

IS-IS Routing Protocol – A Better Routing Mechanism Approach

S. P. Pingat
BCUD, University of Pune

A. A. Deshmukh
BCUD, University of Pune

ABSTRACT

IS-IS and OSPF are link state routing protocols. Both protocols suggest different procedures for this purpose. Since the procedures are quite different, but targeted to achieve the same goal; it is worthwhile comparing the two according to performance criteria common to both. The following criteria are chosen.

- The longest arrival time of an LSU packet at all the routers. Following this measure it can also find if all the LSA packets arrive successfully at all the routers before new instances of the packets are generated.
- The average arrival time of LSA packets at all the routers.
- The total required bandwidth in each scheme.
- The number of memory accesses a router performs in each scheme, which is evidence of the amount of internal work it performs.

Clearly, a scheme in which more LSA packets arrive at all the routers more quickly, with a smaller amount of internal work and with the use of a smaller amount of bandwidth and within the routers, is more efficient. In terms of arrival times of routing update packets it find that in our model of broadcast networks the method suggested in OSPF is giving less performance than that of IS-IS. In particular, the OSPF performance in consideration with average arrival time of routing update packets is 2–10 times longer than in IS-IS. There are scenarios where OSPF outperforms IS-IS and vice versa, in terms of the bandwidth each scheme consumes as well as the number of routers memory accesses perform in each scheme, IS-IS outperforms OSPF.

1. INTRODUCTION

Finding a path from any source to any destination in the network it uses a process called routing. Routing protocols are used to accomplish routing. Routing protocols are divided into two classes: 1) Distance vector protocols 2) Link state protocols. There are main two classes of routing protocols: distance vector protocols and link state protocols. In distance vector protocols a router maintains a vector of its distances to all destinations and periodically transmits this vector to its neighbors. Also, a router uses the distance vectors it receives from its neighbors to update its own distance vector. This process is guaranteed to converge at each router to the correct distances to all destinations. On the other hand, in link state protocols a router maintains a topology map of the routing domain in which it is located.

A routing protocol operates within a routing domain, which contains interconnected routers and networks. There are two types of networks, point-to-point and multi-access. A point-to-point network is a serial transmission link connecting two

routers. Multi-access networks are networks to which more than two routers are connected. There are two types of multi-access networks, broadcast and nonbroadcast. Broadcast networks are networks in which a router/host can send a single copy of a packet that will be received by all other routers/hosts attached to the network, such as Ethernet or fiber distributed data interface (FDDI). These networks typically consist of a single shared transmission link. Nonbroadcast networks are composed of an arbitrary topology of point-to-point serial links connecting packet switches. Examples of such networks are X.25, frame relay, and asynchronous transfer mode (ATM).

In what follows it only concentrate on link state protocols. In such protocols, in order to obtain and maintain up-to-date topological maps, routers exchange routing packets known as link state advertisements (LSAs). Every router generates and sends a local LSA that describes its connections to the other routers or networks. The topology map of a routing domain contains nodes connected by links. Usually, the map contains the routers and multi-access networks as nodes. Router interfaces to broadcast and nonbroadcast multi-access networks, and point-to-point networks, are considered links. It is also possible to consider routers connected to a single nonbroadcast multi-access network as connected to each other by point-to-point links in the topological map. This is known as the point-to-multipoint network model.

In this article it only consider multi-access broadcast networks and recall that such networks are considered nodes in the topological map. Thus, two questions arise: how the routers' connectivity over a broadcast network is described, and who generates LSAs on behalf of a broadcast network.

In link state protocols, LSAs are generated periodically or when topological changes occur. In order to ensure identical topological maps at all the routers, every LSA needs to arrive at every router. Flooding does the dissemination of LSAs, which is a method in which a router transmits an LSA to all its neighbor routers except to the one from which the LSA arrived. In this context, an important question is how LSAs are disseminated reliably over broadcast networks. In two main standard widely used link state protocols, Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS), suggested by the Internet Engineering Task Force (IETF) [2] and International Organization for Standardization (ISO) [3], the DR has a major role in dissemination, but in different ways. In this article it compare the procedures of OSPF and IS-IS according to the following performance measures:

- The longest time it takes for a LSA packet to arrive at all routers
- The average arrival time of LSA packets at all routers

- The total required bandwidth in each scheme
- The number of memory accesses a router performs in each scheme, which is evidence of the amount of internal work it performs.

In terms of arrival times of routing update packets it find that in our model of broadcast networks the method suggested in OSPF is giving less performance than that of IS-IS. In particular, the OSPF performance in consideration with average arrival time of routing update packets is 2–10 times longer than in IS-IS. There are scenarios where OSPF outperforms IS-IS and vice versa, in terms of the bandwidth each scheme consumes as well as the number of routers memory accesses perform in each scheme, IS-IS outperforms OSPF.

2. LSA PACKET DISSEMINATION OVER BROADCAST NETWORKS IN IS-IS AND OSPF

In this section it describe the procedures for reliable dissemination of LSA packets over broadcast networks in IS-IS and OSPF. The description of the OSPF scheme is based on RFC 2178. That of IS-IS is based on RFCs 1142 and 1195. In the description of IS-IS it assume that IS-IS operates in a TCP/IP environment, so the information carried in IS-IS is slightly changed to be suitable for the TCP/IP environment.

2.1 The Dissemination of LSA Packets in IS-IS

Reliable delivery of LSA packets over broadcast networks in IS-IS is based on the DR periodically announcing which LSA packets it has in its database. All the other routers compare their database content with that of the DR and either request from the DR LSAs they do not have or transmit to the DR LSAs it does not have. More specifically, when a router generates a new LSA or receives a new one on one of its other connections (not the one to the BN) it transmits the LSA to all the routers attached to the BN. In parallel, the DR periodically transmits CSNP packets in which it lists all the LSAs it has in its topological database. Every router R_i compares the information in every received CSNP packet with its routing database and does the following: if it notices that the DR does not have an LSA it itself has, or that the copy of the DR is older, R_i schedules its copy of the LSA for transmission into the BN. If R_i notices that the DR has an LSA it does not have, or notices that the DR's copy of the LSA is newer than the one it has, it requests the LSA from the DR by transmitting a PSNP packet. When the DR receives a PSNP it schedules the requested LSA packets for transmission into the BN. The DR transmits its CSNP packets once every time interval of completeSNPInterval time units. A router R_i transmits a PSNP packet once every time interval of partialSNPInterval time units if it needs to send one (i.e., if it notices it is missing LSA packets). All the packets are sent to the ALLSPFRouters multicast address.

In addition, when a router receives an LSA on the interface to the BN that is older than the copy it itself has in its routing database; it schedules the more recent LSA for transmission into the BN.

Several routers may notice that the DR does not have a certain LSA and all schedule the LSA for transmission. In this case, if a router later receives the LSA successfully on the BN, it

cancel the scheduling of the LSA transmission because it assumes that the DR has also received this copy of the LSA.

Finally, as mentioned, the DR transmits the LSA on behalf of the BN. The actions performed for this task are similar to those every router performs concerning its local LSAs.

2.2 The Dissemination of LSA Packets in OSPF

In OSPF the reliable dissemination of LSA packets is guaranteed via acknowledgments. There are three methods to acknowledge LSAs. The first method is by delayed ACKs. A delayed ACK is generated by a router upon successful reception of an LSA; it is not always transmitted immediately upon generation, but is sometimes delayed to enable the receiving router to collect and send several acknowledgments together.

The second method is by directed ACKs. A directed ACK is submitted for transmission immediately upon generation. The third method is by implied ACKs. Consider the case where a router R_i transmits an LSA to another router R_j and waits for an Ack. If R_i later receives the same LSA from R_j , the received LSA serves as an implied ACK since it signals R_i that R_j has the considered LSA. R_i maintains a retransmission list for every adjacent neighbor router R_j . In this list R_i keeps all the LSAs that were sent to R_j but have not yet been acknowledged. When R_i transmits an LSA to R_j , it starts a timer that expires after RxmtInterval time units. If within this timeout interval an acknowledgment of any type is not received from R_j for the LSA, the LSA is retransmitted. When an ACK is received, the LSA is deleted from the retransmission list.

The major idea behind the scheme for the reliable dissemination of LSA packets is that every router, when generating or receiving a new copy of an LSA, is trying to guarantee that all its adjacent neighbors also receive this LSA. Therefore, every LSA that a router R_i wants to flood over the BN is sent to the DR and BDR only. R_i then waits to receive an ACK from both. Then it is the responsibility of the DR to flood the LSA reliably over the BN to the rest of the routers. Second, the DR and BDR are responsible to reliably flood over the BN any new LSAs they generate or receive from other connections (i.e., from connections not to the BN). As mentioned, the DR has a special role in the dissemination procedure. When it generates or receives a new or more recent copy of an LSA on another connection, not to the BN, it generates a flooding packet and tries to disseminate the LSA to all other routers on the BN. The LSA is also inserted into the retransmission list of every other router R_i so that if an ACK is not received from R_i within the RxmtInterval timeout interval, the LSA is retransmitted to R_i . When the DR receives a new or more recent copy of an LSA on the connection to the BN, if it is received from the BDR, the DR transmits a delayed ACK and inserts the LSA into the retransmission list of every other router R_i except that of the BDR. If an ACK is not received from R_i within the RxmtInterval timeout interval, the DR will retransmit the LSA to R_i (i.e., it helps the BDR to disseminate the LSA). If the LSA is received from a router R , the DR is responsible for the dissemination of the LSA over the BN. For this purpose it generates a flooding packet and inserts the LSA into the retransmission list of every router except that of R .

The flooding packet itself serves as an implied ACK to router R . Finally, if a duplicate copy of an LSA is received on the connection to the BN and does not serve as an implied ACK,

the DR acknowledges its reception immediately. The LSA is considered a retransmission from a router which assumes that the DR did not yet receive the LSA. The BDR, as the DR, is responsible for the reliable dissemination of a new or more recent copy of an LSA it generates or receives from another connection, not to the BN. Therefore, in this case the BDR has the same rules as the DR. Should such an LSA arrive on the connection to the BN, if it arrives from the DR, the BDR acknowledges the LSA with a delayed ACK and inserts the LSA into the retransmission lists of all adjacent routers, in order to ensure reliable dissemination if the DR fails. If the LSA arrives from a router R, the BDR again tries to ensure reliable dissemination by inserting the LSA into the retransmission lists of all adjacent routers. However, it does not flood the LSA immediately, since it leaves this role to the DR. Notice that the BDR does not transmit an ACK immediately in this case. An ACK will be sent to all the routers when the LSA arrives again, as a duplicate, from the DR. The reason for this rule is described below.

If a duplicate copy of an LSA arrives from the DR and serves as an implied ACK, the BDR transmits a delayed ACK to ALLSPFRouters. This occurs when the BDR received the LSA before, as a new LSA, from a router R. At that time, the BDR had to know if the DR received the LSA also, and to acknowledge its reception to all the routers. However, it is not clear if all the routers received the LSA yet. Therefore, the BDR waits and transmits an ACK packet only after receiving the LSA again, as a duplicate, from the DR. At this time it is guaranteed that all the routers received the LSA by flooding from the DR, and they wait for an ACK from the BDR.

Finally, if the LSA was not an implied ACK, it was sent again by the DR or a router R. It shall be acknowledged immediately since the sending router assumes that the BDR did not receive the LSA yet. When considering a router R_i we must remember that its adjacent neighbors are the DR and BDR only. Therefore, when a router generates or receives a new or a more recent LSA from another connection, not to BN, it only tries to disseminate the LSA to the DR and BDR. When a new or more recent LSA arrives on a connection to the BN, R_i only refers to such LSAs that arrive from the DR or BDR. In this case R_i transmits a delayed ACK to the transmitter of the LSA and inserts the LSA into the retransmission list of the other router in order to ensure that both the DR and BDR receive the LSA. In this case R_i takes part in the effort to reliably disseminate routing information.

Finally, when a duplicate LSA is received on the connection to the BN from the DR or BDR and it does not serve as an implied ACK, R_i transmits a directed ACK to the transmitter of the LSA. It is assumed that R_i received the LSA again by a retransmission from the DR or BDR, so it shall be acknowledged.

3. SIMULATION AND COMPARISON

In our simulation we assume that time is divided into time slots of length LSRefreshTime seconds each. At the beginning of each slot all the routers attached to the BN generate new LSU packets for transmission into the BN. As we show later for both OSPF and IS-IS, all the routing packets arrive at all the routers within a slot when LSRefreshTime is on the order of several minutes (e.g., 30 min), as is the case in real implementations of OSPF and IS-IS. Therefore, it is enough to simulate a single slot, and recall that we assume that only routing traffic is transmitted in the network and there are no transmission errors (i.e., all packets arrive at their destinations). The above-mentioned performance criteria are

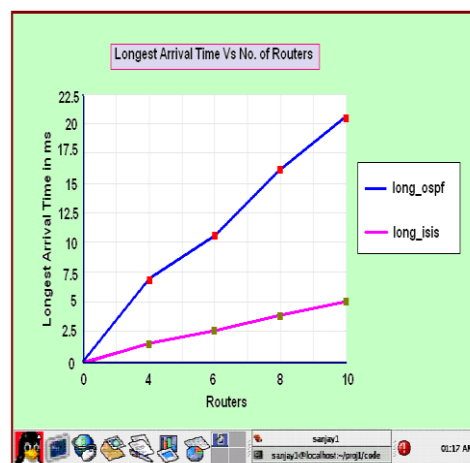
therefore defined as follows: The longest arrival time of an LSU packet at all the routers is the time since the beginning of a slot and until the last LSU arrives at all the routers.

- The average arrival time of LSU packets at all the routers is computed as follows: for every LSU generated at the beginning of a slot and arriving at all the routers during the slot, it measures the time elapsed since the beginning of the slot until the LSU arrives at all the routers. Then, assume it measures these time intervals for 1 LSU packets. Then divide the sum of all the arrival times by 1; the result is the average arrival time.
- The bandwidth consumed in every scheme is computed by measuring the fraction of transmission time each scheme consumes in a slot. The fraction of transmission time in each scheme is computed by summing the transmission time of all the routing update packets transmitted within a slot and dividing this sum by the slot length, LSRefreshTime.
- The number of memory accesses a router performs in each scheme is the number of memory accesses it performs in a slot when handling the routing traffic. The exact definition of the cases considered to be memory accesses is given later.

4. RESULT ANALYSIS USING GRAPH

Longest Arrival Time

The longest arrival time for OSPF and IS-IS is compared by taking four different networks models. The first network model has 4 routers, second network model has 6 routers, third network model has 8 routers and fourth network model has 10 routers.

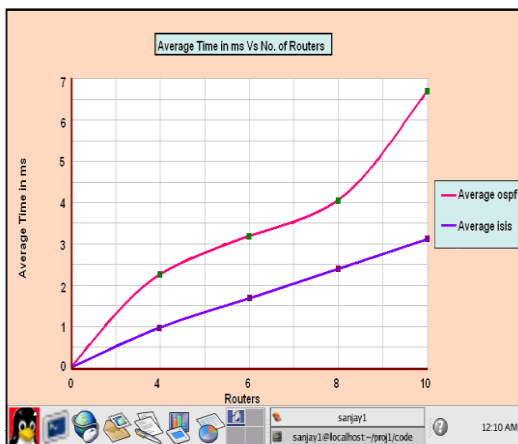


Graph 1

From Graph 1, it is observed that the longest arrival time for OSPF is larger than IS-IS. The difference for longest arrival time of OSPF and IS-IS increases as number of routers increases.

Average Arrival Time

The average arrival time for OSPF and IS-IS is compared by taking four different networks models. The first network model has 4 routers, second network model has 6 routers, third network model has 8 routers and fourth network model has 10 routers.

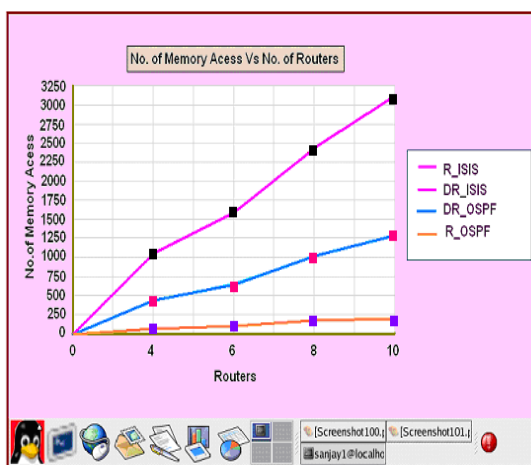


Graph2

From Graph 2, it is observed that the average arrival time for OSPF is larger than IS-IS. The difference for average arrival time of OSPF and IS-IS increases as number of routers increases.

Number of Memory Accesses

The Number of Memory Accesses for OSPF and IS-IS is compared by taking four different networks models. The first network model has 4 routers, second network model has 6 routers, third network model has 8 routers and fourth network model has 10 routers.

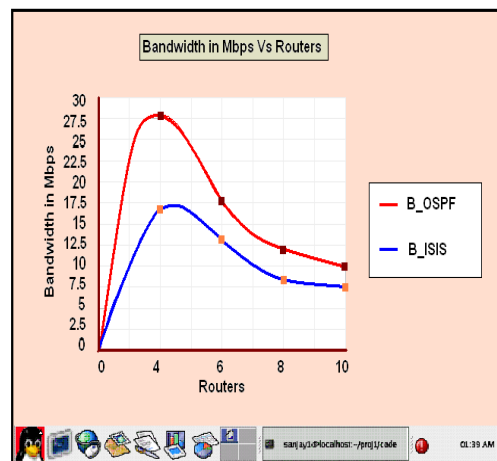


Graph 3

From Graph 3, it is observed that the Number of Memory Accesses for IS-IS is larger than OSPF. The difference between the Number of Memory Accesses for OSPF and IS-IS increases as number of routers increases. The Number of Memory Accesses of router and designated routers for IS-IS are same. The Number of Memory Accesses of router and designated routers for OSPF are different. In OSPF, the Number of Memory Accesses of designated routers are more than that of routers.

Bandwidth

The bandwidth for OSPF and IS-IS is compared by taking four different networks models. The first network model has 4 routers, second network model has 6 routers, third network model has 8 routers and fourth network model has 10 routers.



Graph 4

From Graph 4, it is observed that the bandwidth required for OSPF is larger than IS-IS. The bandwidth of OSPF decreases as the number of routers increases. The bandwidth of IS-IS also decreases as the number of routers increases but the rate of decrease in IS-IS is less than OSPF.

5. CONCLUSION

IS-IS outperforms OSPF in the arrival times of LSU packets because of two reasons. First, the routers transmit LSack packets, which delay the arrival of the token at the various routers. Second, in OSPF LSU packets are first sent to the DR, and then the DR disseminates the LSU packets to all the routers.

The arrival times of routing update packets are considered the most important performance criterion. Therefore, a conclusion from the study in this article is that it is worthwhile to change the scheme for routing information dissemination over broadcast networks currently defined in OSPF to that of IS-IS. Also, routers that are connected to broadcast networks and capable of running both IS-IS and OSPF shall be configured to run IS-IS. It would like to mention at this point that in many aspects it have checked a worst case model where the transmissions of the LSU packets are synchronized, so the load on the network comes in bursts. This is a bad scenario because routers need to wait the longest time before they can transmit their LSU packets. This is apparent in the times required to disseminate routing information. Also, the topology of the area makes the broadcast network a bottleneck, and LSA packets from different parts of the area

arrive at the broadcast network only through one router. If they can arrive at the broadcast network through several routers, they can arrive at all the routers attached to the broadcast network more quickly.

6. FUTURE WORK

In future IS-IS routing protocol can be used in major cases instead of OSPF since IS-IS outperforms OSPF in the arrival times of LSU packets.

The arrival times of routing update packets are considered the most important performance criterion. It is worthwhile to change the scheme for routing information dissemination over broadcast networks currently defined in OSPF to that of IS-IS. Also, routers that are connected to broadcast networks and capable of running both IS-IS and OSPF shall be configured to run IS-IS.

7. REFERENCES

- [1] J. Moy, *OSPF -- Anatomy of an Internet Routing Protocol*, Addison-Wesley, 1998.
- [2] J. Moy, "OSPF Version 2," RFC 2178, July 1997.
- [3] D. Oran, "OSI IS-IS Intra-domain Routing Protocol," RFC 1142, Feb. 1990.
- [4] C. Huitema, *Routing in the Internet*, Prentice Hall, 1995.
- [5] R. Perlman, *Interconnection -- Bridges and Routers*, Addison-Wesley, 1995.
- [6] "Intermediate System to Intermediate System Intra-domain Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)," ISO-DP-10589, Feb. 1990.
- [7] R. Callon, "Use of OSI IS-IS for Routing in TCP/IP and Dual Environment," RFC 1195, Dec 1990.
- [8] D. Katz, "Transmission of IP and ARP over FDDI Networks," RFC 1390, Jan. 1993.
- [9] S. Floyd, V. Jacobson, "The Synchronisation of Periodic Routing Messages," *SIGCOMM*, San Francisco CA, 1993, pp.33-44.
- [10] J. H. Fong, A. C. Gilbert, S. Kannan, and M. J. Strauss, "Better alternatives to OSPF routing," *Algorithmica*, vol. 43, no. 1-2, pp. 113–131, 2005.
- [11] D. Xu, M. Chiang, and J. Rexford, "DEFT: Distributed exponentiallyweighted flow splitting," in *INFOCOM'07*, Anchorage, AK, May 2007.