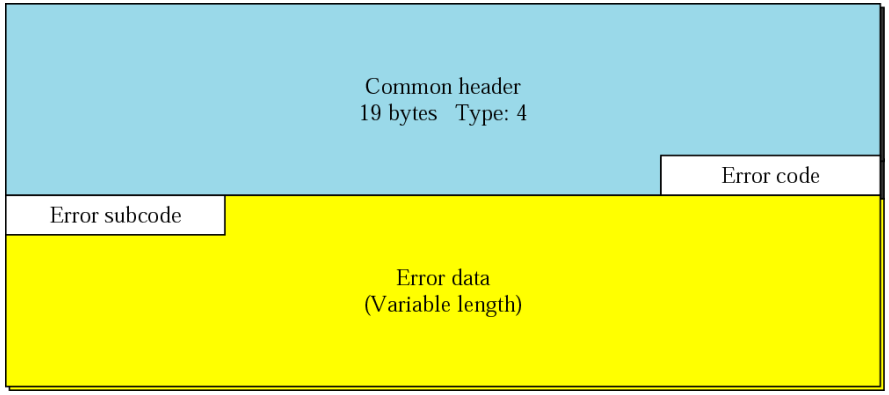

Open Shortest Path First (RFC 1247)

- Uses IP, has a value in the IP Header (8 bit protocol field)
- Interior routing protocol, its domain is also an autonomous system
- Special routers (autonomous system boundary routers) or backbone routers responsible to dissipate information about other AS into the current system.
- Divides an AS into areas
- Metric based on type of service
  - Minimum delay (rtt), maximum throughput, reliability, etc..

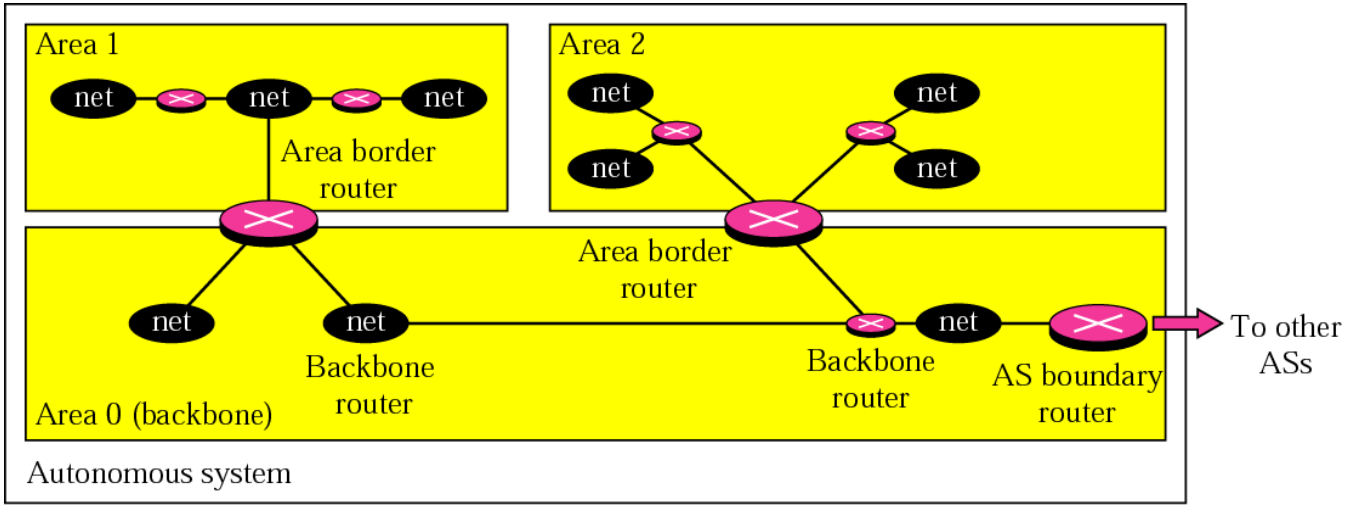

OSPF (type of links)

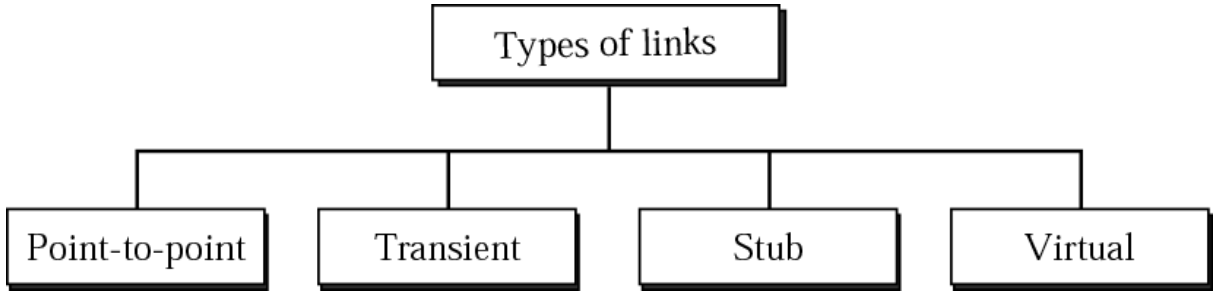

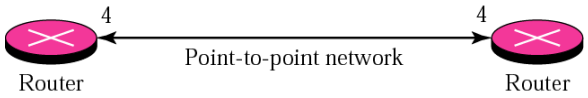

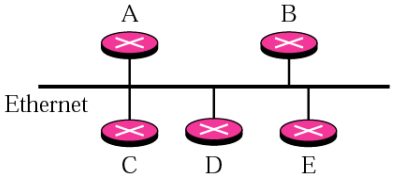

a. Transient network

a. Stub network b. Representation

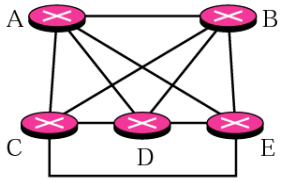

b. Unrealistic representation

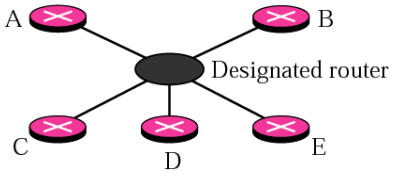

c. Realistic representation

OSPF (link state advertisement)

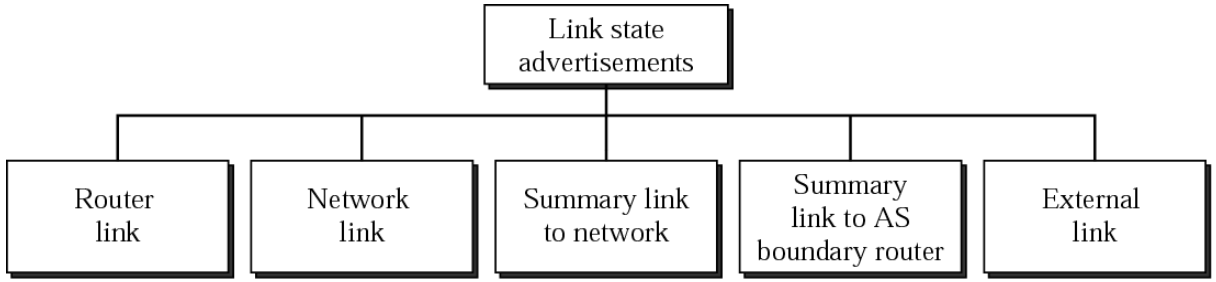

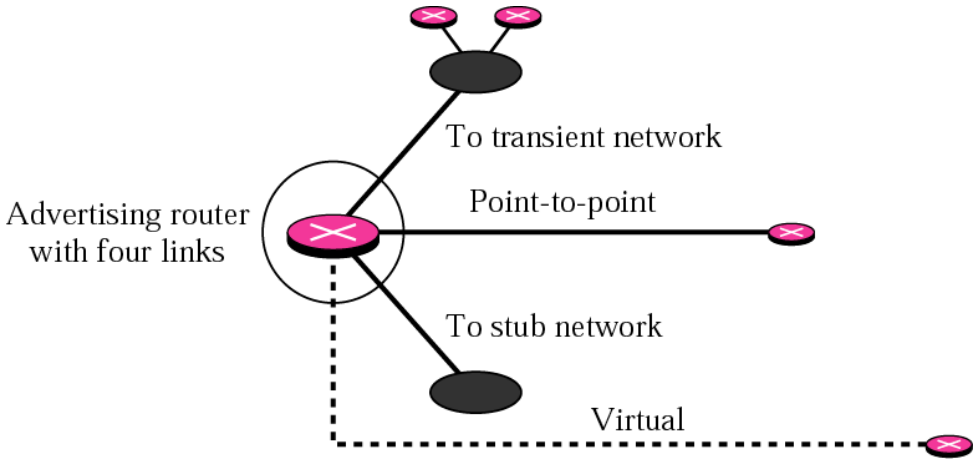

Router Link

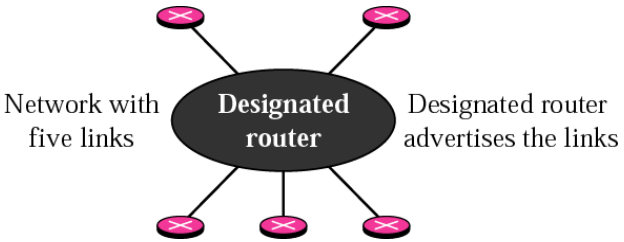

Network Link

OSPF (LSA cont.)

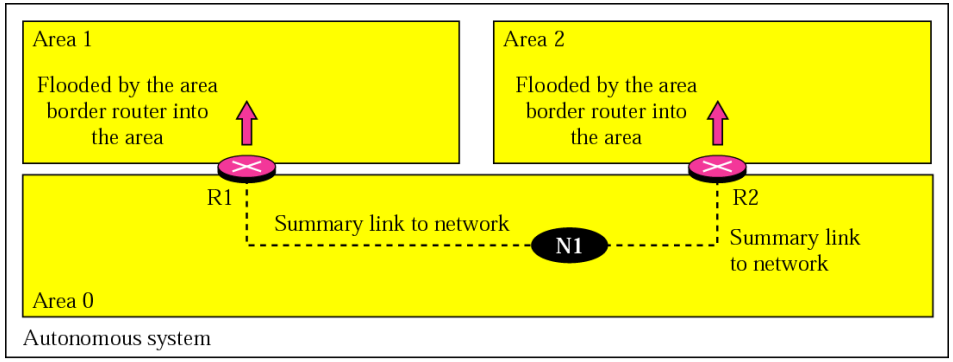

Summary link to Network

Summary link to AS boundary router

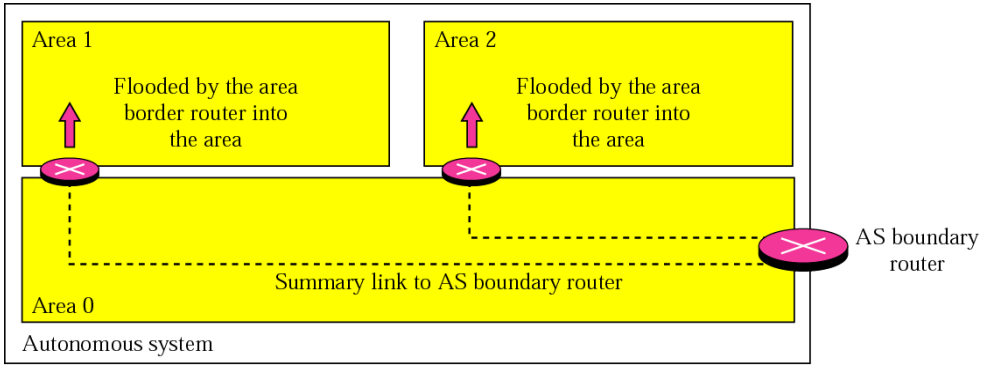

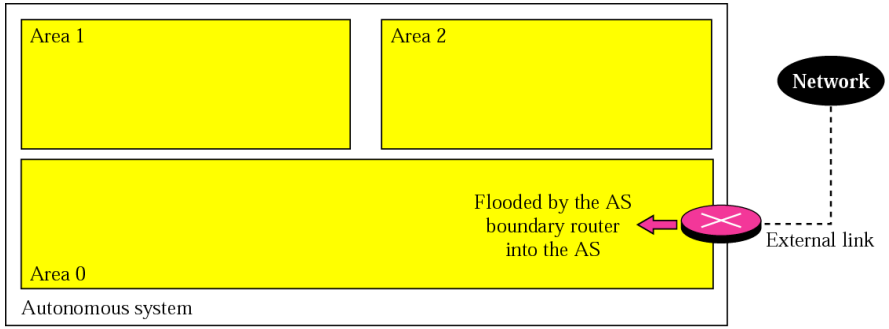

External Link

Shortest Path Calculation

a. Start with A

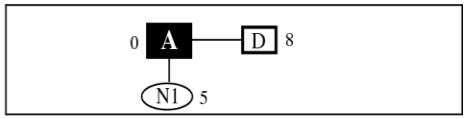

b. Make A permanent, add its neighbors

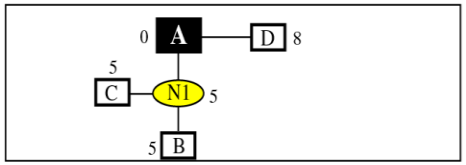

c. Make N1 permanent, add its neighbors

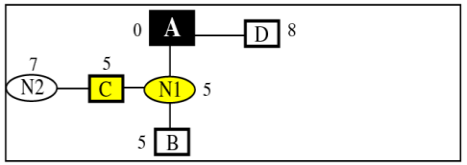

d. Make C permanent, add its neighbors

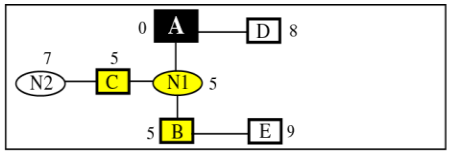

e. Make B permanent, add its neighbors

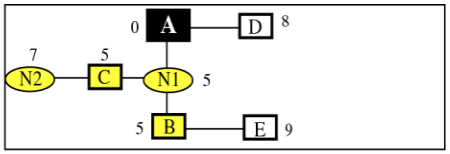

f. Make N2 permanent

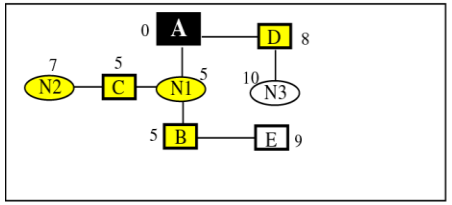

g. Make D permanent, add its neighbors

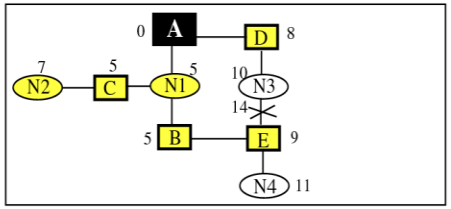

h. Make E permanent, add its neighbors

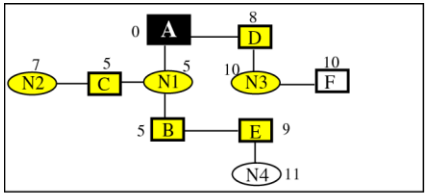

i. Make N3 permanent, add its neighbors

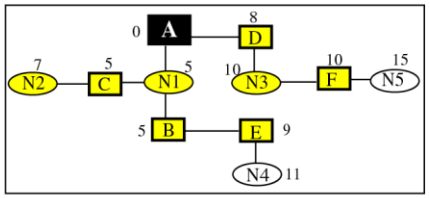

j. Make F permanent, add its neighbors

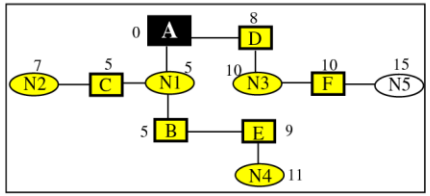

k. Make N4 permanent

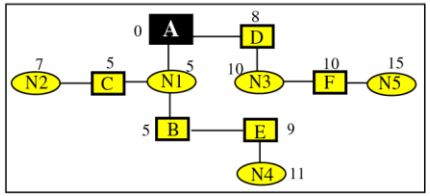

l. Make N5 permanent

Types of OSPF packets and header format

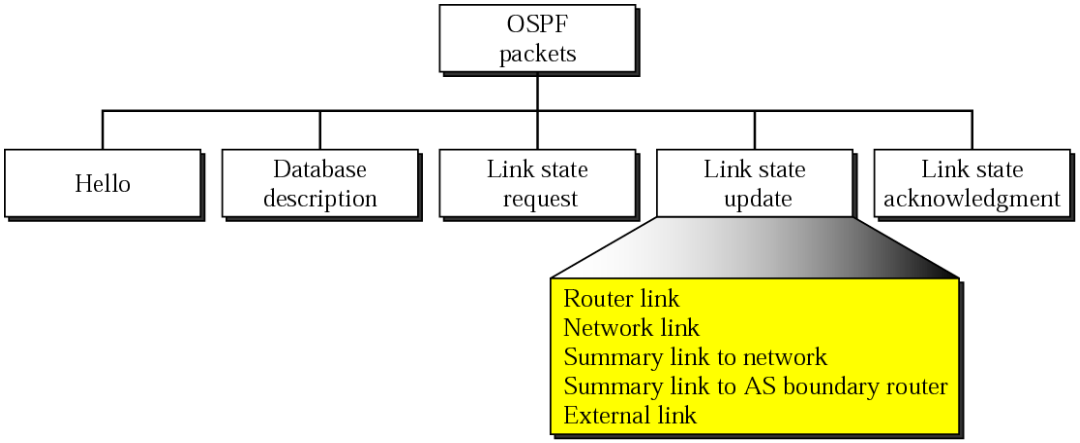

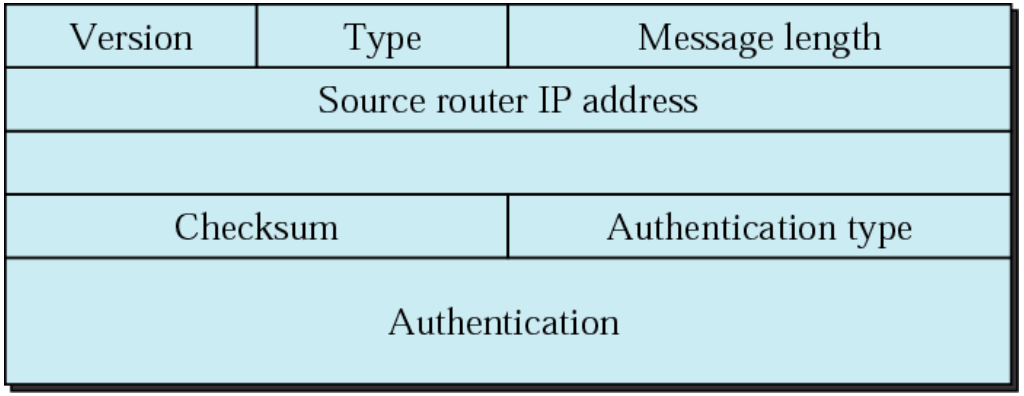

Link State Update Packet
A router link example

LSA header not covered
Refer to RFC 1247

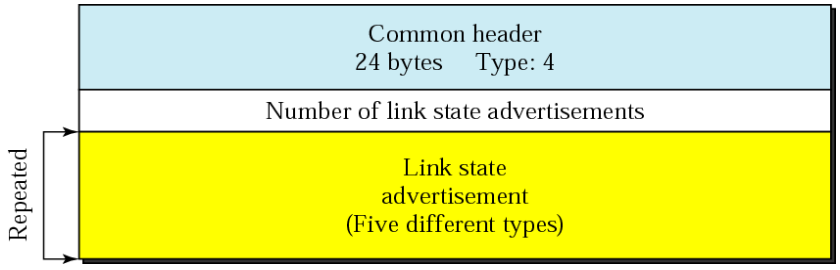

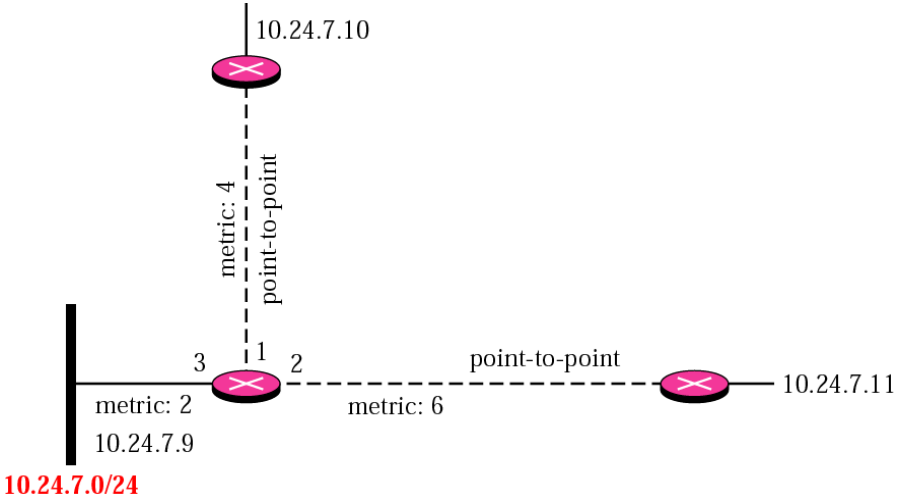

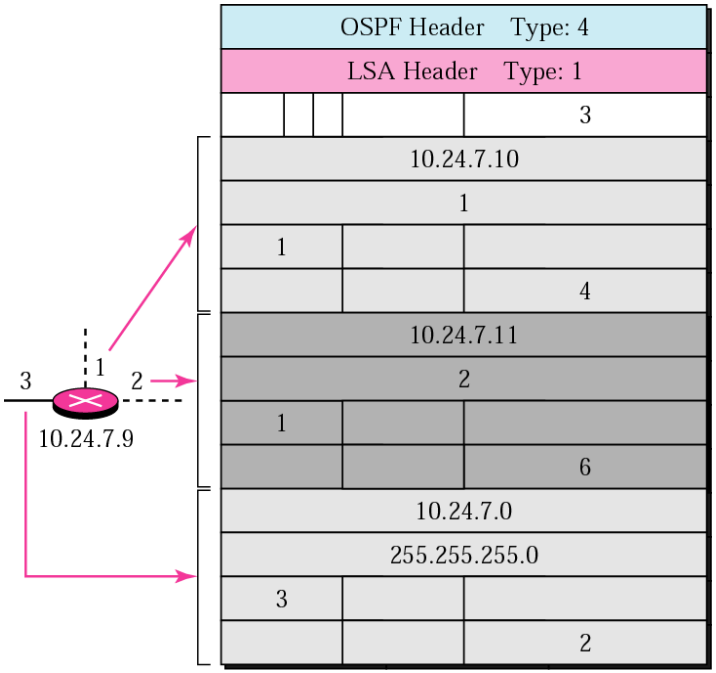

A Network Link Example

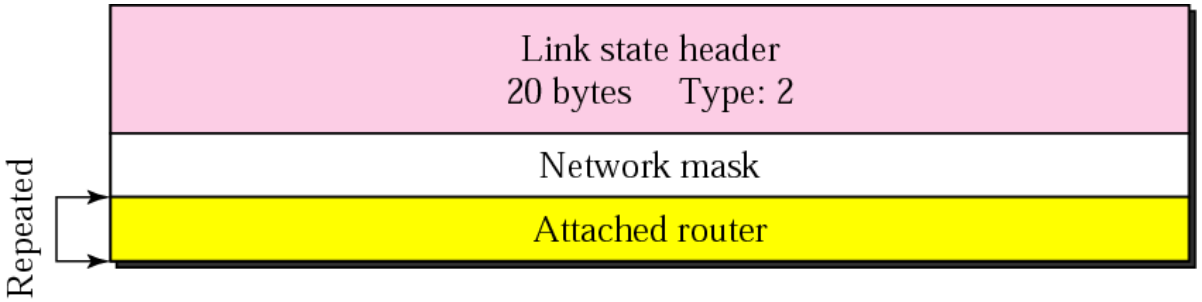

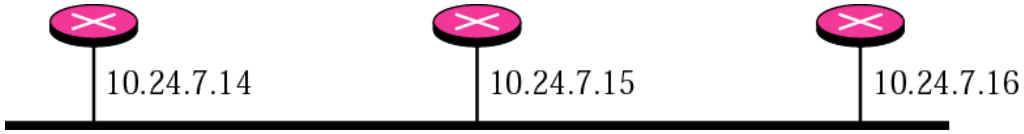

OSPF Header Type: 4
LSA Header Type: 2
255.255.255.0
10.24.7.14
10.24.7.15
10.24.7.16

Summary Links state Advertisements

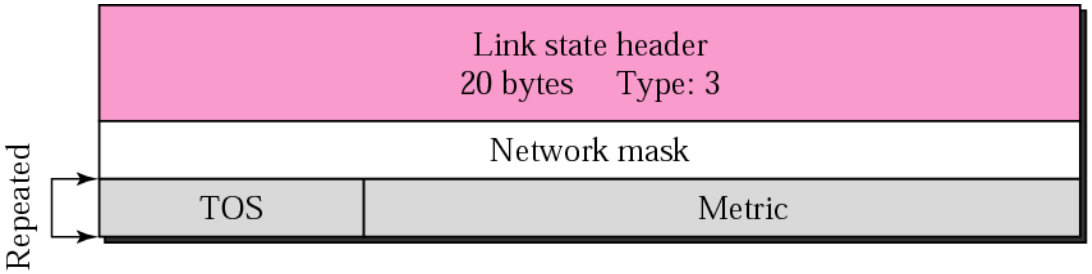

Summary link to network

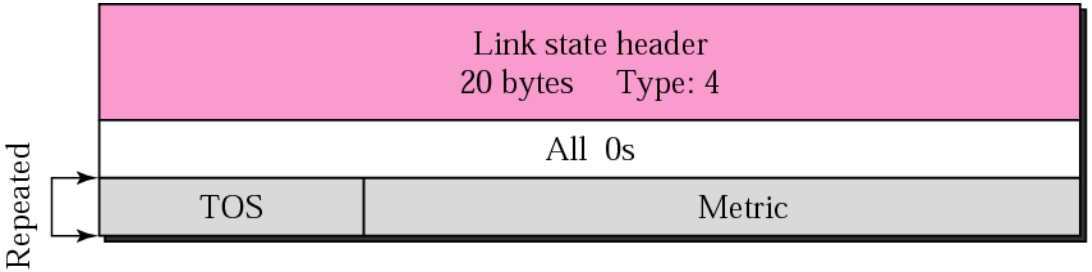

Summary link to AS boundary

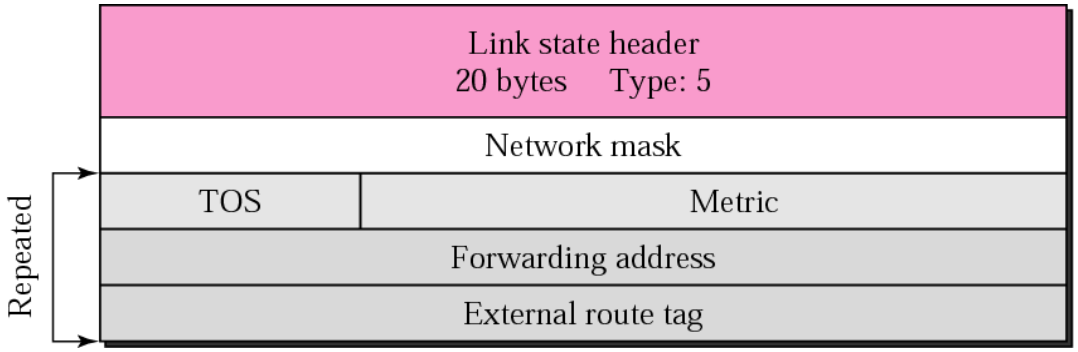

External Link