AALTO-YLIOPISTO
Tietoliikenne- ja tietoverkkotekniikan laitos
S-38.2121 Reititys tietoliikenneverkoissa
Nicklas Beijar

Model solutions and grading guidelines for the exam 13.12.2011

This document describes the guidelines for grading and provides model solutions indicating the main points that were expected from the answer. It should be seen as a guideline for what is expected from the answer. It is not a strict requirement list.

A good answer must clearly show that the subject is understood. The given points can vary from the guidelines if the answer is incomlete, insufficient or does not show an underlying understanding. Serious errors showing misunderstanding decrease the points, i.e. give negative points in a subquestion. However, a minor error in some detail does not give negative points. Extra information related to the question may give additional points.

For each question, a grade between 0 and 6 is given. A plus sign "+" denotes 0.25 points while a minus sign "-" denotes -0.25 points.

Question 1

Selitä puhelinkeskuksen numeroanalyysin periaate.
Explain the principle of digit analysis in a telephone exchange.

Model solution and grading

Part of the number that is analyzed: The number is hierarchical. In originating and transit exchanges, only the leading digits are analyzed to find the route to the terminating exchange. The terminating exchange analyzes all digits to find the identity of the subscriber's physical interface. (1p) *An exceptionally thorough explanation (e.g. the relation to network hierarchy) can give 2 points. However, things generally considered known by most people (such as area codes) do not give points.*

Analysis implementation: The digit analysis is implemented using a analysis tree. (1p)
The final pointer in the tree points to a bucket in a bucket file. The bucket describes what to do with the call, e.g. a set of alternative paths. (1p) *Also a figure is accepted.*

Analysis input: The digit analysis is based on the dialed digits, the incoming circuit group, the origin, the subscriber category (e.g. operator), etc. (1p)

Analysis output: The digit analysis produces a set of alternative paths or a translated number. The analysis may also return additional information that may be needed in outgoing signaling for the call. (1p)

Number translation: Number portability and service numbers require a number translation. This can be implemented using a Service Control Point (SCP) or directly in the exchange. The analysis may then be repeated with the translated number as input. (1p)

Related slides

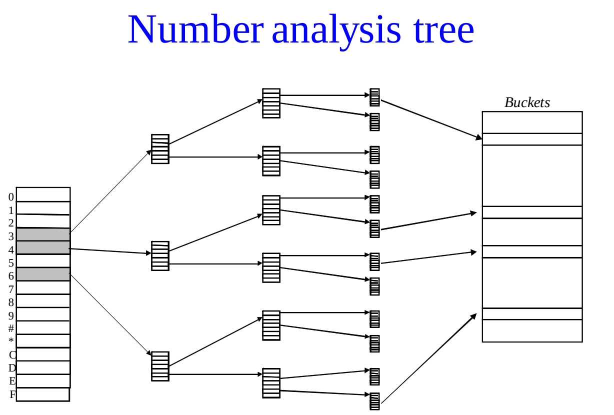

Number analysis tree

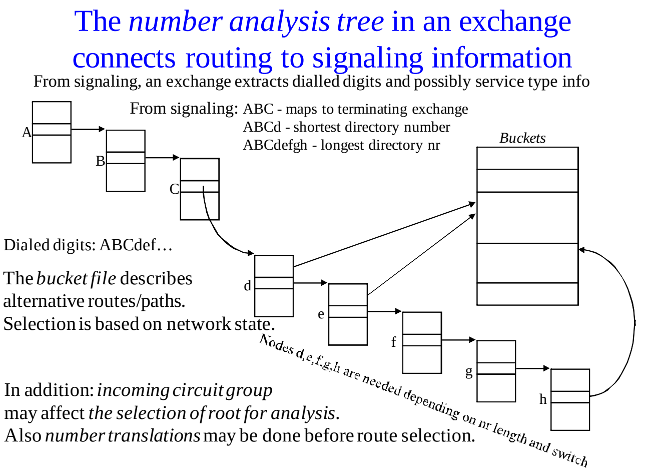

The *number analysis tree* in an exchange connects routing to signaling information

From signaling, an exchange extracts dialled digits and possibly service type info

From signaling: ABC - maps to terminating exchange
ABCd - shortest directory number
ABCdefgh - longest directory nr

Dialed digits: ABCdef…

The *bucket file* describes alternative routes/paths. Selection is based on network state.

In addition: *incoming circuit group* may affect *the selection of root for analysis*. Also *number translations* may be done before route selection.

Nodes d,e,f,g,h are needed depending on nr length and switch

A
B
C
d
e
f
g
h
Buckets

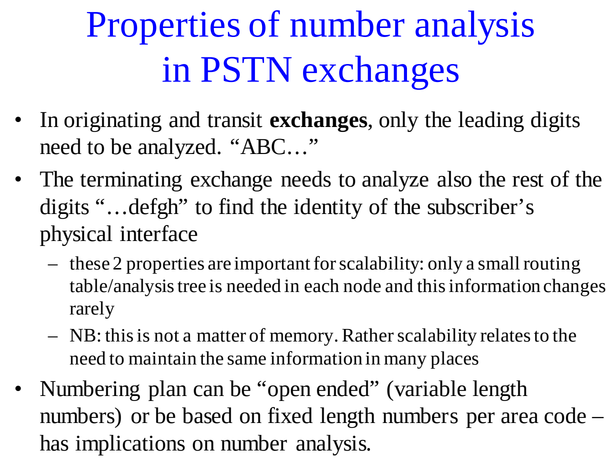

Properties of number analysis in PSTN exchanges

- In originating and transit **exchanges**, only the leading digits need to be analyzed. "ABC…"
- The terminating exchange needs to analyze also the rest of the digits "…defgh" to find the identity of the subscriber's physical interface
  - these 2 properties are important for scalability: only a small routing table/analysis tree is needed in each node and this information changes rarely
  - NB: this is not a matter of memory. Rather scalability relates to the need to maintain the same information in many places
- Numbering plan can be "open ended" (variable length numbers) or be based on fixed length numbers per area code – has implications on number analysis.

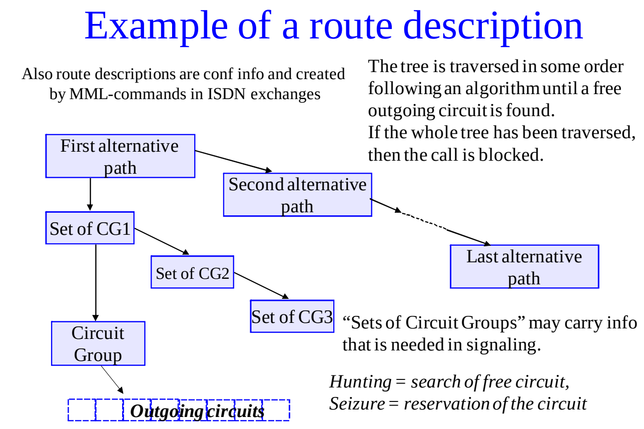

Example of a route description

Also route descriptions are conf info and created by MML-commands in ISDN exchanges

The tree is traversed in some order following an algorithm until a free outgoing circuit is found. If the whole tree has been traversed, then the call is blocked.

"Sets of Circuit Groups" may carry info that is needed in signaling.

Hunting = search of free circuit,
Seizure = reservation of the circuit

Grading comments

Seems to be more difficult this year than previous years. Most of the answers were at a very high (generic) level, not showing deeper understanding. Therefore, the scale was changed to give up to 2 point for the understanding the hierarchical numbering.

Question 2

Näytä, että silmukka on mahdollinen vaikka etäisyysvektoriprotokolla käyttää jaettua horisonttia ja myrkytettyjä vektoreita.

Show that a routing loop is possible even if the distance vector protocol uses split horizon and poisonous vectors.

Model solution and grading

Model solution: see example “Three-node loops are still possible” in slides.

2p for proving understanding what poisonous vectors are (e.g. by indicating them in the example) (e.g. telling that poisonous vectors only prevent 2-hop loops gives 1p).

1p for proving understanding of the cause of a loop (e.g. loss of packet)

2p for showing and example or describing text of a case where a loop appears even though poisonous vectors are used

1p for explanations of what happens in the given example (scaled according to the correctness of the example, e.g. a good explanation of an incomplete example worth 1p gives ½p for explanation)

Related slides

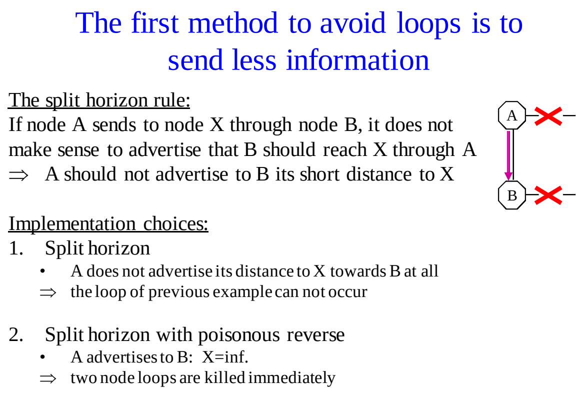

The first method to avoid loops is to send less information

The split horizon rule:
If node A sends to node X through node B, it does not make sense to advertise that B should reach X through A
⇒ A should not advertise to B its short distance to X

Implementation choices:

1. Split horizon
   - A does not advertise its distance to X towards B at all
   - ⇒ the loop of previous example can not occur
2. Split horizon with poisonous reverse
   - A advertises to B: X=inf.
   - ⇒ two node loops are killed immediately

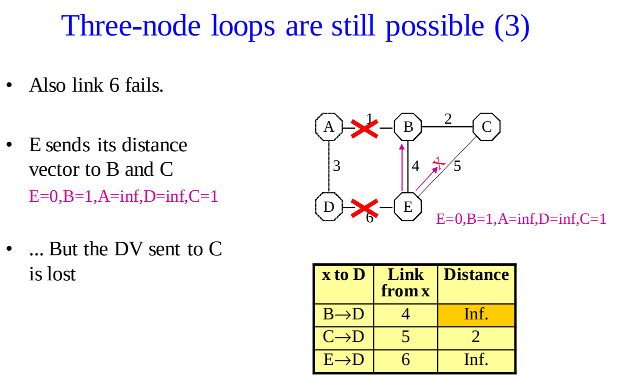

Three-node loops are still possible (3)

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

x to D	Link from x	Distance
B→D	4	Inf.
C→D	5	2
E→D	6	Inf.

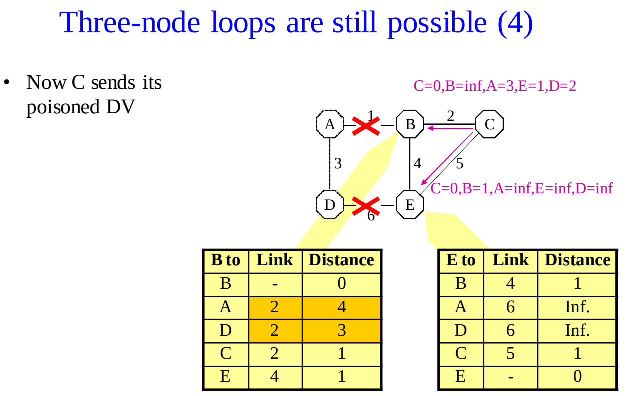

Three-node loops are still possible (4)

- Now C sends its poisoned DV

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

E to	Link	Distance
B	4	1
A	6	Inf.
D	6	Inf.
C	5	1
E	-	0

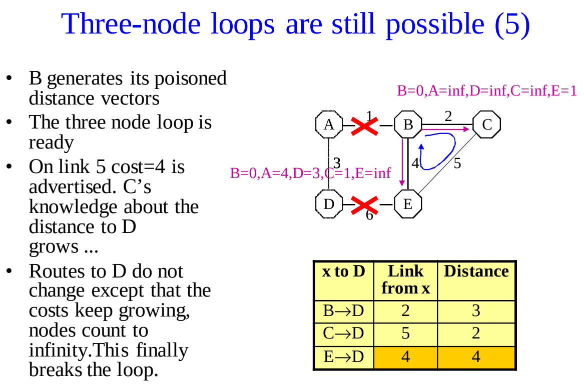

Three-node loops are still possible (5)

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

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

Question 3

Luettele OSPF:n osa-protokollat. Kuvaa lyhyesti jokaisen osa-protokollan tehtävät ja toimintaperiaatteet.

List the sub-protocols of OSPF. Describe shortly the tasks and operating principles of each sub-protocol.

Model solution and grading

For each sub-protocol 2p (½p for the name, 1p for the task, ½p for the operational principles)

Hello protocol (½p)

- Task: Checks if the link is working bidirectionally (½p), selects designated router and backup designated router (½p)
- Operational principle: Periodical sending of Hello messages on all links (½p)

Database Exchange protocol (½p)

- Task: Synchronizes link databases when a link starts working (1p)
- Operational principle: Each end of the link describes its link database with database description packets, the other side of the link requests differing and new records (½p)

Flooding protocol (½p)

- Task: Updating the local LSAs to all routers in the area (1p)
- Operational principle: Update messages are sent on all links (a) periodically and (b) when the topology has been modified. The Update messages are distributed with flooding (by repeating an unseen message on all interfaces except the incoming interface) so that it is received by all other routers in the area. (½p)

Related slides

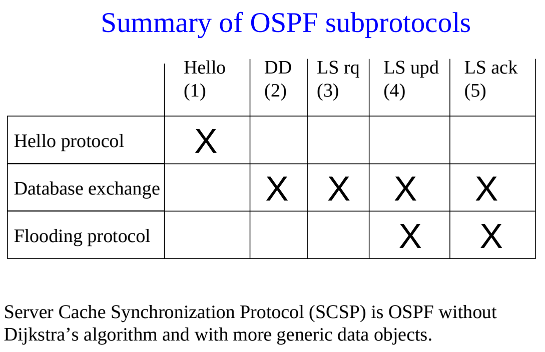

Summary of OSPF subprotocols

	Hello (1)	DD (2)	LS rq (3)	LS upd (4)	LS ack (5)
Hello protocol	X				
Database exchange		X	X	X	X
Flooding protocol				X	X

Server Cache Synchronization Protocol (SCSP) is OSPF without Dijkstra's algorithm and with more generic data objects.

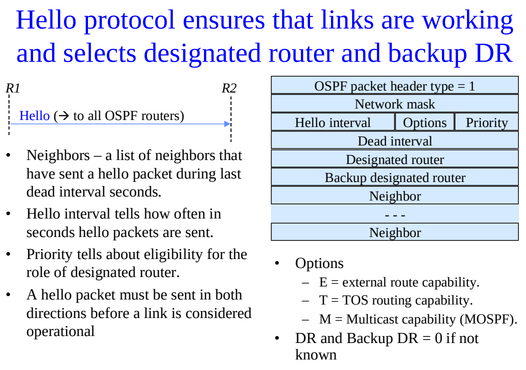

Hello protocol ensures that links are working and selects designated router and backup DR

- Neighbors – a list of neighbors that have sent a hello packet during last dead interval seconds.
- Hello interval tells how often in seconds hello packets are sent.
- Priority tells about eligibility for the role of designated router.
- A hello packet must be sent in both directions before a link is considered operational

- Options
  - E = external route capability.
  - T = TOS routing capability.
  - M = Multicast capability (MOSPF).
- DR and Backup DR = 0 if not known

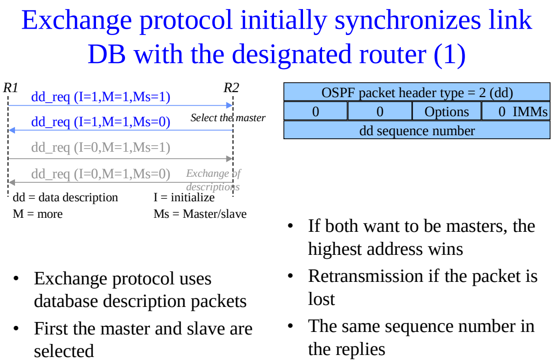

Exchange protocol initially synchronizes link DB with the designated router (1)

- Exchange protocol uses database description packets
- First the master and slave are selected
- If both want to be masters, the highest address wins
- Retransmission if the packet is lost
- The same sequence number in the replies

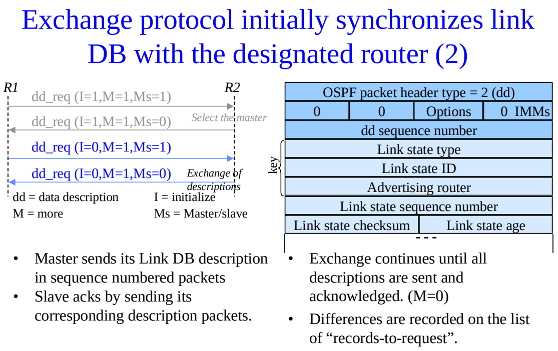

Exchange protocol initially synchronizes link DB with the designated router (2)

- Master sends its Link DB description in sequence numbered packets
- Slave acks by sending its corresponding description packets.
- Exchange continues until all descriptions are sent and acknowledged. (M=0)
- Differences are recorded on the list of “records-to-request”.

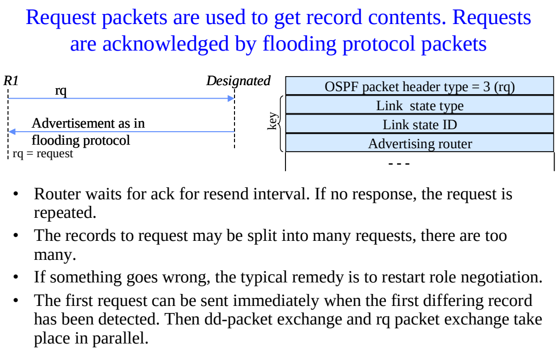

Request packets are used to get record contents. Requests are acknowledged by flooding protocol packets

- Router waits for ack for resend interval. If no response, the request is repeated.
- The records to request may be split into many requests, there are too many.
- If something goes wrong, the typical remedy is to restart role negotiation.
- The first request can be sent immediately when the first differing record has been detected. Then dd-packet exchange and rq packet exchange take place in parallel.

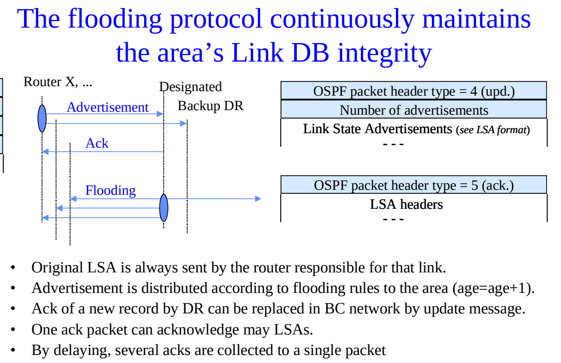

The flooding protocol continuously maintains the area’s Link DB integrity

Router X, ...
Designated
Backup DR
Advertisement
Ack
Flooding
OSPF packet header type = 4 (upd.)
Number of advertisements
Link State Advertisements (see LSA format)
- - -
OSPF packet header type = 5 (ack.)
LSA headers
- - -

- Original LSA is always sent by the router responsible for that link.
- Advertisement is distributed according to flooding rules to the area (age=age+1).
- Ack of a new record by DR can be replaced in BC network by update message.
- One ack packet can acknowledge may LSAs.
- By delaying, several acks are collected to a single packet

Question 4

Ethernet-verkko yhdistää viisi reititintä. Miten OSPF laskee reitittimien väliset kustannukset? Millaisia linkkitilatietueita muodostetaan?
An Ethernet network connects five routers. How does OSPF calculate the costs between the routers? What link state records are generated?

Model solution and grading

Calculating the routes (max 3p of the following)

- Ethernet is a broadcast network. In a broadcast network the five routers can send traffic directly to each other over a single hop. However, modeling all connections between the five routers would create 5(5-1) = 20 links. (1p for understanding the background)
- To reduce the number of links to model, OSPF calculates the routes using the concept of a virtual router in the network to which all five actual routers are connected (1p for describing the use of virtual routers)
- The cost between two routers A and B is calculated by summing the costs of the link from router A to the virtual router and from the virtual router to router B. (1p for describing how the costs are summed)
- The link from the virtual router to the actual router has a cost of 0 and the link from the actual router to the virtual router has the actual cost of the network (1p for describing the zero cost)

Generated link state records (max 3p)

- Each of the five routers generates a router LSA (1p for indicating the router LSA) containing a description of the link going from the router to the network. (½p for explaining the purpose of the router LSA)
- The designated router generates a network LSA (1p for indicating the network LSA) that lists the five routers connected to the network. (½p for explaining the purpose of the network LSA)

Related slides

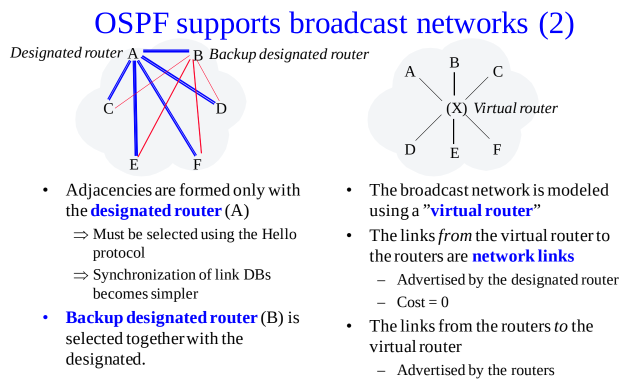

OSPF supports broadcast networks (2)

- Adjacencies are formed only with the **designated router** (A)
  - ⇒ Must be selected using the Hello protocol
  - ⇒ Synchronization of link DBs becomes simpler
- **Backup designated router** (B) is selected together with the designated.

- The broadcast network is modeled using a ”**virtual router**”
- The links *from* the virtual router to the routers are **network links**
  - Advertised by the designated router
  - Cost = 0
- The links from the routers *to* the virtual router
  - Advertised by the routers

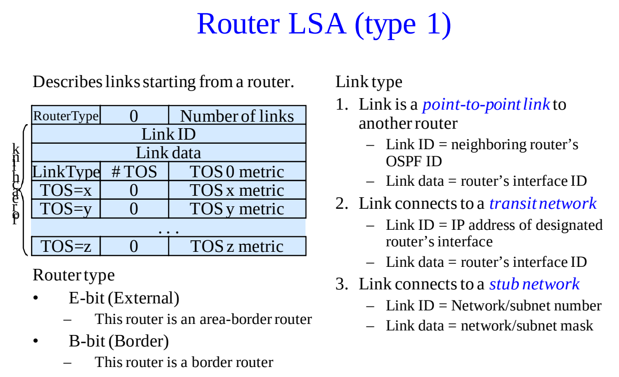

Router LSA (type 1)

Describes links starting from a router.

RouterType	0	Number of links
Link ID		
Link data		
LinkType	# TOS	TOS 0 metric
TOS=x	0	TOS x metric
TOS=y	0	TOS y metric
. . .		
TOS=z	0	TOS z metric

Router type

- E-bit (External)
  - This router is an area-border router
- B-bit (Border)
  - This router is a border router

Link type

1. Link is a *point-to-point link* to another router
   - Link ID = neighboring router’s OSPF ID
   - Link data = router’s interface ID
2. Link connects to a *transit network*
   - Link ID = IP address of designated router’s interface
   - Link data = router’s interface ID
3. Link connects to a *stub network*
   - Link ID = Network/subnet number
   - Link data = network/subnet mask

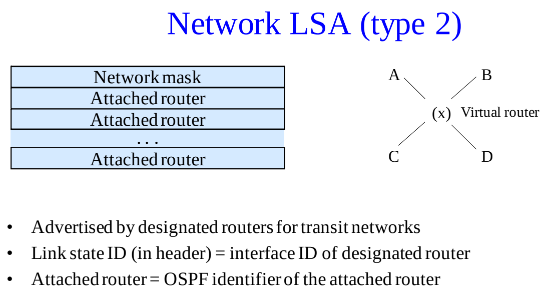

Network LSA (type 2)

Network mask
Attached router
Attached router
. . .
Attached router

- Advertised by designated routers for transit networks
- Link state ID (in header) = interface ID of designated router
- Attached router = OSPF identifier of the attached router

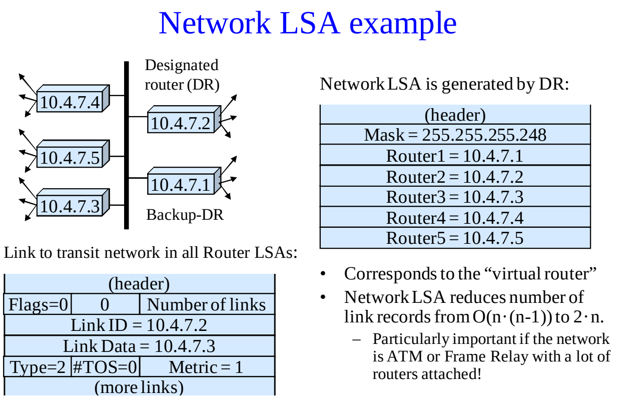

Network LSA example

Link to transit network in all Router LSAs:

(header)		
Flags=0	0	Number of links
Link ID = 10.4.7.2		
Link Data = 10.4.7.3		
Type=2	#TOS=0	Metric = 1
(more links)		

Network LSA is generated by DR:

(header)
Mask = 255.255.255.248
Router1 = 10.4.7.1
Router2 = 10.4.7.2
Router3 = 10.4.7.3
Router4 = 10.4.7.4
Router5 = 10.4.7.5

- Corresponds to the “virtual router”
- Network LSA reduces number of link records from O(n·(n-1)) to 2·n.
  - Particularly important if the network is ATM or Frame Relay with a lot of routers attached!

Question 5

Kuvaa AODV:n (Ad hoc On-demand Distance Vector) -reititysprotokollan toiminataperiaatteet ja järjestysnumeroon liittyvät toiminnot.
Describe the operating principles and the use of sequence numbers in the AODV (Ad hoc On-demand Distance Vector) routing protocol.

Model solution and grading

Operating principles (max 3p)

- Ad-hoc On-Demand Distance Vector (AODV) is a **reactive** routing protocol, i.e. the route is generated when it is needed. Routes are generated and maintained only between active senders and receivers. (1p for telling that AODV is reactive)
- To find the route, AODV **floods the network** with a route request packet, using an expanding ring search (1p for describing the route request).
- The reply can be generated by an intermediate node (if a satisfying route is available) or by the destination itself (if no intermediate node replied). The reply travels the reverse path to the requesting node. (1p for describing the route reply)

Use of sequence numbers (max 3p)

- The route request contains the minimum sequence number requred by the requesting node. (1p for describing the sequence number in route requests)
- An intermediate node may reply if it has a sequence number higher than the requested one. The sequence number in the reply is the sequence number known by the intermediate node. (1p for describing the sequence number when an intermediate node replies)
- If a route request containing the destination’s current sequence number reaches the destination , the current sequence number is incremented. The sequence number is thus generated by the destination node. (1p for describing the sequence number when the destination replies)
- Sequence numbers are used to prevent routing loops and avoid old and broken routes (½p for explaining the generic purpose of sequence numbers)

Related slides

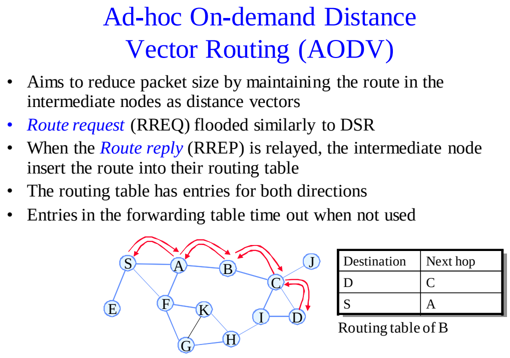

Ad-hoc On-demand Distance Vector Routing (AODV)

- Aims to reduce packet size by maintaining the route in the intermediate nodes as distance vectors
- *Route request* (RREQ) flooded similarly to DSR
- When the *Route reply* (RREP) is relayed, the intermediate node insert the route into their routing table
- The routing table has entries for both directions
- Entries in the forwarding table time out when not used

Destination	Next hop
D	C
S	A

Routing table of B

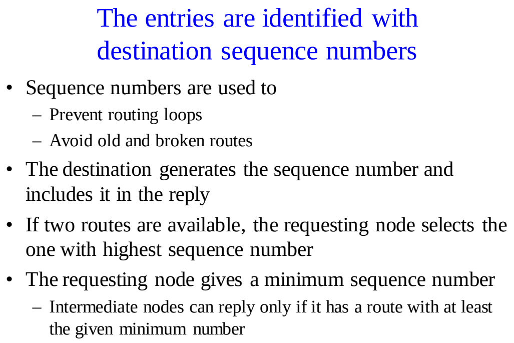

The entries are identified with destination sequence numbers

- Sequence numbers are used to
  - Prevent routing loops
  - Avoid old and broken routes
- The destination generates the sequence number and includes it in the reply
- If two routes are available, the requesting node selects the one with highest sequence number
- The requesting node gives a minimum sequence number
  - Intermediate nodes can reply only if it has a route with at least the given minimum number

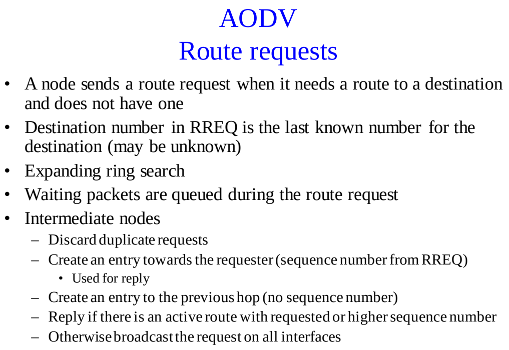

AODV Route requests

- A node sends a route request when it needs a route to a destination and does not have one
- Destination number in RREQ is the last known number for the destination (may be unknown)
- Expanding ring search
- Waiting packets are queued during the route request
- Intermediate nodes
  - Discard duplicate requests
  - Create an entry towards the requester (sequence number from RREQ)
    - Used for reply
  - Create an entry to the previous hop (no sequence number)
  - Reply if there is an active route with requested or higher sequence number
  - Otherwise broadcast the request on all interfaces

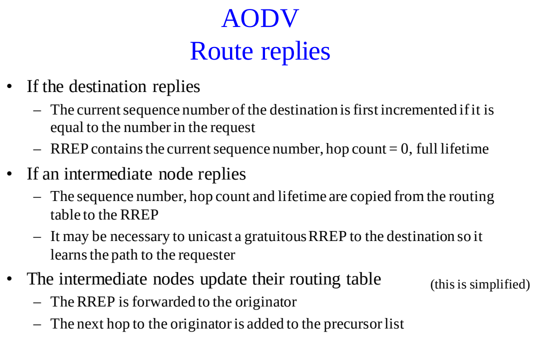

AODV Route replies

- If the destination replies
  - The current sequence number of the destination is first incremented if it is equal to the number in the request
  - RREP contains the current sequence number, hop count = 0, full lifetime
- If an intermediate node replies
  - The sequence number, hop count and lifetime are copied from the routing table to the RREP
  - It may be necessary to unicast a gratuitous RREP to the destination so it learns the path to the requester
- The intermediate nodes update their routing table (this is simplified)
  - The RREP is forwarded to the originator
  - The next hop to the originator is added to the precursor list

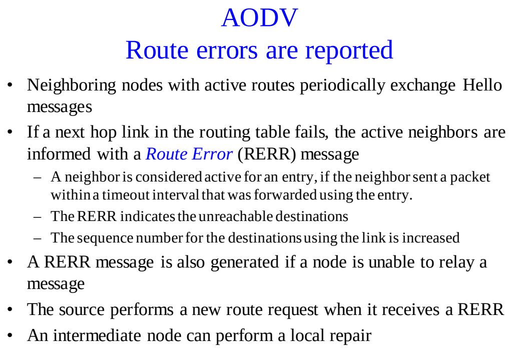

AODV
Route errors are reported

- Neighboring nodes with active routes periodically exchange Hello messages
- If a next hop link in the routing table fails, the active neighbors are informed with a *Route Error* (RERR) message
  - A neighbor is considered active for an entry, if the neighbor sent a packet within a timeout interval that was forwarded using the entry.
  - The RERR indicates the unreachable destinations
  - The sequence number for the destinations using the link is increased
- A RERR message is also generated if a node is unable to relay a message
- The source performs a new route request when it receives a RERR
- An intermediate node can perform a local repair

Question 6

Kuvan 1 verkossa reititin R3 toimii kohtaamispisteenä (rendezvous point) monilähetysprotokollan ollessa PIM-SM. Kuvaa mitä tapahtuu kun (a) solmut G1 ja G2 liittyvät monilähetysryhmän vastaanottajiksi, (b) solmu S lähettää paketteja kyseisen ryhmän vastaanottajille ja (c) otetaan lähettäjäkohtaiset puut käyttöön.

In the network of Figure 1 router R3 is the rendezvous point when PIM-SM is used for multicasting. Describe what happens when (a) nodes G1 and G2 join the multicast group as receivers, (b) node S sends packets to the group, and (c) source-based trees are enabled.

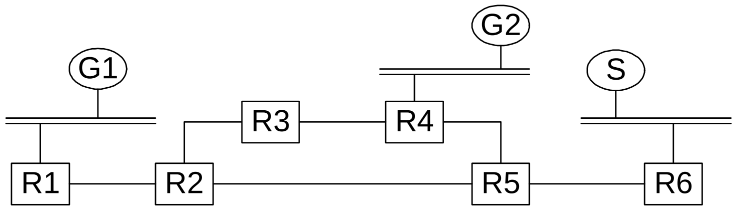

Kuva 1 / Figure 1

Model solution and grading

Nodes G1 and G2 join the multicast group as receivers (max 2p of the following)

- Router R1 detects the member G1 using IGMP and router R4 detects the receiver G2 using IGMP. (extra ½p)
- Routers R1 and R4 each send a **Join** (*, G) request towards R3 (=RP). (1p) *can be replaced with a figure indicating the Join messages*
- Each router (R1, R2, R4) forwards the message on the link that is on the shortest route to R2 according to their routing table. (½p)
- Each router (R1, R2, R3, R4) creates the tree by adding a routing table entry for all sources "*" in the group G indicating the downstream and the upstream interfaces on the multicast tree. (½p) *can be replaced with a figure*
- A router that is already on the tree does not need to forward the message (does not happen in this case). (½p)

Node S sends packets to the group (max 3p)

- Node S sends **packets to the rendezvous point R3 which resends** the packets on the multicast tree, so that the packets are received by G1 and G2. (1p) *can be replaced by a figure*
- The packets send in this first stage are encapsulated in a **Register** message and sent using unicast (1p)
- When R3 detects continuous traffic from R3 it sends a **source specific Join** (S, G) request to S. The intermediate nodes R4, R5 and R6 add routing table entries for source S in the group indicating the downstream and upstream interfaces for the tree. (½p)
- When R3 starts receiving packets on the tree from S it sends a **Register Stop** message to S. (½p)

Source-based trees are enabled (max 1p of the following)

- R1 and R4 detect lots of traffic from source S and send a **source specific Join (G, S) request directly towards S**. The intermediate routers R1, R2, R5, R6, R4 add routing table entries for source S in the group indicating the downstream and upstream interfaces for the tree. (1p) *½p for showing the tree as a figure but not mentioning Join messages*
- R1 and R4 also sends a **Prune** (*, G) message tree once messages are received on the (G, S) tree. (½p)
- Copies of the messages are still sent to R3. (½p)

Related slides

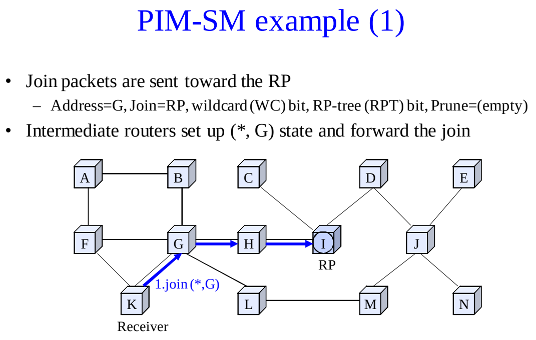

PIM-SM example (1)

- Join packets are sent toward the RP
  - Address=G, Join=RP, wildcard (WC) bit, RP-tree (RPT) bit, Prune=(empty)
- Intermediate routers set up (*, G) state and forward the join

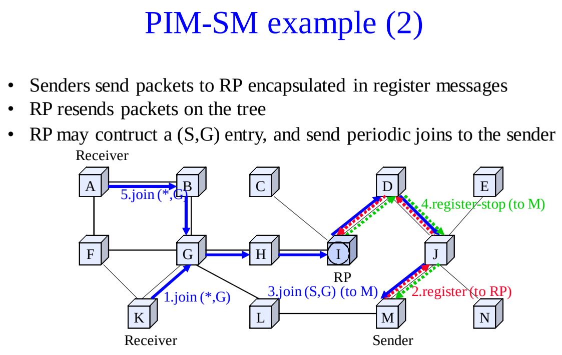

PIM-SM example (2)

- Senders send packets to RP encapsulated in register messages
- RP resends packets on the tree
- RP may contruct a (S,G) entry, and send periodic joins to the sender

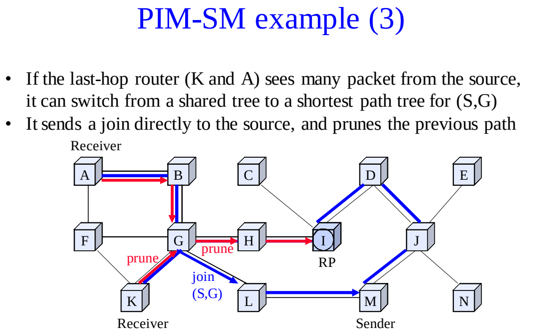

PIM-SM example (3)

- If the last-hop router (K and A) sees many packet from the source, it can switch from a shared tree to a shortest path tree for (S,G)
- It sends a join directly to the source, and prunes the previous path

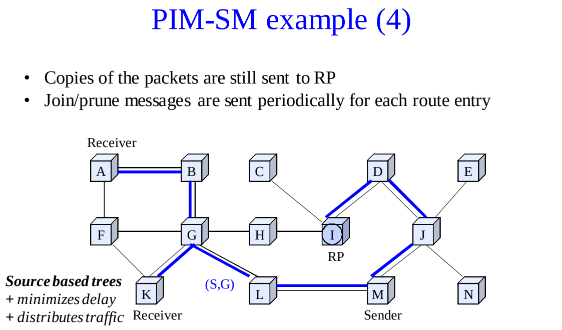

PIM-SM example (4)

- Copies of the packets are still sent to RP
- Join/prune messages are sent periodically for each route entry

Source based trees
+ minimizes delay
+ distributes traffic