

Avoiding Loops and Packet Losses in ISP Networks

V.Sunilkumar¹

¹ GIT/ GITAM UNIVERSTIY

Received: 28 July 2011 Accepted: 17 August 2011 Published: 27 August 2011

Abstract

Even in well managed Large ISP networks, failures of links and routers are common. Due to these failures the routers update their routing tables. Transient loops can occur in the networks when the routers adapt their forwarding tables. In this paper, a new approach is proposed that lets the network converge to its optimal state without loops and the related packet lossless. The mechanism (OUTFC-Ordered Updating Technique with Fast Convergence) is based on an ordering of the updates of the forwarding tables of the routers and fast convergence. Typically we have chosen a Network consisting of routers and Link costs for simulation. Link failures are simulated. Avoiding transient loops in each case is demonstrated, by constructing a Reverse Shortest PathTree (RSPT).

Index terms— ISP networks, OUTFC (Ordered Updating Technique with Fast Convergence), Link Failures, Reverse Shortest Path Tree (RSPT).

1 INTRODUCTION

he link-state intra domain routing protocols that are used in ISP network [1] [2], were designed when IP networks were research networks carrying besteffort packets. The same protocols are now used in large commercial LSPs with stringent Service Level Agreements (SLA). Furthermore, for most Internet Service Providers, fast convergence in case of failures is a key problem that must be solved. Today, customers are requiring 99.99% reliability or better and providers try to avoid all packet losses.

Transient loops can be occurred due the topological change in the network when links failure occurred in network [1]. A network typically contains point-to-point links and LAN Point-to-point links are typically used between Points of Presence (POPs) while LANs are mainly used inside POPs. When a point-topoint link fails, two cases are possible. If the link is not locally protected, the IGP should converge as quickly as possible. Another source of changes in IP networks are the IGP metrics. Today, network operators often change IGP metrics manually to reroute some traffic in case of sudden traffic increase. Second type of important events is those that affect routers. Routers can fail abruptly, but often routers need to be rebooted for software upgrades.

Author ? ? ? : Department of IT, E-mails : avrohit2000@gmail.com , sunil.veee@gmail.com, smkr-ishna@ieee.org, Author : Department of CSE, GITAM University E-mail : siddhartha.gundapaneni@gmail.com To avoid transient loops during the convergence of link-state protocols, we propose to force the routers to update their FIB by respecting an ordering that will ensure the consistency of the FIB of the routers during the whole convergence phase of the network [1]. In the context of a predictable maintenance operation, the resources undergoing the maintenance will be kept up until the routers have updated their FIB and no longer use the links to forward packets. In the case of a sudden failure of a link that is protected with a Fast Reroute technique, the proposed ordering ensures that a packet entering the network will either follow a consistent path to its destination by avoiding the failed component or reach the router adjacent to the failure and will be deviated by the Fast Reroute technique to a node that is not affected by the failure, so that it will finally reach its destination.

2 II.

3 OUR APPROACH

Studies on the occurrence of failures in a backbone network have shown that failures of links and routers are common even in a well managed network [1]. On the other hand, an increasing number of users and services are relying on the Internet and expecting it to be always available. In order to ensure high availability in spite of failures, a routing scheme needs to quickly restore forwarding to affected destinations. Traditional routing schemes such as OSPF trigger link state advertisements in response to a change in topology, and cause network-wide recomputation of routing tables. Such a global rerouting incurs some delay before traffic forwarding can resume on alternate paths. During this convergence delay, routers may have inconsistent views of the network, resulting in forwarding loops and dropped packets [2].

OUTCF [7] was recently proposed to address the above concerns and achieves three interconnected objectives: 1) loop-free forwarding; 2) minimal convergence delay. At no time can a forwarding loop happen with OUTCF in the case of a single failure. OUTCF also reduces the period of disruption when packets are dropped due to the lack of valid routes. Lastly, OUTCF minimizes the convergence delay, i.e., packets are forwarded along optimal paths and the network is ready to absorb another change as soon as possible. The drawback of OUTCF, however, is that it requires each packet to carry the cost of the remaining

4 RELATED WORK

The problem of avoiding transient loops during IGP convergence has rarely been studied in the literature although many authors have proposed solutions to provide loop-free routing. An existing approach to loop-free rerouting in a link state IGP [8] requires that the rerouting routers take care of routing consistency for each of their compromised destinations, separately. In fact, those mechanisms were inspired by distance-vector protocols providing a transiently loopfree convergence [7]. With this kind of approach, a router should ask and wait clearance from its neighbours for each destination for which it has to reroute. This implies a potentially large number of messages exchanged between routers, when many destinations are impacted by the failure. Every time a router receives clearance from its neighbours for a given destination, it can only update forwarding information for this particular one. This solution would not fit well in a Tier-1 ISP topology where many destinations can be impacted by a single topological change. Indeed, in such networks, it is common to have a few thousands of prefixes advertised in the IGP [5]. Note that those solutions do not consider the problem of traffic loss in the case of a planned link shutdown. In [6], a new type of routing protocol allowing improving the resilience of IP networks was proposed. This solution imposes some restrictions on the network topology and expensive computations on the routers. Moreover, they do not address the transient issues that occur during the convergence of their routing protocol. In [4], extensions to link-state routing protocols are proposed to distribute link state packets to a subset of the routers after a failure. This fastens the IGP convergence, but does not solve the transient routing problems and may cause suboptimal routing.

In [2][3], transient loops are avoided when possible by using distinct FIB states in each interface of the routers. Upon a link failure, the network does not converge to the shortest paths. Based on the new topology. Indeed, the failure is not reported. Instead, the routers adjacent to the failed link forward packets along alternate links, and other routers are prepared to forward packets arriving from an unusual interface in a consistent fashion towards the destination. As such, the solution is a Fast Reroute technique. Our solution is orthogonal to [9] as our goal is to let the network actually converge to its optimal forwarding state by avoiding transient forwarding loops when a Fast Reroute mechanism has been activated, or when the failure is planned.

IV.

5 METHOD TO HANDLE LOOPS

Each router will maintain one waiting list associated with each link being shut down during the RSPT computations. A rerouting router R will update its FIB for a destination (which means that its paths to contain one or more links of the SRLG) once it has received the completion messages that unlock the FIB update in for one of the links being shut down. When updating its FIB, selects the outgoing interfaces for destination according to the new topology, i.e., by considering the removal or the metric increase of all the affected links. The meaning of a completion message concerning a link sent by a router is that has updated its FIB for all the destinations that it was reaching via before the event [8]. Let us now show that if a packet with destination reaches a rerouting router that has not performed its FIB update for destination, then all the routers on its paths to cannot have performed a FIB update for. If has not updated its FIB for destination, it cannot have sent a completion message for any of the failing links that it uses to reach. The failing links that a router on uses to reach are used by to reach, so that cannot have received all the necessary completion messages for any of those links. In other words, did not send a completion message for the links that it uses to reach. Thus, locks the FIB update for those links long its paths towards them. We provide the pseudo code that implements the ordering with completion messages. To process the metric increase (or shutdown) of a set of link , a router will compute the reverse SPT rooted on each link belonging to , that it uses in its current, outdated SPT. During this computation, it will obtain the rank associated with. It will then record the next-hops that it uses to reach in a

list. These are the neighbors to which it will send a completion message concerning link. If the rank associated with a link is equal to zero, then updates its FIB directly for the destinations that it reaches via this link, and it sends a completion message to the corresponding next-hops. In the other cases, builds the waiting list associated with, containing the neighbors that are using to reach, and it starts the timer considering the rank associated with this link [9]. Once a waiting list for a link becomes empty or its associated timer elapses, can update its FIB for all the destinations that it reached via this link and send its own completion message towards the neighbors that it used to reach the link. We consider a network to explain how to avoid the transient loops occur in the network by converging link state routing protocol. The Indian cities are connected in this net work like Mumbai (MUM), New Delhi(ND), Hyderabad(HYD), Madras etc.

6 // Computation of the

To understand this problem, let us consider the Internet2/Abilene backbone. Fig. ?? shows the IGP topology of this network. Assume that the link between MUM and BPL fails but was protected by an MPLS tunnel between BPL and MUM via JPR and BHU. When JPR receives a packet with destination BAN, it forwards it to BPL, which forwards it back to JPR, but inside the protection tunnel, so that MUM will decapsulate the packet, and forwards it to its destination, BAN.

This suboptimal routing should not last long, and thus after a while the routers must converge, i.e., adapt to the new shortest paths inside the network, and remove the tunnel. As the link is protected, the reach ability of the destinations is still ensured and thus the adaptation to the topological change should be done by avoiding transient loops rather than by urging the updates on each router. The new LSP generated by BPL indicates that BPL is now only connected to RAJ and JPR. Before the failure, the shortest path from LUK to MUM, BAN, CHE and HYD was via ND, RAJ and BPL. After the failure, ND will send its packets to MUM, BAN, CHE and HYD via LUK, JPR and BHU. During the IGP convergence following the failure of link MUM?BPL, transient loops may occur between ND and LUK depending on the order of the forwarding table updates performed by the routers. If ND updates its FIB before LUK, the packets sent by ND to MUM via LUK will loop on the LUK-ND link. To avoid causing a transient loop between LUK and ND, LUK should update its FIB before ND for this particular failure. A detailed analysis of the Internet2 topology shows that transient routing loops may occur during the failure of most links, except CHE?BAN and CHE?HYD. The duration of each loop will depend on how and when the FIB of each router is updated. Measurements on commercial routers have shown that updating the FIB may require several hundred of milliseconds. Transient routing loops of hundred milliseconds or more are thus possible and have been measured in real networks. As shown with the simple example above, the transient routing loops depend on the ordering of the updates of the FIBs. In the remainder of this paper, This proof is constructive as we give an algorithm that routers can apply to compute the ranks that let them respect the proposed ordering.

7 CONVERGENCE TIMES IN ISP NETWORKS

In this section, we analyze by simulations the convergence time of the proposed technique, in the case of a link down event. The results obtained for link up events are very similar. Indeed, the updates that are performed in the FIB of each router for the shutdown of a link impact the same prefixes for the linkup of the link. The only difference in the case of a link up is that the routers do not need to compute a reverse Shortest Path Tree. As no packet are lost during the convergence process.

8 CONCLUSION

In the proposed work, we have initially described the various types of topological changes that can occur in large IP networks. When failures occurs in the network the routers updates routing tables. Those updates may cause transient loops and each loop may cause packet losses or delays. Large ISPs require events. To protect the network from transient loops, we propose OUTFC method that it is useful to define an ordering on the updates of the FIBs. We have proposed an ordering applicable for the failures of protected links and the increase of a link metric and another ordering for the establishment of a new link or the decrease of a link metric. We have shown by simulations that our method avoids the loops and converges network to its optimal state. ¹

¹© 2011 Global Journals Inc. (US) Global Journal of Computer Science and Technology Volume XI Issue XV Version I



Figure 1:

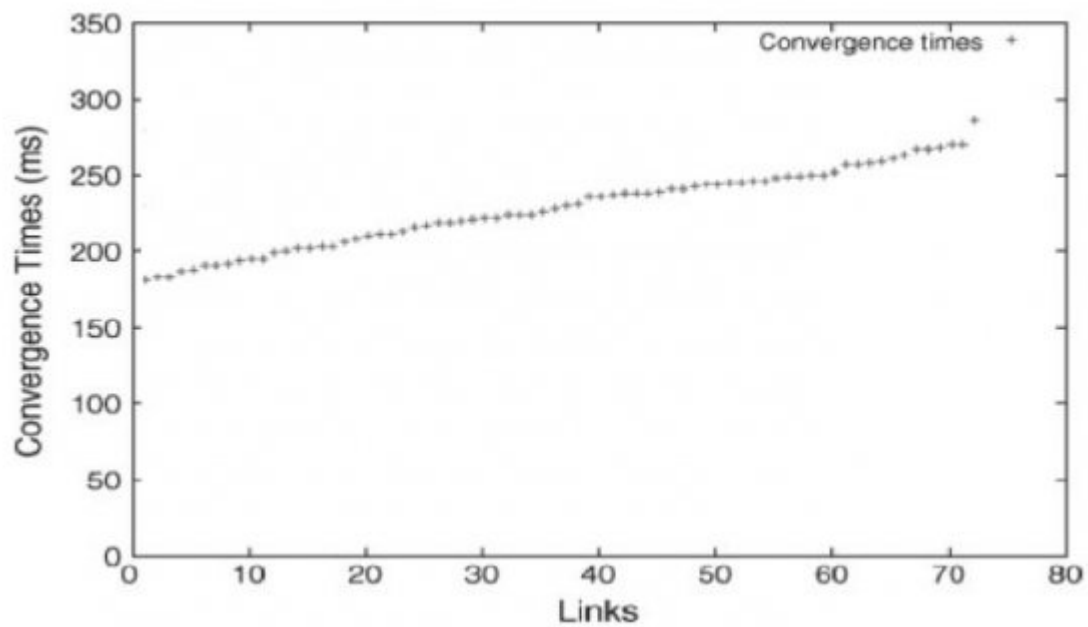


Figure 2:

148 [Iso/Iec (2002)] , Iso/Iec . 10589:2002. April 2002. (Tech. Rep.)

149 [Francois et al. (2005)] ‘Achieving subsecond IGP convergence in large IP networks’. P Francois , C Filsfils , J Evans

150 , O Bonaventure . *Computer Communication Review* 2005. 35 (3) p. .

151 [Shaikh et al.] ‘Avoiding Instability during Graceful Shutdown of OSPF’. A Shaikh , R Dube , A Varma .

152 *REFERENCES REFERENCES REFERENCIAS*,

153 [Francois and Bonaventure (2005)] ‘Avoiding transient loops during IGP Convergence in IP Networks’. P Francois

154 , O Bonaventure . *Proc. IEEE INFOCOM*, (IEEE INFOCOM) March 2005.

155 [Avoiding Transient Loops during the Convergence of Link-State Routing Protocols Pierre Francois, Member, IEEE]

156 *Avoiding Transient Loops during the Convergence of Link-State Routing Protocols Pierre Francois, Member,*

157 *IEEE*, Bonaventure, Member, IEEE.

158 [Pan et al. (2005)] *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, P Pan , G Swallow , A Atlas . Internet

159 RFC 4090. May 2005.

160 [Iannaccone et al. (2004)] ‘Feasibility of IP restoration in a tier-1 backbone’. G Iannaccone , C Chuah , S

161 Bhattacharyya , C Diot . *IEEE Network Magazine* January-February 2004.

162 [Intermediate system to intermediate system routing information exchange protocol for use in conjunction with the protocol for p

163 *Intermediate system to intermediate system routing information exchange protocol for use in conjunction*

164 *with the protocol for providing the connectionless-mode network service*, ISO (iso 8473)

165 [Shand and Bryant (2006)] *IP Fast Reroute Framework*, M Shand , S Bryant . October 2006.

166 [Alaettinoglu et al. (2000)] *Towards millisecond IGP convergence*, C Alaettinoglu , V Jacobson , H Yu .

167 ISISconvergence- 00. November 2000. (internet draft, draft-alaettinoglu)