

Distance vector protocols

Distance Vector Routing Principles
Routing loops and countermeasures to loops
Bellman-Ford route calculations
RIP

Distance Vector Routing Principles

RIP - Routing Information Protocol is a basic protocol for interior routing


- RIP is a distance vector protocol
 - Based on the Bellman-Ford algorithm
- Routing table contains information about other known nodes
 - distance
 - link
- The nodes periodically send distance vectors based on the routing tables on all their links
- The nodes update their routing table with received distance vectors

S-38.121 S-02 RKa, NB

RIP-3

Let us study DV protocol principles

Example network with nodes A, B, C, D, E and links 1, 2, 3, 4, 5, 6.



Initial state: Nodes know their own addresses and interfaces, nothing more.

Node A creates its routing table:


From A to node	Link	Cost
A	local	0

The corresponding DV is: A=0

S-38.121 S-02 RKa, NB

RIP-4

Generation of routing tables starts when all routers send their DVs on all interfaces



Let's look at reception in Node B:

From node B to	Link	Cost
B	local	0


1. B receives the DV A=0
2. B increments the DV + 1 => A=1 and
3. B looks for the result in its routing table, no match
4. B adds the result to its RT, result is

From node B to	Link	Cost
B	local	0
A	1	1

B creates its own DV and sends it to all neighbors

A to	Link	Cost
A	-	0
B	1	1

$A=2 > A=0$




C to	Link	Cost
C	-	0
B	2	1
A	2	2

E to	Link	Cost
E	-	0
B	4	1
A	4	2

D sends its distance vector to all neighbors

A to	Link	Cost
A	-	0
B	1	1
D	3	1



D=0, A=1

C to	Link	Cost
C	-	0
B	2	1
A	2	2


E to	Link	Cost
E	-	0
B	4	1
A	4	2
D	6	1

A=2 == A=2 ⇒ no change

Nodes with changed RTs create DVs and send them to neighbors


B to	Link	Cost
B	-	0
A	1	1
D	1	2
C	2	1
E	4	1

D to	Link	Cost
D	-	0
A	3	1
B	3	2
E	6	1



A=0, B=1, D=1

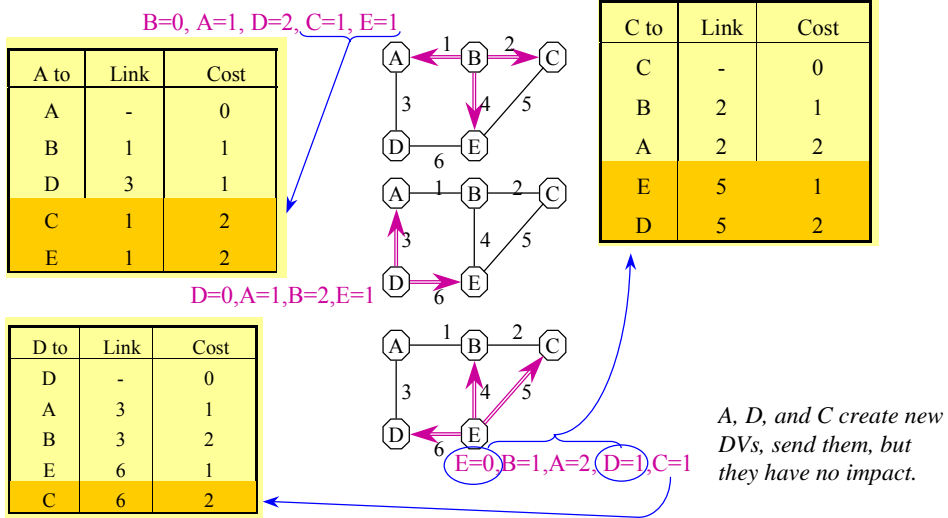
C=0 B=1, A=2



E=0 B=1, A=2, D=1

E to	Link	Cost
E	-	0
B	4	1
A	4	2
D	6	1
C	5	1

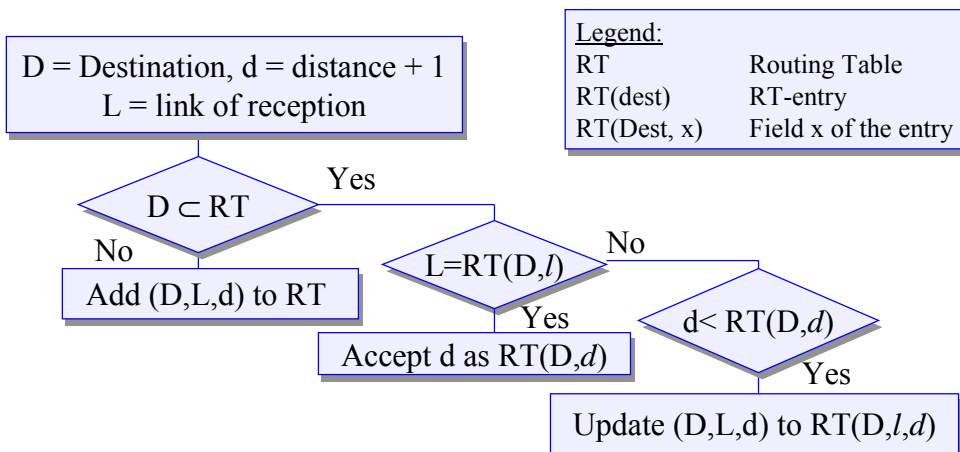
Again the changes are sent ...



S-38.121 S-02 RKa, NB

RIP-9

Processing of Received Distance Vectors



Note: this is simplified, shows only the principle!

S-38.121 S-02 RKa, NB

RIP-10


A link breaks...

A round of updates starts on link failure

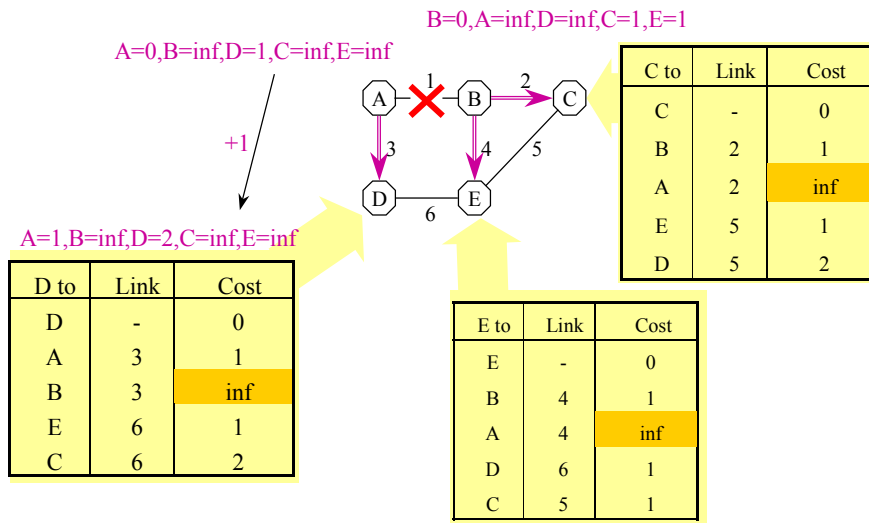
A to	Link	Cost
A	-	0
B	1	inf.
D	3	1
C	1	inf.
E	1	inf.

A gives an infinite distance to the nodes reached through link 1

B to	Link	Cost
B	-	0
A	1	inf
D	1	inf
C	2	1
E	4	1




D, E ja C update their Routing Tables



S-38.121 S-02 RKa, NB

RIP-13

D, C, E generate Distance Vectors...



S-38.121 S-02 RKa, NB


RIP-14

A, B, D, E generate their Distance Vectors

A to	Link	Cost
A	-	0
B	3	3
D	3	1
C	3	3
E	3	2

B to	Link	Cost
B	-	0
A	4	3
D	4	2
C	2	1
E	4	1


C to	Link	Cost
C	-	0
B	2	1
A	5	3
E	5	1
D	5	2



The result is that all nodes are able to communicate with all other nodes.

Routing loops

DV -protocol may create a transient routing loop




Let's assume that cost of loop 5 is 8.
A stable initial state for routes to C would be:

x to C	Link from x	Cost
A->C	1	2
B->C	2	1
C->C	-	0
D->C	3	3
E->C	4	2

Let's just look at the first link of each route.

Link 2 fails



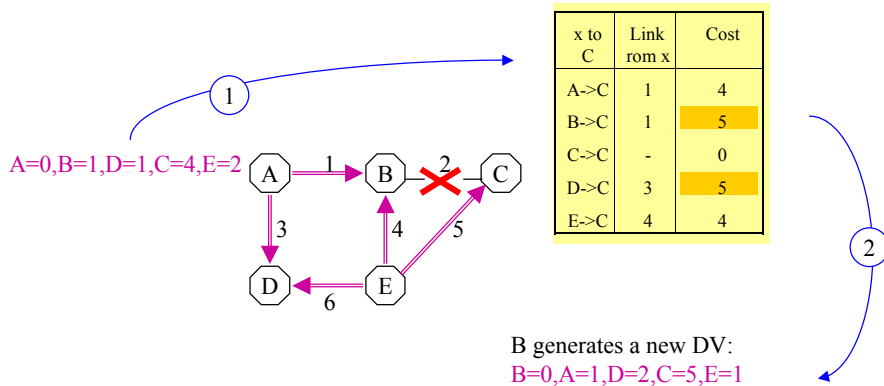
All packets to C are sent to B.
B sends them to A. A sends them back to B... until TTL=0.
(Bouncing effect)

x to C	Link from x	Cost
A->C	1	2
B->C	2	inf
C->C	-	0
D->C	3	3
E->C	4	2

Intermediate state

x to C	Link from x	Cost
A->C	1	4
B->C	1	3
C->C	-	0
D->C	3	3
E->C	4	4

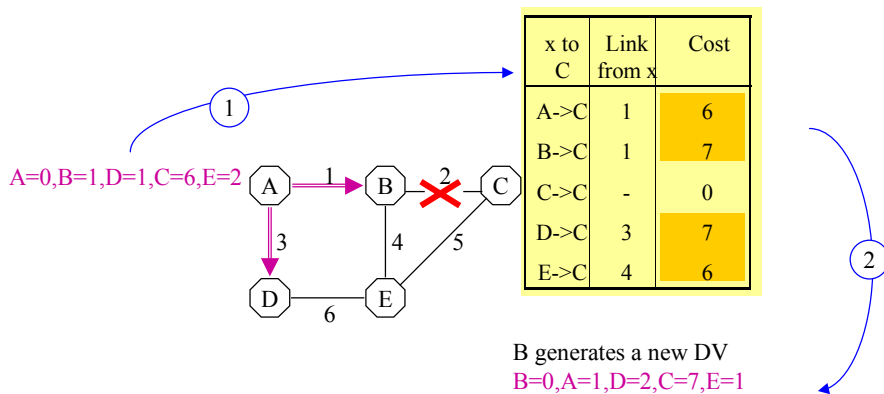
A and E send their Distance Vectors



⇒ Distance seen by A to C grows to 6

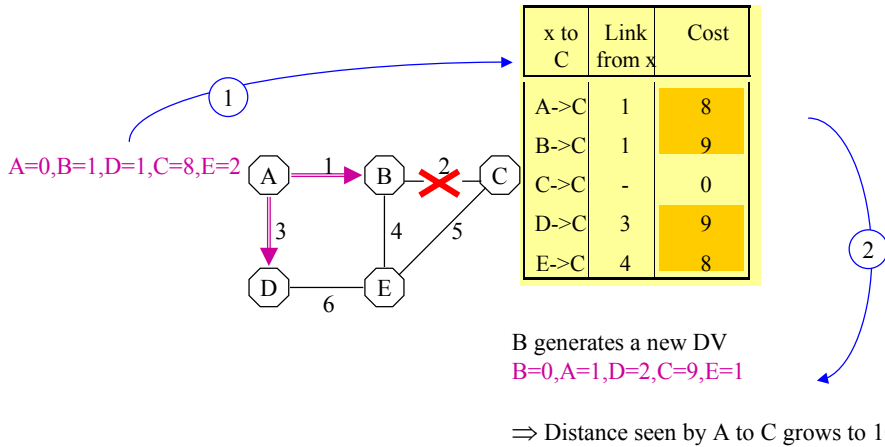
Distance vectors sent by C do not change anything because of high link cost

A sends a new DV




⇒ Distance seen by A to C grows to 8


A sends a new DV



A sends a new DV




E sends a new DV



x to C	Link from x	Cost
A->C	1	10
B->C	4	9
C->C	-	0
D->C	6	9
E->C	5	8

B send its DV but the Tables are already OK

x to C	Link from x	Cost
A->C	1	10
B->C	4	9
C->C	-	0
D->C	6	9
E->C	5	8




- Each update round improved the costs by 2
- The Process progresses in a random order, because it is genuinely parallel in nature.
- During the process, the state of the network is bad. DV-packets may be lost due to the overload created by bouncing user messages

Counting to Infinity occurs when failures break the network to isolated islands (1)

- Linkki 1 is broken, and the network has recovered.
- All link costs = 1

A to	Link	Cost
D	3	1
A	-	0
B	3	3
E	3	2
C	3	3




D to	Link	Cost
D	-	0
A	3	1
B	6	2
E	6	1
C	6	2

Counting to Infinity occurs when failures break the network to isolated islands (2)

- Also link 6 breaks.
- D has not yet sent its distance vector.

A to	Link	Cost
D	3	1
A	-	0
B	3	3
E	3	2
C	3	3



D to	Link	Cost
D	-	0
A	3	1
B	6	inf
E	6	inf
C	6	inf


Counting to Infinity occurs when failures break the network to isolated islands (3)

- A sends its distance vector first.

A=0,B=3,D=1,C=3,E=2

- D adds the information sent by A into its routing table.

A to	Link	Cost
D	3	1
A	-	0
B	3	3
E	3	2
C	3	3



D to	Link	Cost
D	-	0
A	3	1
B	3	4
E	3	3
C	3	4


S-38.121 S-02 RKa, NB

RIP-27

Counting to Infinity occurs when failures break the network to isolated islands (4)

- The result is a loop. Costs are incremented by 2 on each round.
- An agreement is needed: Cost greater than any route cost is = inf.

A to	Link	Cost
D	3	1
A	-	0
B	3	5
E	3	4
C	3	5



D to	Link	Cost
D	-	0
A	3	1
B	3	4
E	3	3
C	3	4

S-38.121 S-02 RKa, NB

RIP-28

Loops can be often avoided if less info is sent and by generating DVs immediately on RT change

Split horizon rule:

If node A sends to node X thru node B, it does not make sense for B to try to reach X thru A


⇒ A should not advertise to B its short distance to X

Implementation choices:

1. A does not advertise its distance to X towards B at all
⇒ the loop of previous example can not occur
2. A advertises to B: $X = \text{inf.}$ (“split horizon with poisonous reverse”)
⇒ two node loops are killed

Three node loops are still possible (1)


- Linkki 1 is broken, and the network has recovered.
- All link costs = 1



x to D	Link from x	Cost
B→D	4	2
C→D	5	2
E→D	6	1

Three node loops are still possible (2)

- Also link 6 fails.
- E sends its distance vector to B and C
 $E=0, B=1, A=\text{inf}, D=\text{inf}, C=1$




x to D	Link	Cost
B→D	4	2
C→D	5	2
E→D	6	inf

S-38.121 S-02 RKa, NB

RIP-31

Three node loops are still possible (3)

- Also link 6 fails.
- E sends its distance vector to B and C
 $E=0, B=1, A=\text{inf}, D=\text{inf}, C=1$
- ... But the DV sent to C is lost



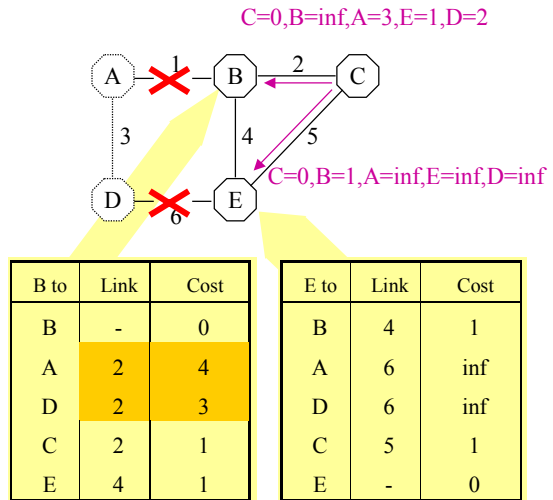
x to D	Link from x	Cost
B→D	4	inf
C→D	5	2
E→D	6	inf

S-38.121 S-02 RKa, NB

RIP-32

Three node loops are still possible (4)

- Now C sends its poisoned DV




S-38.121 S-02 RKa, NB

RIP-33

Three node loops are still possible (5)

- B generates its poisoned distance vectors
- The three node loop is ready
- Routes to D do not change except that the costs keep growing, nodes count to infinity. This finally breaks the loop: on link 5 cost=4 is advertised. C's knowledge about the distance to D grows ...



x to D	Link from x	Cost
B→D	2	3
C→D	5	2
E→D	4	4

S-38.121 S-02 RKa, NB

RIP-34

When should a DV-protocol advertise

Time of advertisement is a compromise:

- + immediate delivery of change info
- + recovery from packet loss
- + need to monitor the neighbors
- sending all changes at the same time
- traffic load created by the protocol

+ = Faster
- = Slower

Event triggered updates improve the functioning of RIP

- Entries in the Routing Tables have refresh and obsolescence timeouts.
- RIP advertises when the refresh timer expires and when a change occurs in an entry.
- Triggered updates speed up counting to infinity and reduce the probability of loops

The Bellman-Ford algorithm

Bellman-Ford algorithm (1)

- DV-protocols are based on the Bellman-Ford algorithm
- Centralized version:
 1. Let N be the number of nodes and M the number of links.
 2. L is the link table with M rows, $L[l].m$ - link cost,
 $L[l].s$ - link source
 $L[l].d$ - link sink
 3. D is $N \times N$ matrix, such that $D[i,j]$ is the distance from i to j
 4. H on $N \times N$ matrix, such that $H[i,j]$ is the link i uses to send to j

D

1	1	..	i	..	N
⋮					
j		distance			
⋮		from i			
N		to j			

Both directions are presented separately in the Link table!
A Column \equiv DV of the corresponding node

Bellman-Ford algorithm (2)

- Initialized Distance and Link matrices

$$D = \begin{matrix} & 1 & .. & & .. & N \\ \begin{matrix} 1 \\ : \\ : \\ : \\ N \end{matrix} & \left| \begin{array}{ccccc} \mathbf{0} & \infty & \infty & \infty & \infty \\ \infty & \mathbf{0} & \infty & \infty & \infty \\ \infty & \infty & \mathbf{0} & \infty & \infty \\ \infty & \infty & \infty & \mathbf{0} & \infty \\ \infty & \infty & \infty & \infty & \mathbf{0} \end{array} \right. \end{matrix}$$

Distance matrix D

$$H = \begin{matrix} & 1 & .. & & .. & N \\ \begin{matrix} 1 \\ : \\ : \\ : \\ N \end{matrix} & \left| \begin{array}{ccccc} -1 & -1 & -1 & -1 & -1 \\ -1 & \mathbf{-1} & -1 & -1 & -1 \\ -1 & -1 & \mathbf{-1} & -1 & -1 \\ -1 & -1 & -1 & \mathbf{-1} & -1 \\ -1 & -1 & -1 & -1 & \mathbf{-1} \end{array} \right. \end{matrix}$$

Link matrix H

NB: Link vector has both directions of a link separately.
 First in D-matrix appear one hop link distances, second two hop link distances etc.

Bellman-Ford algorithm (3)

1. Initialization: If $i=j$, then $D[i,j] = 0$, else $D[i,j] = \text{inf}$.
 Initialize $\forall H[i,j] = -1$. (previous slide)
2. $\forall l$ and \forall destinations k set $i = L[l].s$, $j = L[l].d$ and calculate $d = L[l].m + D[j,k]$
3. If $d < D[i,k]$, set $D[i,k] = d$; $H[i,k] = l$.
4. If at least one $D[i,k]$ changed, GOTO 2, else END.

Bellman-Ford algorithm (4)

- Number of steps $\leq N$
- Complexity $O(M.N^2)$

RIP protocol

RIP-protocol properties (1)

- Simple protocol. Used before standardization.
- RIP version 1 – RFC 1058.
- RIP is used inside an autonomous system
- RIP works both on shared media (Ethernet) and in point-to-point networks.
- RIP runs on top of UDP and IP.

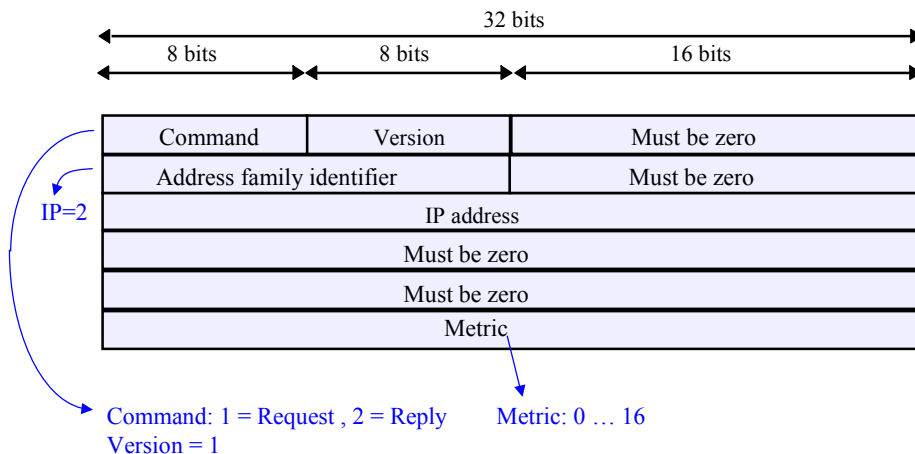
RIP-protocol properties (2)

- An Entry in the Routing Table represents a host, a network or a sub-net
 - <netid,subnetid,host> represents a host
 - <netid,subnetid,0> represents a sub-net
 - <netid,0,0> represents a network
 - <0.0.0.0> represents a route from the Autonomous System
- Information sent to the neighboring subnet is aggregated.

RIP-protocol properties (3)

- Distance = hop count = Nrof links on a path (route),
 - No other metrics
- RIP advertises once in 30s
 - if an entry is 180s old \Rightarrow distance is set to inf
- Distance 16 = infinite
- Timer triggered advertisements must be randomized to avoid bursts of RIP updates. 1-5 s.
- RIP uses poisoned vectors

RIP message format



RIP Routing Table

A routing table entry contains

- Destination IP address
- Distance to destination
- Next hop IP address
- “Recently” updated flag
- Several timers (refresh, obsolescence...)

RIP reply messages

- Distance vectors are sent in reply messages
- 30 seconds period
 - All routing table entries
 - Different DV on different links because of poisoned vectors
- Update messages after changes
 - Changed entries
 - 1-5 seconds delay, so that the message contains all updates that are related to the same change
- Destinations with infinite distance can be omitted if the next hop is same as before.

RIP request messages

- The router can request routing tables from its neighbors at startup
 - Complete list
- Partial routing table
 - For debugging

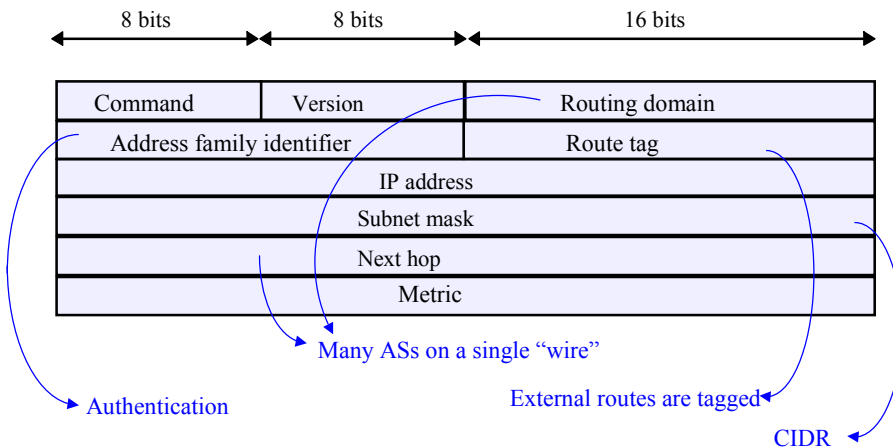
Silent nodes

- When only RIP was used, host could listen to routing traffic and maintain their own routing tables
 - Which router is closest to the destination?
 - Which link, if several available?
- These were "silent nodes", that only listened to routing traffic without sending
- Nowadays there are too many routing protocols
 - RIP-2, OSPF, IGRP, ...

RIP versio 2

- RFC-1388 (1387,1389)
- Why?
 - Simple and lightweight alternative to OSPF and IS-IS
- RIP-2 is a partially interoperable update with v1
 - RIP-1 router understand some of what a RIP-2 router is saying.
- Improvements
 - Authentication
 - Next hop –field
 - Subnet mask
 - External routes
 - Updates with multicast

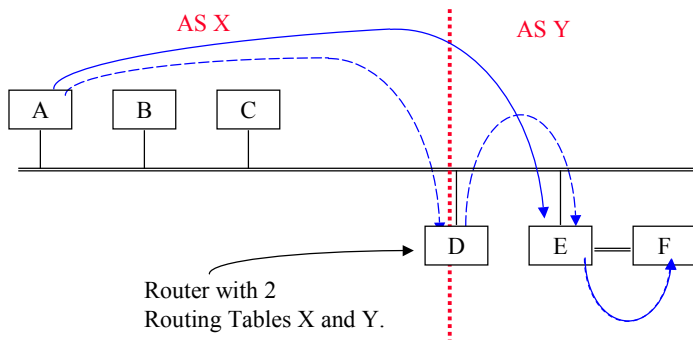
RIP versio 2 - sanomat



Routing from one sub-net to another

- In RIP-1 sub-net mask is not known outside the sub-net, only netid is sent in an advertisement out from a sub-net
 - ⇒ A host and a sub-net can not be distinguished
 - ⇒ All sub-nets must be interconnected with all other sub-nets and exterior traffic is received in the nearest router independent of the final destination inside our AS
- RIP-2 corrects the situation by advertising both the sub-net and the sub-net mask

Routing Domain and Next hop



Next hop ==> D advertises in X: the distance to F is f and the next hop is E!

Observations about RIP

- Routers have a spontaneous tendency to synchronize their send times. This increases the probability of losses in the net. Therefore, send instants are randomized between 15s ... 45s.
 - Reason: send interval = constant + time of message packing + processing time of messages that are in the queue.
- When RIP is used on ISDN links
 - A new call is established/30s \Rightarrow expensive.
- Slow network \Rightarrow queue length are restricted. RIP sends its DVs 25 entries/message in a row \Rightarrow RIP messages may be lost.
- A Correction proposal: ack all DVs: no periodic updates
 - \Rightarrow If there are no RIP message: assume that neighbor is alive and reachable
 - \Rightarrow Info on all alternative routes is stored.