S-38.2121 Routing in Communication Networks

Model solutions for the exam 30.10.2006

Question 1

1. Vertaile lokaalin ja globaalin tiedon käyttöä reitityksessä. Anna esimerkki PSTN reititysalgoritmista joka käyttää vain lokaalia tietoa. Anna esimerkki myös piirikytkentäisen verkon reititysjärjestelmästä, joka käyttää globaalia tietoa.
   Compare the use of global and local information in routing. Give an example of a PSTN routing algorithm that uses only local information. Give also an example of a routing system in the circuit switched network that uses global information.

Solution

Comparison between global and local information (max 3p of the following)

- Scalability: Using global information requires processing of lots of information and the collection and distribution is difficult. Using local information is more scalable. (1p)
- Vulnerability: If the use of global information is implemented as a centralized system, it is more vulnerable. For local information, failure has only local significance. (1p)
- Performance: Global information allows optimal routes to be generated, which leads to efficient use of the network. Algorithms using local information also try to be near optimal, but the limited available information reduces the efficiency. (1p)
- Speed: Algorithms using local information responds quickly to changes in the network. It takes time to distribute global information. (1p)

PSTN routing algorithm using local information (1,5p)

- E.g. FHR

PSTN routing algorithm using global information (1,5p)

- E.g. RCAR, TINA

Selection of route may be based on

Global information

- most efficient use of netw
- a lot of info, real time collection and distribution is difficult
- vulnerable if centralized
- E.g. TINA architecture

Local information

- solution is distributed and nodes are autonomous
- scales to a network of any size
- the goal is to find algorithms that are near optimal

RCAR - Residual Capacity Adaptive Routing is used in Canada

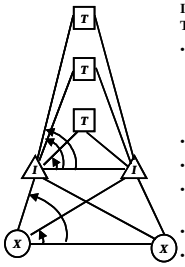

Implementation name DCR - dynamic call routing/ Telecom Canada

- **Info about outgoing circuit reservation status, nrof blocked calls and CPU load is collected each 10s to a centralized network management center. The Center calculates and downloads new routing tables for *I* and *T* switching nodes.**
- **Idea is to choose the path with most free circuits.**
- **Improves network performance significantly.**
- **Adapts quickly to unusual traffic patters and to link and node failures.**
- **Benefits relate to time difference between coasts.**
- **Vulnerable to failure of Mgt Center. Falls back to FHR model, if the Center stops.**

Common problems

- Any algorithm that has some centralized functionality (such as “manual” generation of routing tables) is not considered as an algorithm using global information. In strict sense, we only consider algorithms able to react in near real-time on global information about the network state (e.g. load).

Question 2

2. Kuvaa etäisyysvektoriprotokollan vastaanottoalgoritmi. Kuvaa etäisyysvektoriprotokollan toimintaperiaate pienen esimerkkiverkon avulla (verkossa ei ole vikoja ja kaikkien linkkien painot ovat 1).
   Describe the reception algorithm of a distance vector protocol. Describe the operational principles behind the distance vector protocol using a small example network (there are no faults in the network and the weight of every link is 1).

Solution

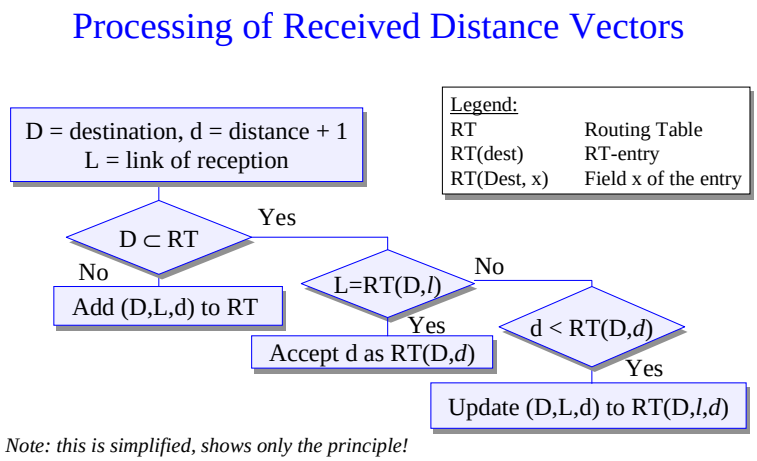

Note: this is simplified, shows only the principle!

Description of the reception algorithm

- Add the link cost to received distance (1p)
- Add the received DV to the routing table if the destination does not already exist in the routing table (1p)
- If the DV was received over the same link that the current routing table uses for the destination, then update even if the received distance is higher (1p)
- If the received DV has a smaller distance, then update the routing table. (1p)

Correct example that shows the principle (2p)

Common problems

- Many forgot to mention that the link cost is added to the received distance (-1p)

- More than half of the students forgot to check if the DV was received over the same link that currently is in the routing table for the destination. This step is critical in reacting to changes in the network. (-1p)
- The examples were generally very good.

Question 3

3. Esitä linkkitilatietueiden tyypit ja niiden käyttö OSPF:ssä.
   Present the types of link state records and their usage in OSPF.

Solution

Router LSA (2p)

- Describes the point-to-point links, transit networks and stub networks attached to a router.

Network LSA (1,5p)

- Describes which routers are connected to a transit (BC or NBMA) network.

Two types of Summary LSA (1,5p)

- Describes (sub)networks in another area or in an external network. These are reachable through an area-border router or a border router.

External LSA (1p)

- Describes destinations in an external network imported from some other routing protocol, e.g. BGP-4.

Group Membership LSA (extra 1p)

- Describes members of a multicast group used in MOSPF.

Not So Stubby Area LSA (extra 1p)

- Describes external destinations added to a stub area.

Maximum points 6p. For each LSA type, the name gives 0,5p and the description gives the rest.

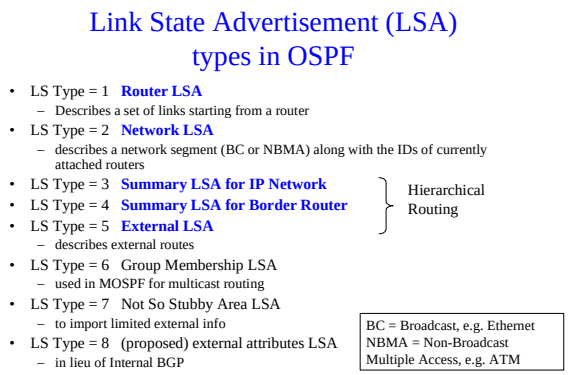

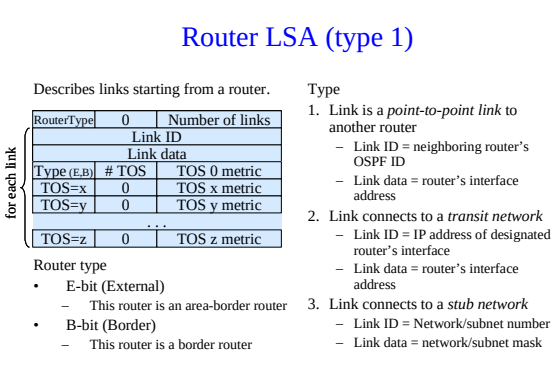

Network LSA (type 2)

Network mask
Attached router
Attached router
. . .
Attached router

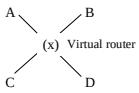

- Advertised by designated routers for transit networks
- Link state ID (in header) = interface ID of designated router
- Attached router = OSPF identifier of the attached router

Summary Link LSA (type 3,4)

Network mask		
0	0	TOS 0 metric
TOS=x	0	TOS x metric
TOS=y	0	TOS y metric
. . .		
TOS=z	0	TOS z metric

- For *IP networks* (type 3)
  - Network mask of network/subnet
  - Link state ID (in header) = IP network/subnet number
- For *border routers* (type 4)
  - Network mask = 0xFFFFFFFF
  - Link state ID (in header) = IP address of border router
- One separate advertisement for each destination

External Link LSA (type 5)

Network mask		
E,TOS=0	0	TOS 0 metric
External route tag (0)		
E,TOS=x	0	TOS x metric
External route tag (x)		
. . .		
E,TOS=z	0	TOS z metric
External route tag (z)		

- Advertised by border routers
  - Information from external gateway protocols (BGP-4)
- One destination per record
- Link state ID (in header) = IP network/subnet of destination
- Network mask = network/subnet mask
- E-bit indicates that distance is not comparable to internal metrics
  - Larger than any internal metric
- Route tag is only used by border routers (not used by OSPF)

Question 4

4. Tietueen ikään ja järjestysnumeroon liittyvät toiminnot OSPF:ssä.
   Explain the actions related to the age and the sequence number of the record in OSPF.

Solution

Sequence number (max 3p of the following)

- Sequence number is used in comparing which record is newer in flooding (1p) and database synchronization. (1p)
- Each time the router is flooding a LSA the sequence number is increased. The number follows a “lollipop” sequence space, and therefore circular comparison must be used. (1p)
- Description of how the sequence number is examined in the reception algorithm (1p)
- Description of what happens when a router is restarted. (1p)

Age (max 3p of the following)

- Age = hop count + age in seconds since advertisement. (1p)
- The age is compared if two LSAs have the same sequence number. (1p)
- A LSA with MaxAge (= 60 min) is removed. Removal is synchronized with the other routers. (1p)
- Other timers used in OSPF. (1p)

Flooding protocol distributes information about topology changes

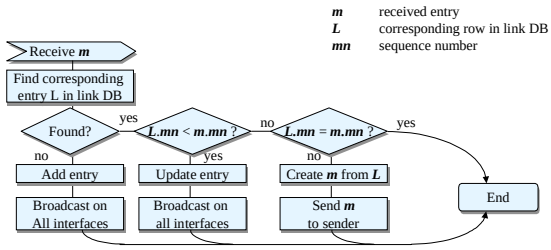

After reconnection of the islands "bringing up adjacencies" is required

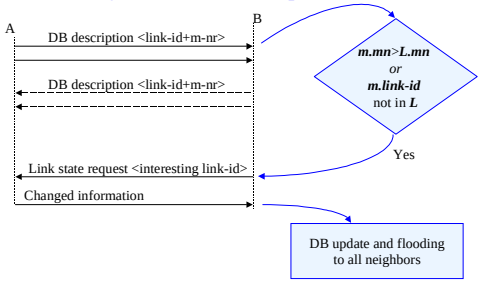

Link records have an age, old/dead ones are removed from Link DB (1)

- Old information must be removed from DB
- Every node must use the same information
  - ⇒ The removals must be synchronized
- The LSAs of OSPF have an age
  - Age = 0 when the advertisement is created
  - Age = number of hops through which the advertisement has traveled + seconds from reception
- Max age is 1 hour
  - Not used in the calculation of routes
  - Must be removed
- Every entry must be advertised at a 30 min interval.
  - The new advertisement zeroes the age and increments the sequence number.

Link records have an age, old/dead ones are removed from Link DB (2)

- When the age reaches MaxAge (= 1 h) the entry is removed
  - The router must send an advertisement to the neighbors when the aged entry is removed
- The flooding algorithm examines the age of the received advertisement
  1. MaxAge advertisement is accepted and flooded – this removes obsolete info.
  2. If the age difference of the advertisement to the DB is small, the advertisement is not flooded to avoid overloading the network with multiple copies of the same info. This is due to normal routing when the entry is received on different paths.
  3. If the age difference is large (>MaxAgeDiff), the newest advertisement is accepted and distributed. In this case, the router has probably been restarted.
  4. If a MaxAge record is not found, advertisement has not impact. The router most likely has already removed the dead LSA.

Common problems

- Most students knew how the lollipop sequence number space works
- Surprisingly many believe that a router can **request** an update of a record that is aging. Requesting records is only possible in database synchronization after reconnecting a link, but that is not triggered by the age.
- The question requires details, but gives much choise in what details are included. Too general answers don't give many points.

Question 5

5. Selitä miten monilähetysryhmään liittyminen ja pakettien lähettäminen ryhmälle tapahtuu PIM-SM protokollassa.
   Describe what happens when a node joins a multicast group in the PIM-SM protocol. Describe also how packets are sent to the group.

Solution

Joining (max 3p of the following)

- The joining node sends join message toward the rendezvous point. (1p)

- The join message travels until it reaches the RP or the first node already belonging to the multicast tree. (1p)
- Every node that forwards the join message creates state information about the group by adding an entry in their forwarding table. This entry indicates the interfaces belonging to the multicast tree. (1p)
- The receiving host joins the multicast group by replying to the local router's IGMP membership queries (not specific to and not part of PIM-SM) (extra 1/2p)
- Normal unicast routing tables are used in forwarding toward the RP. (extra 1/2p)

Sending (max 3p of the following)

- The sender send the first packets encapsulated in Register messages to the RP. (1p)
- The packets are forwarded along the multicast tree by the RP. (1p)
- If RP observes many packets, it sends a join toward the sender to create a source-specific route. Then it sends a Register-Stop when packets start arriving on the source-specific route. (1p)
- A router serving a receiver can switch to a source-specific tree by sending a join directly to the source. (extra 1p)

Common problems

- Many thought that the RP knows the individual receivers and the routes to them. This would mean that the RP unicasts or source-routes the packets to the receivers. This is very different from PIM-SM, where the state information is distributed along the path and is thus not centralized to the RP.
- The join message is sent toward the RP but not necessarily all the way to the RP. If some router is already part of the multicast tree it does not forward it further.
- It was nice to see so many students remembering the switching from shared trees to source-specific trees. For this, an extra point was given to replace some other missing details.

Question 6

6. Miten proaktiiviset ja reaktiiviset reititysmenetelmät eroavat toisistaan? Miten ZRP (Zone Routing Protocol) yhdistää proaktiivista ja reaktiivista reititystä? Mitä hyötyä on tästä yhdistelmästä?
   How do proactive and reactive routing methods differ? How does ZRP (Zone Routing Protocol) combine proactive and reactive routing? What are the advantages of this combination?

Solution

In proactive routing, the routing tables are generated before the packets can be sent. A routing protocol generates and maintains routes to all hosts in the network. (1p)

In reactive routing, the route is generated when it is needed. Routes are generated and maintained only between active senders and receivers. (1p)

ZRP uses proactive routing within a zone and reactive routing outside the zone. A zone of a node is defined as the neigbors that are within a given number of hops (zone radius) from this node. (2p)

Routes to nodes within the zone are immediately available. Most packets are assumed to be sent to nodes located nearby. Still a node does not have to maintan routes to all nodes in the network. (1p)

The proactive routing information can be utilized in the reactive routing as well, a concept called bordercasting. Thus, the query is sent directly to the peripheral nodes (not to all nodes in the zone) since the route to these is known. This makes flooding more efficient and scalable. (1p)

Common problems

- This question generally went well. Almost all understood the difference between proactive and reactive routing, and how ZRP combines them. Some answers had long descriptions of the properties of both approaches, which was positive although not necessary.
- The last sub-question (advantages of combining proactive and reactive routing) was the most difficult one and the one that tested the deeper understanding. There was quite much variation in the answers. Too general descriptions (e.g. “more scalable”) were given lower points because it does not tell the reason and how it relates to combining proactive and reactive routing.

General grading principles

Note that this document only describes the grading principles. The model solutions describe the main points that were expected to be included in the answer. It is not a strict requirement list. A good answer must clearly show that the subject is understood.

Generally small errors do not decrease the points. Serious errors decrease the points. Some extra information **related to the question** may give small extra points.