CPSC 317 COMPUTER NETWORKING

Module 5: Network Layer – Day 4 – Routing

1

READING

- Reading: 5.1, 5.2, 5.2.1, 5.2.2

LEARNING GOALS

Intra-domain routing I

- Explain why different routing algorithms are used within and between ASes
- Explain what link state routing is
- Execute Dijkstra's algorithm on a routing graph

RECAP

The network layer has two sub-layers:

- Data plane a.k.a., forwarding function
  - Local to one router
  - Forwards datagrams based on forwarding table
- Control plane a.k.a., routing function
  - Global across all routers
  - Decides how to route datagrams and updates forwarding table

RECAP: FORWARDING TABLE

- Each router keeps a forwarding table indicating, for each IP range, the link to be used
- Formally, the Forwarding Information Base (FIB)
- Built from aggregating information retrieved from other routers (routing table, or RIB)
- Routers use longest prefix matching to choose the outgoing link based on the destination address for each incoming datagram

KEY CHALLENGE FOR FORWARDING

- Imagine a 32 port (interface) router
- Each interface has 40Gbps bandwidth
- How long does it have to forward each datagram?
- 40Gbps = $5 * 10^9$ bytes per second
- Say the average datagram is 500 bytes
- $5000 * 10^6$ bytes/second / 500 bytes/datagram =
- $10*10^6$ datagrams per second per interface * 32 interfaces
- = $320 * 10^6$ datagrams per second or 3 ns per datagram

ROUTING PROTOCOL

Goal: determine "good" paths (equivalently, routes), from a sending host to a receiving host, through a network of routers

- path: sequence of routers that packets will traverse in going from given initial source host to given final destination host
- "good": least "cost", "fastest", "least congested"

ALTERNATIVES FOR DECIDING ROUTING

- Configured statically by network operators
- Dynamically by having the routers exchange messages (i.e., routing protocol)
  - Construct an entire network map
  - Distributed and decentralized
  - On failure we want to converge to the same view of the network
- Configured centrally by software (Software Defined Networking)

ROUTING

Two levels of routing:

- **Routing within a single AS**, under the control of a single administrative entity
  - Interior Gateway Protocols – IGP
- **Routing between different AS**, where we have no control over the routing policies of the other AS
  - Exterior Gateway Protocols – EGP

INTERIOR GATEWAY PROTOCOLS (IGP)

- Link State
  - Broadcast link information to all routers
  - Create a “map” of the network
- Distance Vector
  - Send local information to neighbouring routers
  - Create “signposts” for routing
- Examples:
  - OSPF, IGRP, RIP

LINK STATE

- Each router tells every other router about all of its links
- This gives every router complete information about the entire network
- Every so often, each router uses Dijkstra's algorithm to find the shortest path to all the other routers
- It then updates its forwarding table

A DIGRESSION – DIJKSTRA'S ALGORITHM

- Represent the network as a graph
- Nodes correspond to routers
- Networks are attached to routers, but not explicitly shown
- Label on the edge corresponds to the "cost" of using a link between routers
- Compute the lowest cost path from a single node in a graph to all other nodes
- Dynamic programming

AN EXAMPLE NETWORK GRAPH

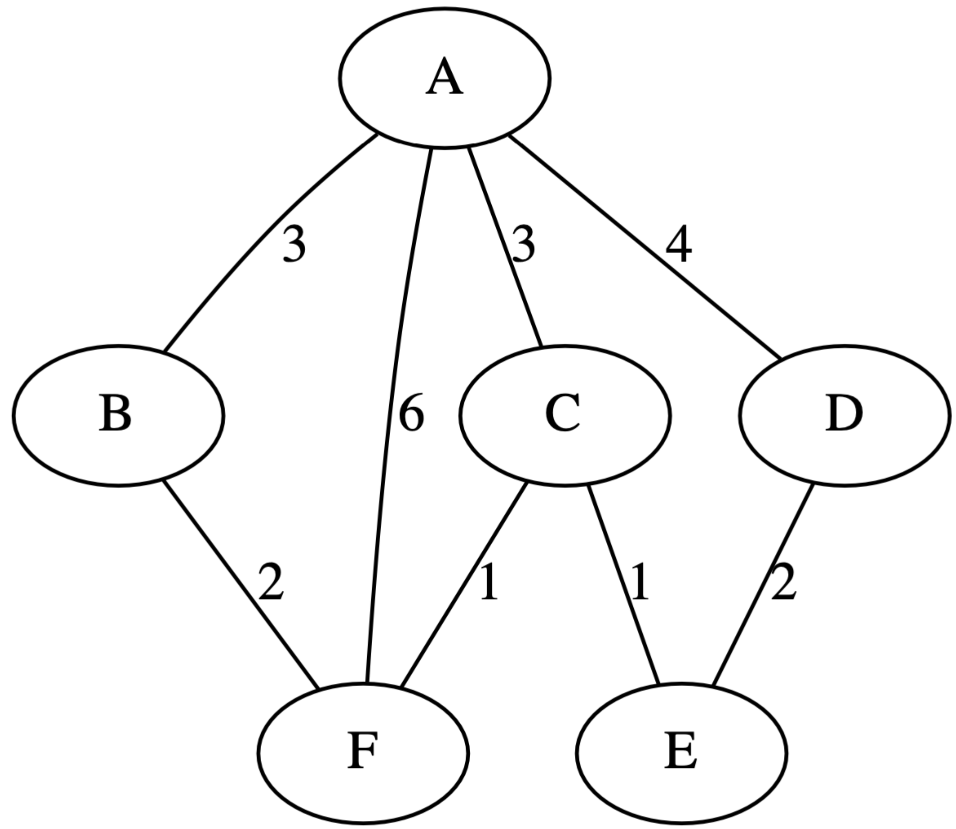

DIJKSTRA'S ALGORITHM

Initialize all Costs to infinity and all Prevs to undefined

For each neighbour of the source, enter Cost and Prev

Loop until all nodes are in Done

    Choose a min cost node not in Done, call it X

    Add X to Done

    For all neighbours of X

        If path through X is cheaper, update Cost and Prev

DIJKSTRA'S: INITIALIZE

Done:

Dest	Cost	Prev
B	-	-
C	-	-
D	-	-
E	-	-
F	-	-

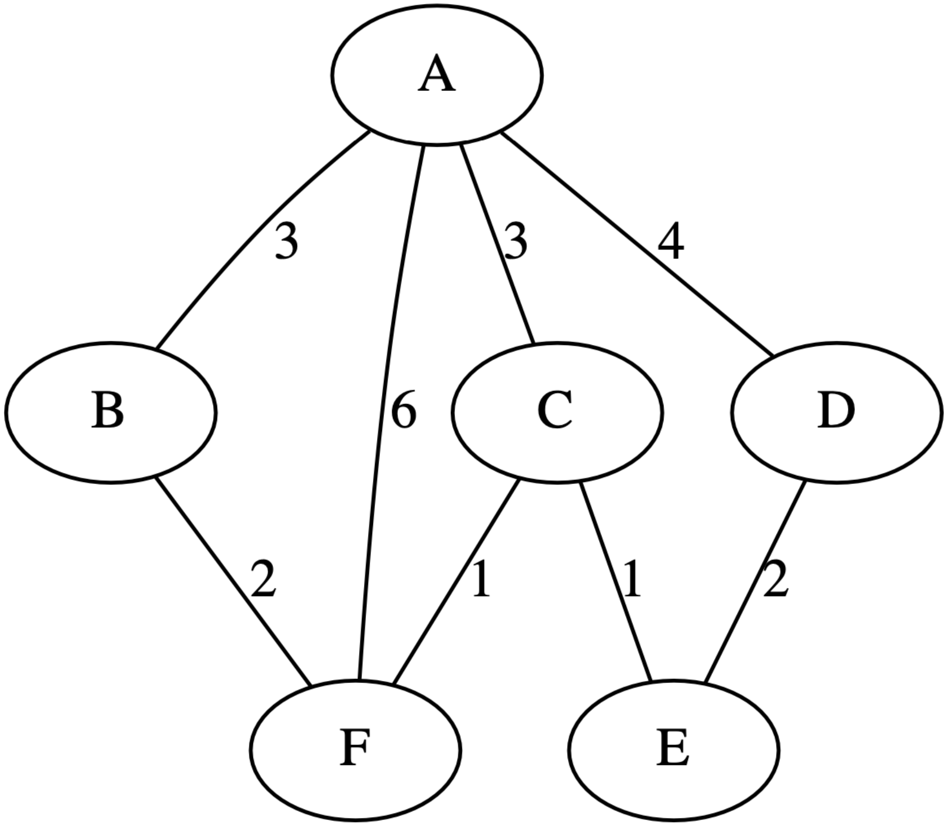

DIJKSTRA'S – DIRECT NEIGHBOURS

Done: A

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	-	-
F	6	A

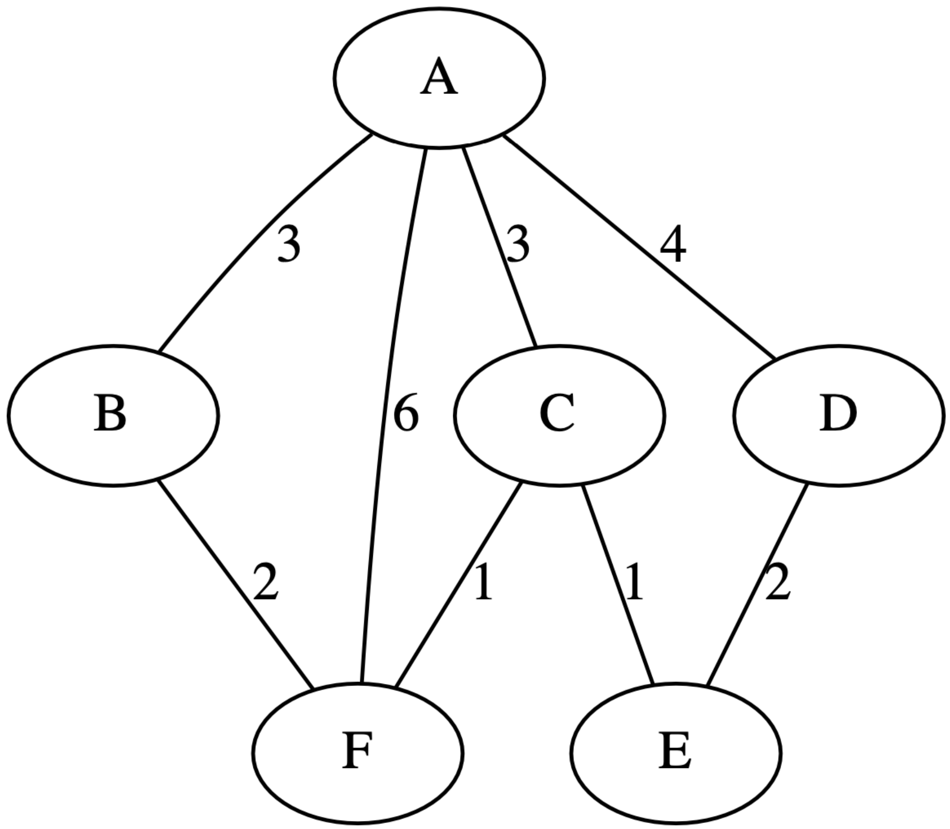

DIJKSTRA'S: ONE STEP – X = B

Done: A B

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	-	-
F	6	A

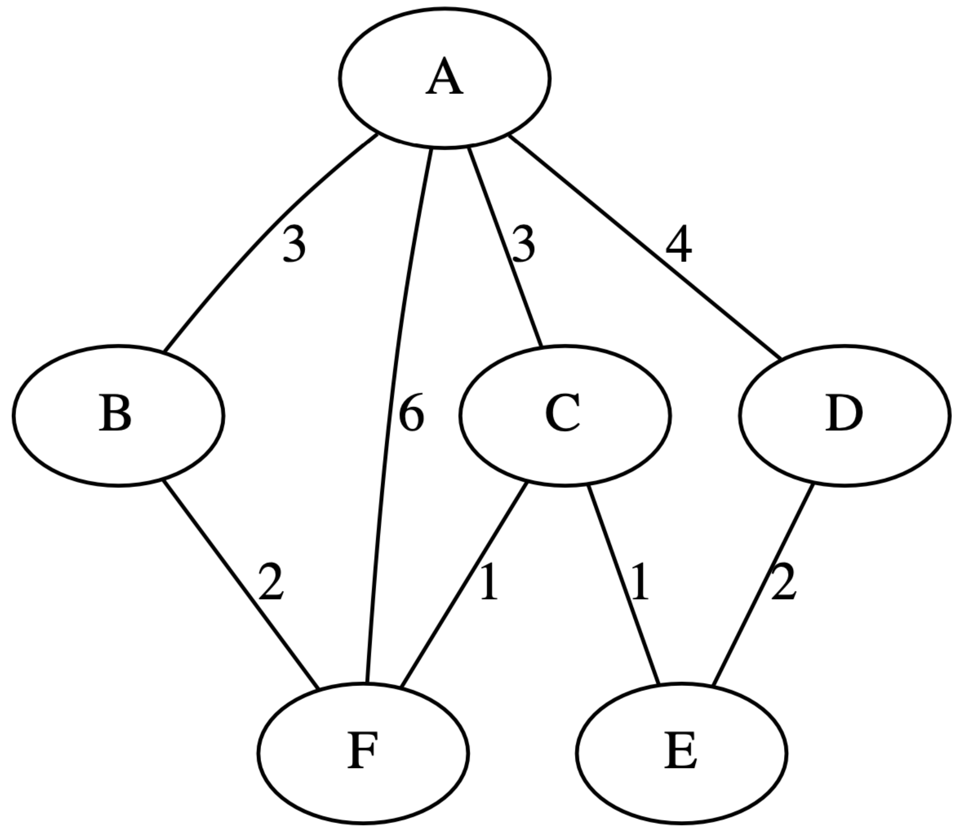

DIJKSTRA'S: ONE STEP – X = B

Done: A B

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	-	-
F	5	B

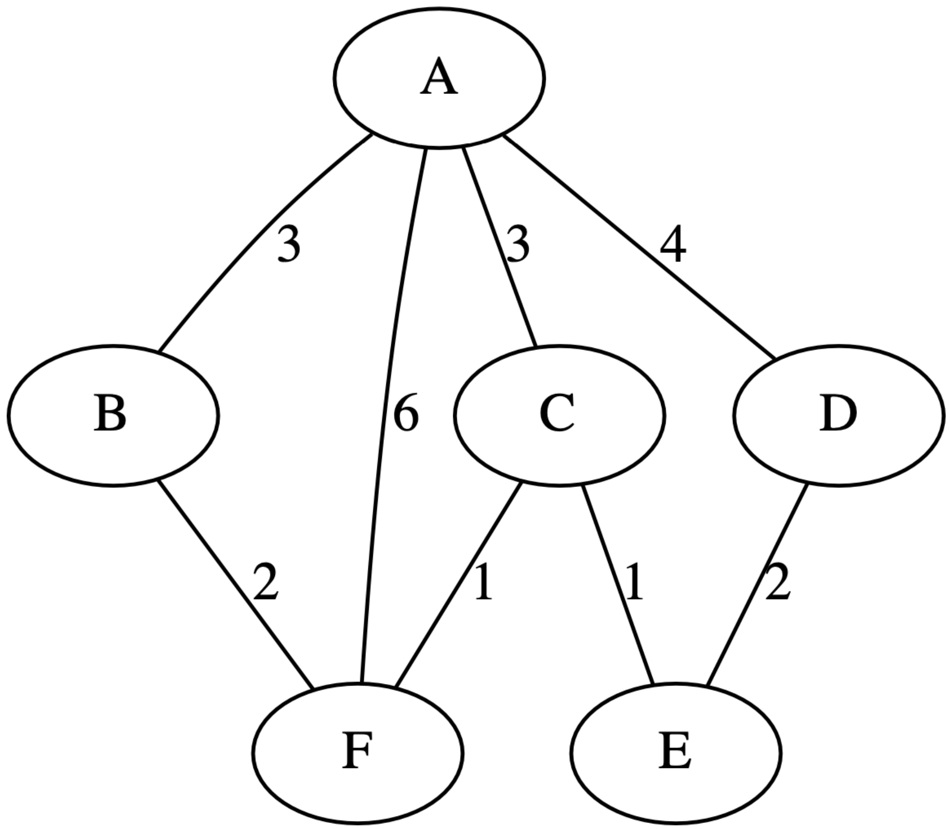

DIJKSTRA'S: ONE STEP – X = C

Done: A B C

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	-	-
F	5	B

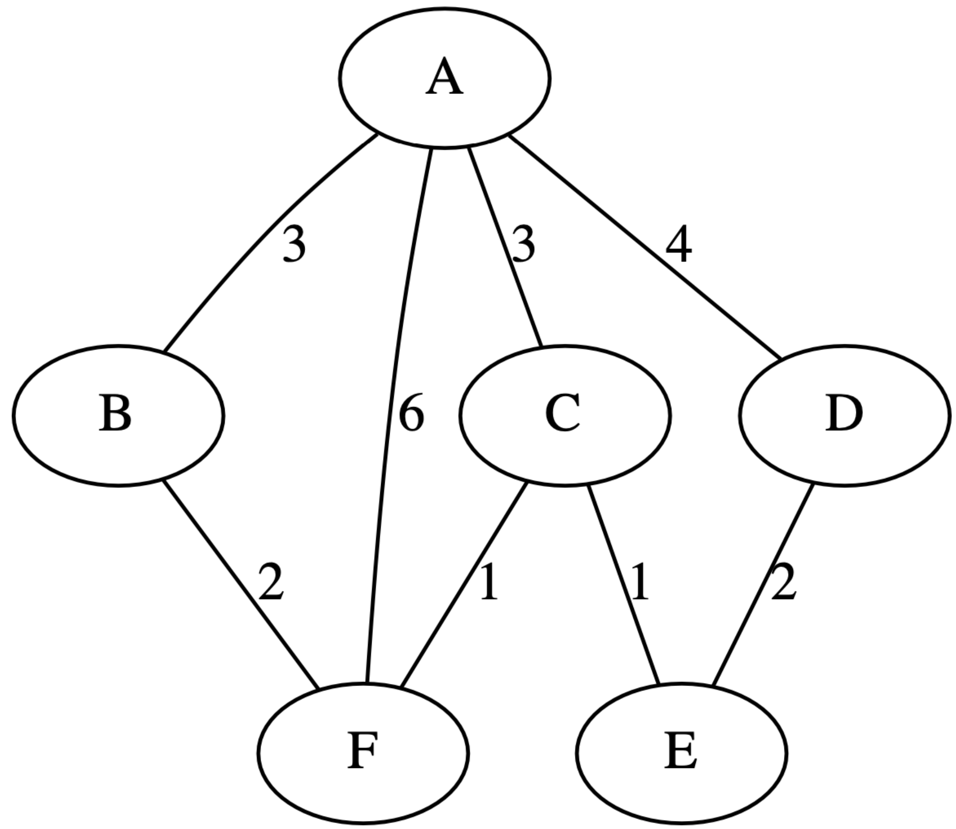

DIJKSTRA'S: ONE STEP – X = C

Done: A B C

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	-	-
F	4	C

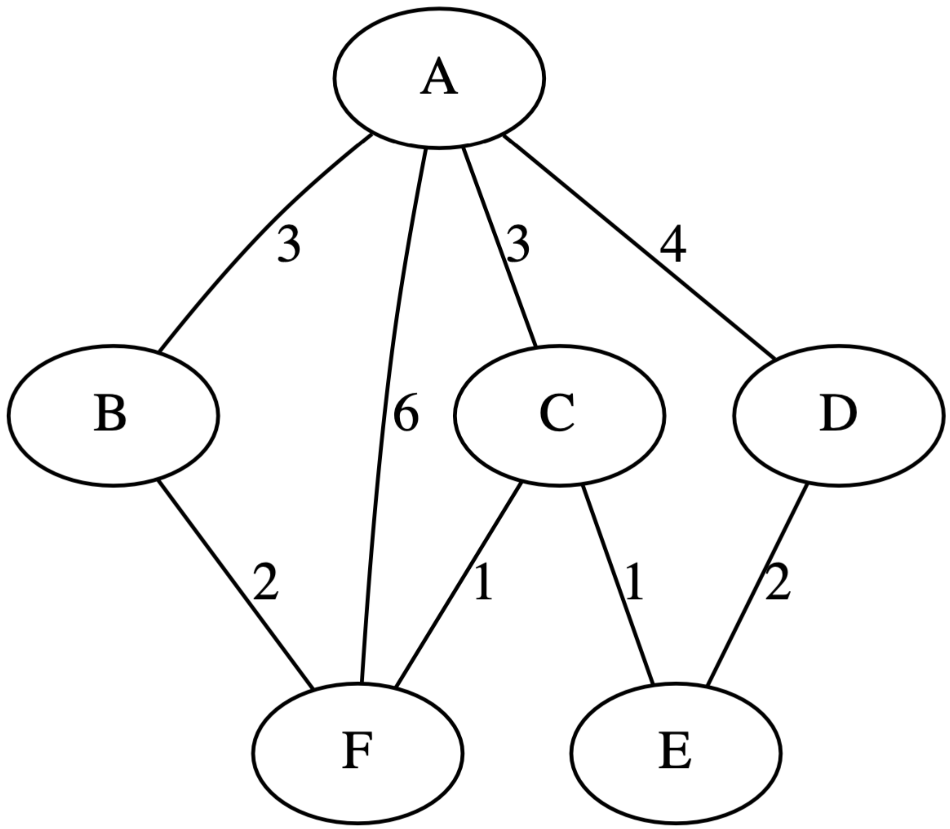

Dijkstra's: One Step – X = C

Done: A B C

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

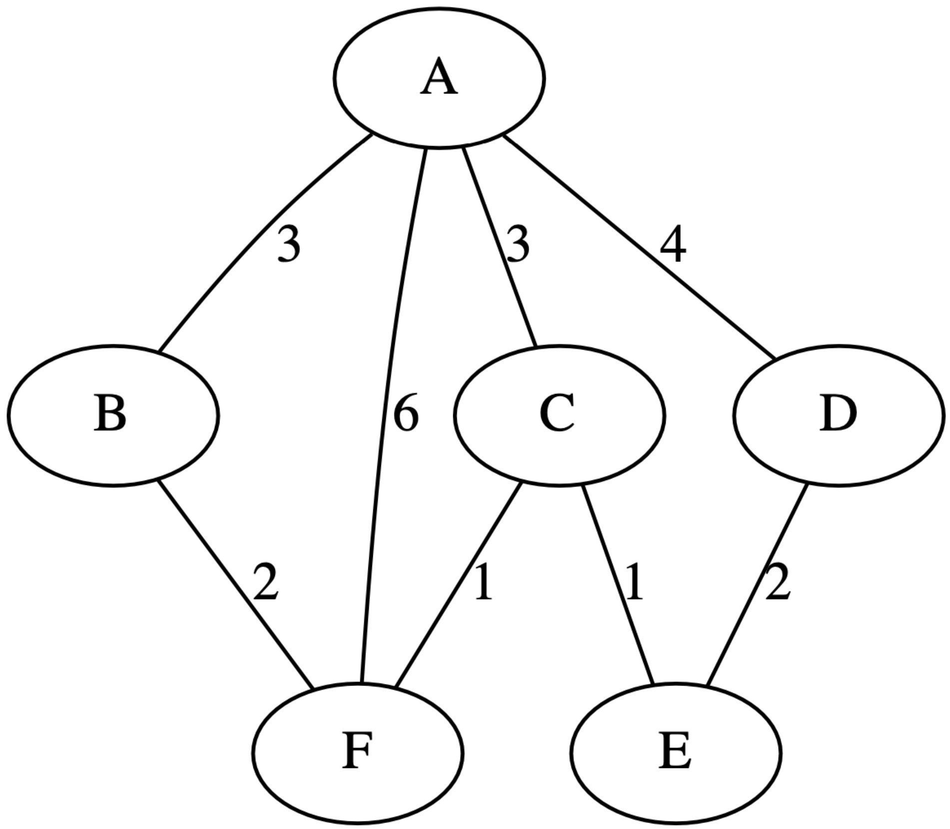

DIJKSTRA'S: ONE STEP – X = F

Done: A B C F

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

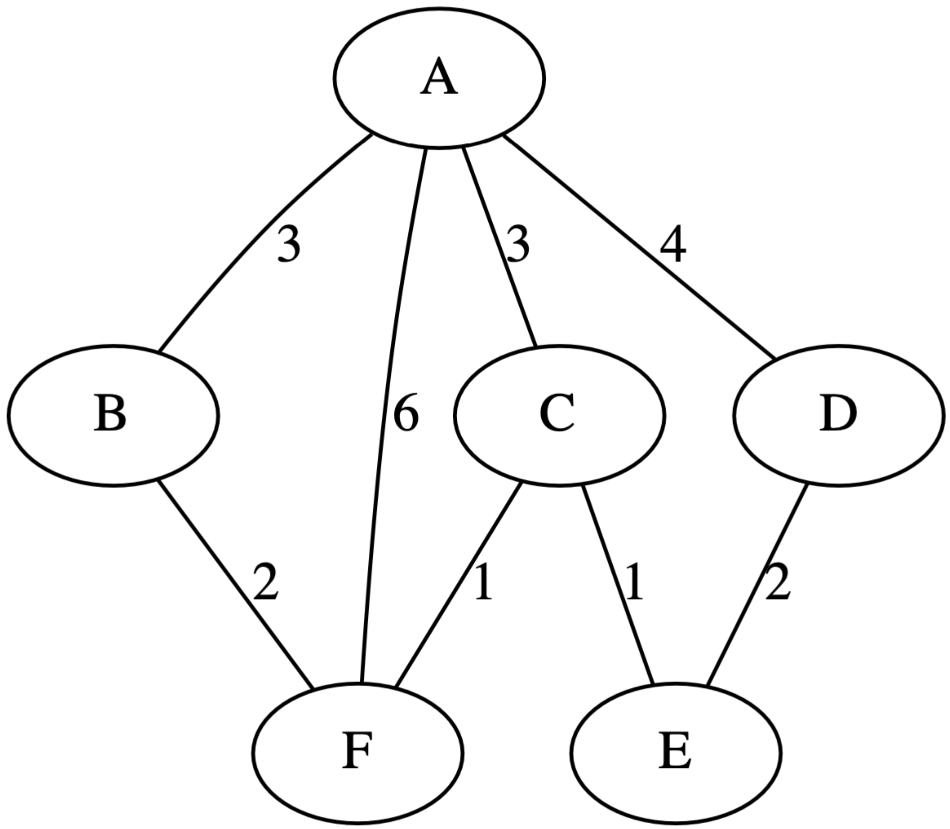

DIJKSTRA'S: ONE STEP – X = D

Done: A B C F D

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

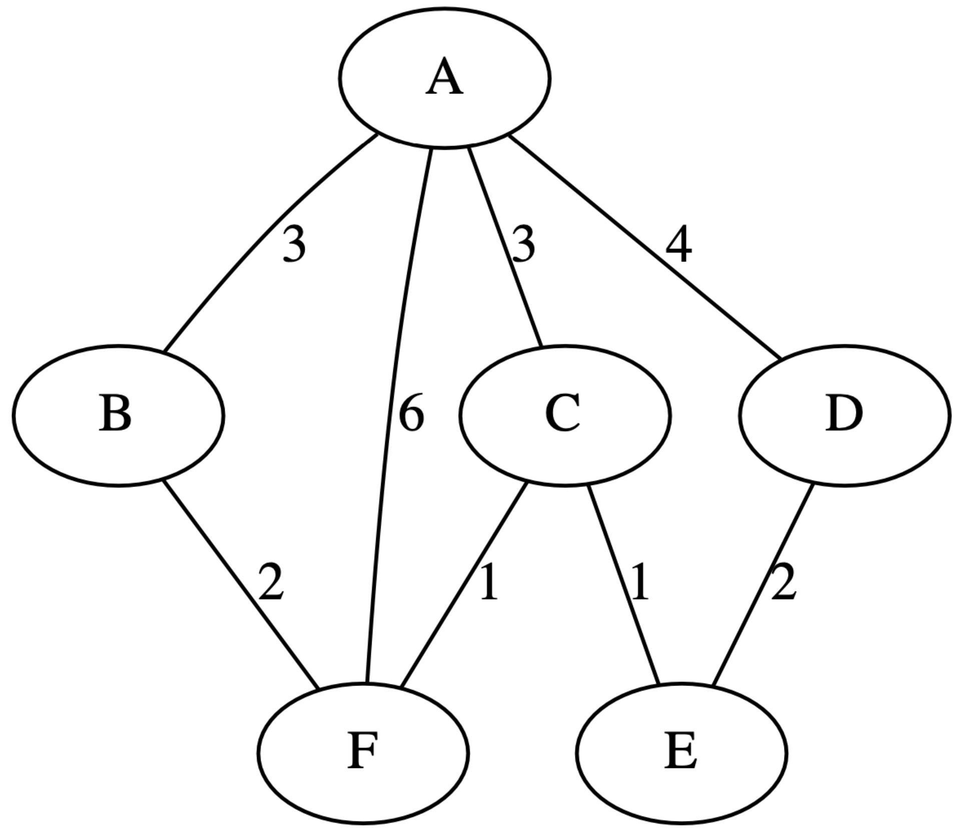

DIJKSTRA'S: ONE STEP – X = E

Done: A B C F D E

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

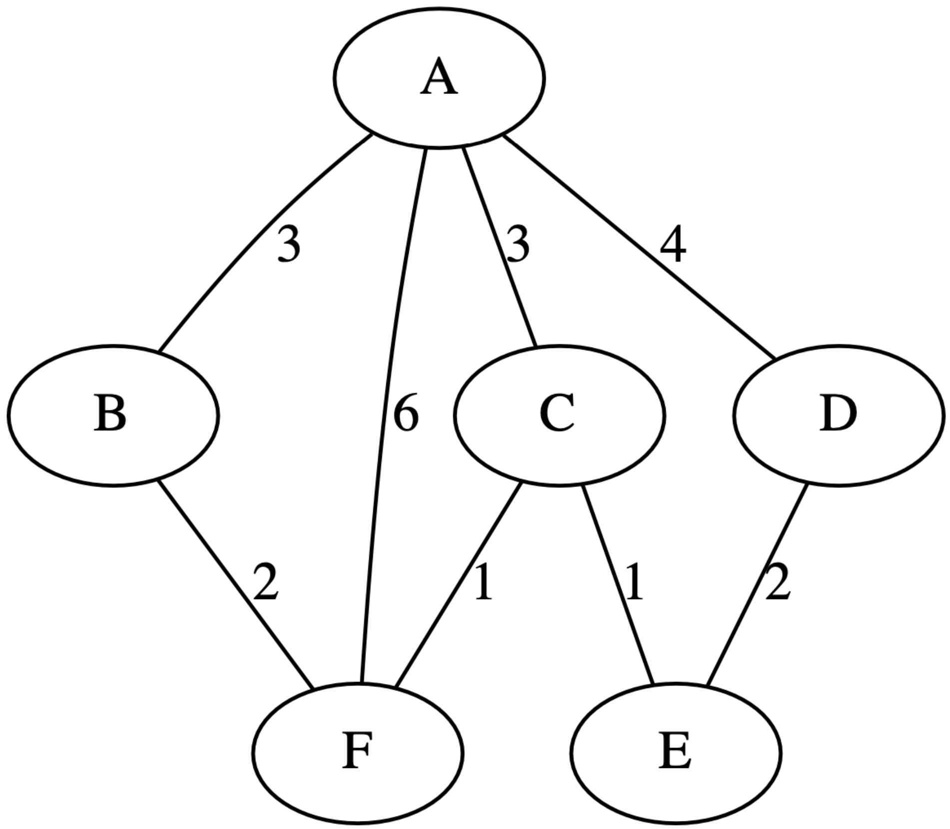

DIJKSTRA'S: DONE

Done: A B C F D E

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

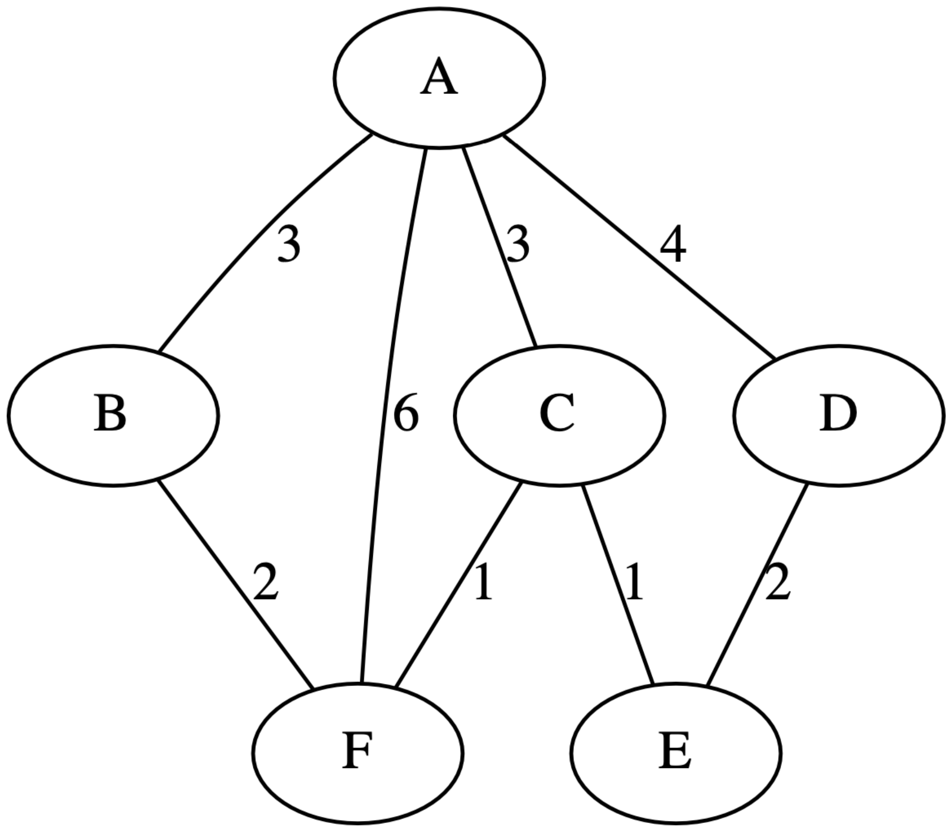

LINK STATE – UPDATE FORWARDING TABLE

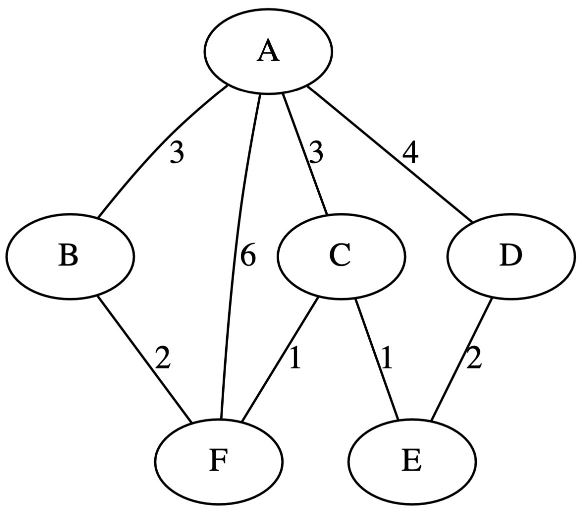

Dest	Cost	Prev
B	3	A
C	3	A
D	4	A
E	4	C
F	4	C

Dest	Link
B	L_B
C	L_C
D	L_D
E	L_C
F	L_C

A BIT MORE REALITY

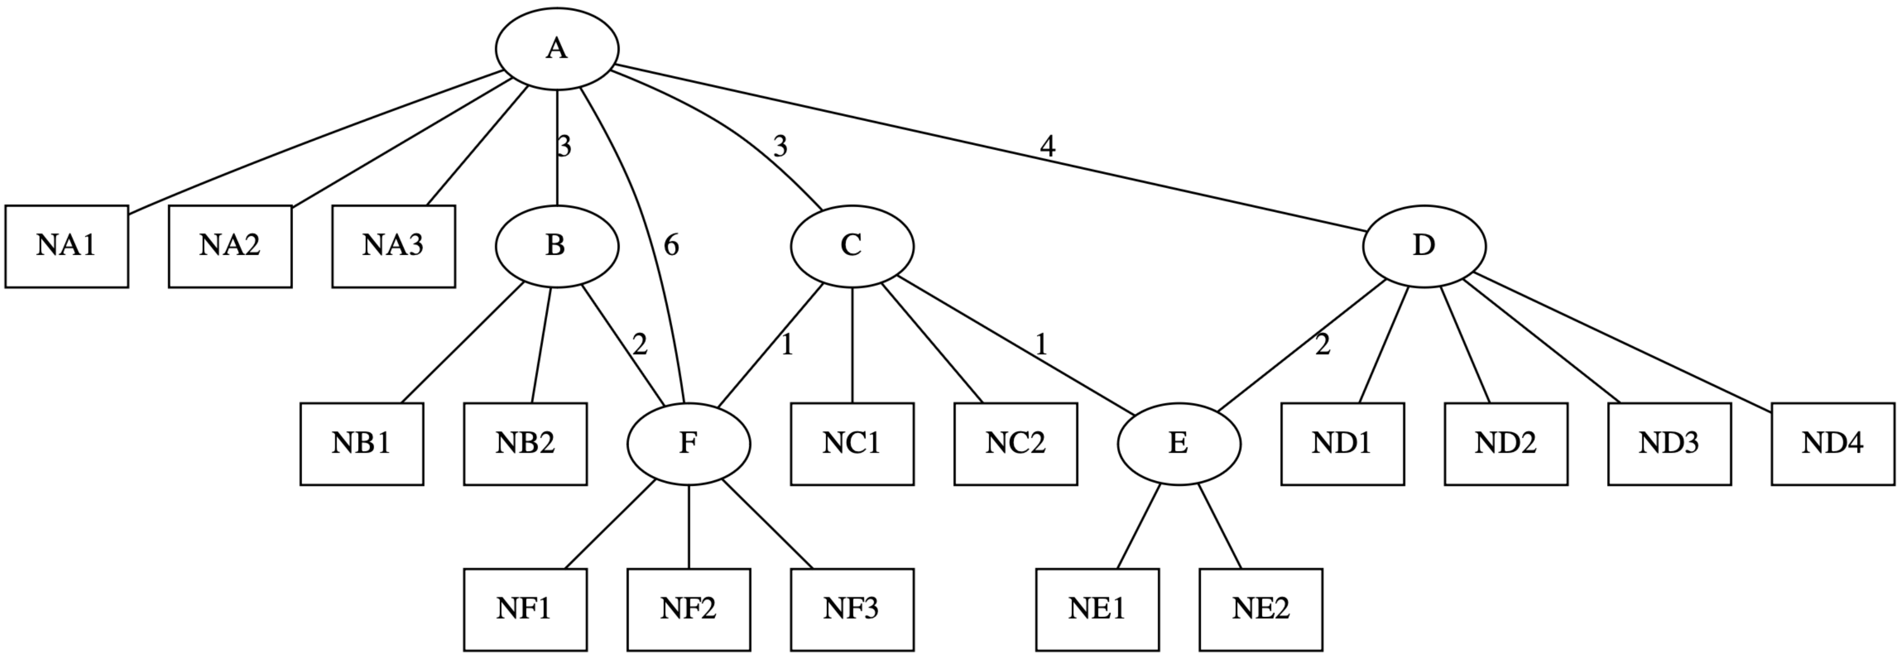

A'S FORWARDING TABLE

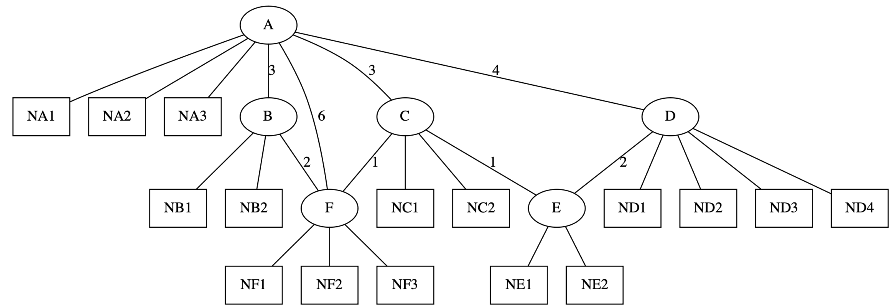

Dest	Link
NA1	L_{A1}
NA2	L_{A2}
NA3	L_{A3}
NB1	L_B
NB2	L_B
NC1	L_C

Dest	Link
NC2	L_C
ND1	L_D
ND2	L_D
ND3	L_D
ND4	L_D
NE1	L_C

Dest	Link
NE2	L_C
NF1	L_C
NF2	L_C
NF3	L_C

CLICKER QUESTION

How expensive is using a link-state routing algorithm?

N is the number of routers in the network

L is the number of links in the network

A) O(N + L)
B) O(N * L)
C) O(N ^ 2)
D) O(L ^ 2)
E) Something else