Author's Copy

Fast BGP Convergence Following Link/Router Failure

Swapan Kumar Ray[1] and Susmit Shannigrahi[1]

Computer Science and Engineering Department, Jadavpur University, Kolkata -700032, India
skray[AT]ieee[DOT]org, susmit[AT]fedoraproject[DOT]org

Abstract. A Modified Border Gateway Protocol (MBGP) has been proposed towards achieving faster BGP convergence in the Internet following link/router/network failures. MBGP adopts the overall strategy of distributed fault detection-cum-identification, fault notification and rediscovery-cum-readvertisement of valid routes. In the assumed simplified model of the Internet, the sole MBGP router in each autonomous system (AS) identifies any failed component using the novel concept of special neighbors and notifies the identity of the failed component to all the MBGP routers in the Internet. Six new messages, including a query-response message pair and four permanent withdrawal messages, have been proposed in MBGP, without changing the BGP message format. The path exploration problem is significantly reduced because some failures cause no path exploration, the others do but only in a small number of nearby routers and, finally, no invalid messages are ever exchanged. Simulation studies have demonstrated significantly faster convergence of MBGP over BGP.

Keywords: BGP Convergence, Slow Convergence in BGP, Fast BGP Convergence, Link or Router Failure in Internet, Path Exploration in BGP, Modified BGP, Special Neighbors.

1 Introduction

The Border Gateway Protocol (BGP) [1] is the de-facto standard for the inter-domain or inter-autonomous system (AS) routing protocol in the Internet. Unfortunately, BGP suffers from the problem of unstable routing and slow convergence following events like the failure of a link or a router, change of AS policies, failure or resetting of the underlying TCP connections, etc. [2]-[4]. This slow convergence of BGP is considered a serious problem for the growth of the Internet because of reasons like excessive loss/delay of packets which hamper the performance of applications like VoIP, streaming video, etc., and cause severe congestion and router overloads in the Internet.

The main reason behind the delayed protocol convergence in BGP is the so-called path exploration phenomenon that is present in all path vector protocols like BGP because they are inherently associated with path dependencies which refers to a recursive path learning phenomenon. The path selected by a router depends on paths learnt by its neighbors; the latter, in turn, depends on what the neighbors have learnt from their neighbors; and so on. Thus, in BGP, following a failure event, some of the paths become invalid so that routers go through a cycle of selecting and propagating invalid paths till all routers in the Internet have learnt valid paths after all obsolete paths have been explored and invalidated. Solving the path exploration problem in BGP is hard and

it is made even harder because BGP allows arbitrary choice of import, export and route selection policies[2]. However, the policy aspects of BGP have not been considered in this paper.

In this paper, we have proposed a method called Modified BGP (MBGP) where each router periodically monitors its immediate neighborhood to detect any failure occurring in any of its neighboring routers, connecting links, or its own internal network. In case a router or a link or a network is found to have failed, the monitoring router first broadcasts, through flooding over the entire Internet, "Permanent Withdrawal (till repair) of the failed component. Immediately, thereafter, it discovers locally optimum alternative valid routes (these replace all invalidated routes and, obviously, avoid the failed component) and advertises them to its neighbors. Upon receipt of the Permanent Withdrawal message from the monitoring node, all routers in the Internet remove, from their routing tables, all routes that pass through the failed component and immediately choose the next best available path vectors from their backup routing tables. Some of the chosen routes may, of course, be later replaced by better routes that might be received from the neighbors. Although the monitoring router announces locally optimum replacement routes for the possible benefit of its neighbors, the latter (as well as their neighbors, and so on) are obviously free to choose some, all or none of them.

It should be noted from the above that because the failure is detected by a router locally and reliably, and no invalid routes are propagated by any router in the Internet, the path exploration will be drastically reduced and the BGP will achieve a fast convergence. Detection of a failed component has been achieved by the novel concept of "special neighbors of a router in the network. This was initially developed in connection with studies on the count-to-infinity and slow convergence problem in distance vector (DV) routing [16][17] and was later applied in some preliminary work on BGP convergence [18][19]. Finally, 6 new routing control messages have been proposed to be incorporated in BGP, without, however, changing any of the existing message formats.

The paper has been organized into seven sections. Following this introductory section, we briefly review some related works in Section 2. Section 3 is devoted to discussing the MBGP basics. The simplified model of the BGP and of an AS network that have been assumed for the present study is presented in Section 4. Section 5 describes the detailed working of the MBGP. Description of the simulation procedure and comparative results of BGP and MBGP have been provided in Section 6. Finally, some concluding remarks have been made in Section 7.

2 Related Works

In their pioneering work, Labovitz and others [2] - [4] showed, through experimental measurements, that the Internet may take a large time, even on the order of tens of minutes, to get back to its stable state operation after a fault has occurred. They observed that the BGP path selection process on the Internet backbone routers mainly caused this delay and the end-to-end internal paths suffered intermittent loss of connectivity, increased packet loss and large latency during this delayed convergence of BGP. Vendor-specific router implementation decisions and ambiguities in BGP specifications[1] were demonstrated as the main reasons for convergence delay [3],[5],[6],[9].

Some studies on BGP convergence problem and its solution were made in [7] - [9] but the suggested ideas were not much practical. An important new direction towards solving the route instability and delayed convergence problem in BGP emerged with the realization that the best way to reduce path exploration is to determine its root cause and then notify the affected routers about it [10]-[12]. However, two unwelcome features in [10]-[12] are the need for modification of the BGP update message format and the considerable processing and memory overhead of the notified routers. Finally, a few papers like [13]-[15] have concentrated on only identifying the root cause of route changes. Unfortunately, the proposed methods are fairly complex and do not appear to be much practicable.

In the remaining portion of this paper, we shall describe the various aspects of our proposed MBGP algorithm including the broad philosophy, the fault sensing mechanism, the various simplifications, assumptions and modifications related to the BGP and the simulation procedure with results for BGP and MBGP. The overall strategy adopted in our method may be broadly described as "distributed fault detection, notification, and rediscovery-cum-readvertisement of alternative valid routes" and it incorporates some insights gained from published research.

3 MBGP Philosophy and Background

3.1 Broad Philosophy of MBGP

A philosophical thought that lies behind our proposed approach towards reducing path exploration in the Internet can be explained with an analogy. We imagine the Internet in its "stable condition" as a vast pool of "calm water". Occurrence of a "component failure" in the Internet which can occur at any time and anywhere is analogous to a "random stone throw" into the vast pool of "calm water". The resultant disturbance in the body of water generates ripples moving in all directions from the "point of disturbance" which is analogous to the "physical location of the failed component". The resultant (radial) movement of ripples may be likened to the "path exploration" phenomenon in the Internet. At the end of the path exploration process, the BGP finally "converges", i.e., the pool of water "again becomes calm". Obviously, a small ripple would die down quickly, disturbing only a small area, whereas a big ripple would remain active for a long time and would disturb a large area. In a similar manner, in the present Internet, some faults cause the path exploration process to last a short duration and result in the exchange of a small number of invalid messages; other faults cause long path exploration, resulting in the exchange of large number of invalid and valid messages before the BGP converges.

Continuing with our above analogy, we endeavor, in the proposed method, to sense any incidence of "random stone throw" as close to its point of occurrence (both in time and in place) as possible and, thereafter, take remedial measures to control the the resultant ripple movement. This would make the disturbed pool of water become calm again with a minimum delay and (as a consequence) with minimum spread of the ripple movement. In order to realize, in practice, this goal of having reduced path exploration, we endow each router in the Internet with some additional intelligence. This allows the router to periodically monitor its neighborhood for sensing the failure of any neighboring component, locate or identify the failed component, notify the neighboring routers

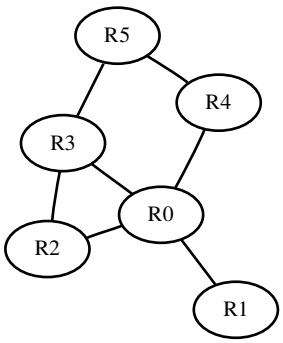

Fig. 1. A network graph to illustrate the four special neighbors.

about the failed component determine (if possible) an alternative valid route that avoids the failed component and, finally, advertise the alternative valid routes to its neighbors. Of course, the BGP neighbors receiving these updates are free to ignore them if they have better routes available. However, it is most important to note that only valid routes are propagated in the Internet to ensure that BGP convergence is achieved much faster.

3.2 Special Neighbors and Their Utilization

Concept of several types of special neighbors (SN) of a router were introduced and utilized in [16]-[19] in connection with studies on DV routing protocol and BGP. In this subsection, we describe four types of SNs which have enabled a BGP router to detect a faulty component in its immediate neighborhood and take appropriate measures towards reducing path exploration in MBGP. We shall use the simple network graph of Fig. 1 to illustrate these SNs neighbors.

1. Singly-Connected Neighbor (SCN): A neighboring router R_y is a SCN of the router R_x if R_x is the sole neighbor of R_y so that R_y can communicate with all other routers in the network only via R_x . In Fig. 1, R1 is a SCN of R0 and is a pendant node in the network. It is obvious that in case of failure of the router R_y or link R_x R_y, R_x can declare R_y to be a Lost Destination(LD) to all routers in the network.
2. Multi-Connected Neighbor (MCN): If a neighboring router R_y of the router R_x is not its SCN, then R_y is a MCN of R_x. In Fig. 1, all neighbors of R_0, except R_1 are its MCNs. It is obvious that in case a router R_x loses its communication with its MCN R_y, because of the failure of the connecting link R_xR_y, then R_x can still communicate with R_y, although in an indirect manner.
3. Co-Neighbor(CN) or Triangle Neighbor (TN): If the MCN neighbor R_y of the router R_x is also a neighbor of another MCN neighbor R_z of R_x, i.e., if R_x, R_y and R_z form a triangle in the network graph and are all mutual neighbors of one another, then R_y is a CN of R_x for R_z and similarly, R_z is a CN of R_x for R_y. In Fig. 1, R_2 and R_3 are CNs of R_0 for R_3 and R_2, respectively. It is obvious that in case the router R_x loses its communication with a neighbor R_y (R_z), where $R_y(R_z)$ is a

CN of R_x for $R_z(R_y)$, then R_x can utilize $R_z(R_y)$ for easily ascertaining whether the link $R_xR_y(R_xR_z)$ or the router $R_y(R_z)$ has failed.

4. Quadrilateral Neighbor (QN): If a router R_x has two MCN neighbors R_y and R_z who have a common neighbor R_w, who is not a neighbour of R_x, i.e, the four routers R_x,R_y,R_z and R_w together form a quadrilateral, then R_y is a QN of R_x for R_z and, similarly, R_z is a QN of R_x for R_y. In Figure 1, R_0, R_3, R_5 and R_4 form a quadrilateral and $R_3(R_4)$ is a QN of R_0 for $R_4(R_3)$. It may be observed that in case the router R_x loses its communication with the QN $R_y(R_z)$, it can still send a message to $R_y(R_z)$ via $R_z(R_y)$ to ascertain whether the link $R_xR_y(R_xR_z)$ or the router $R_y(R_z)$ has failed. One important point that needs to be noted regarding the utilization of a CN and a QN in BGP is that the policies of the concerned routers should not stand in the way of utilizing these special neighbors. MBGP utilizes the above four categories of special neighbors to great advantage as will be described in section 5.

4 Simplified Model of BGP and Some Assumptions

Both the BGP and the Internet architecture are highly complex. In order to study the proposed modification in the BGP, we have assumed a simplified view of the global Internet as an Interconnection of N ASes where each AS has a single BGP speaking router connected to multiple independent IP internetworks, each via a dedicated link to a non-BGP gateway router connected to the internal network. A BGP router within each AS thus peers with one or more BGP speaking routers in other ASes and several non-BGP routers within its own AS, as shown in Figure 2. The nine ASes, AS_0 through AS_8, have their respective BGP routers R_0 through R_8 and their respective pairs of non-BGP internal routers (R_{00},R_{01}) through (R_{80},R_{81}). Only two representative internal networks have been shown within AS_8.

Some of the ASes have been assumed to be stub ASes while the others have been assumed to be transit ASes, there being no multihoming. There is only one stub AS, namely, AS_7, in Figure 2. For simplicity, we have assumed that the transit ASes do not provide geographical store-forwarding of packets but provide store-forward of packets for remote ASes via only e-BGP links.

We assume that the following component failures can occur in the simplified model of the Internet.

- An e-BGP link connecting two neighboring BGP routers
- A BGP router
- A link connecting a BGP router to one of its gateway routers
- A gateway router
- An internal network to the link connecting a gateway router to it.

However, in the context of the above possible faults, we shall make the fairly reasonable assumption that only one fault can occur at a time.

Next, we assume that the two BGP routers sharing each e-BGP link maintain a reliable TCP connection over the link and periodically exchange KEEPALIVE messages. Similarly, within each AS, the BGP router and the gateway router communicate using

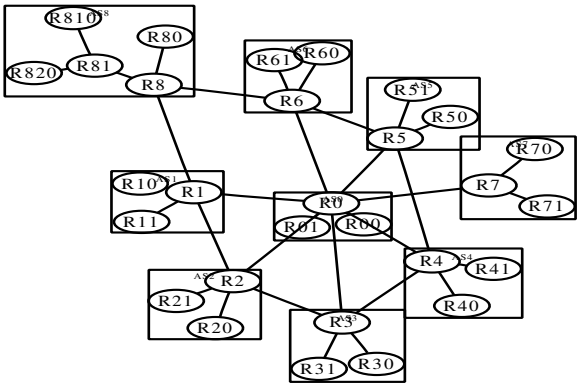

Fig. 2. Simplified model of the Internet adopted for our study.

an intra-AS protocol like RIP-2 and exchange periodic updates of their routing tables. Thus in case either a neighboring BGP router or a gateway router fails or the connecting link fails, a BGP router will detect the failure within a limited delay by the absence of the expected KEEPALIVE or periodic update message. However, if a link connecting a gateway router to its internal network fails, the gateway router will let the BGP router within the AS know about the failure, again within a limited delay.

Finally, we make the following two assumptions regarding the tables that each BGP router maintains.

1. We assume that each router, after joining the Internet, receives the Path Vector Advertisement Table (PVAT) of each of its neighbors and stores them in a Composite Path Vector Table (CMP_PVT) which thus contains multiple routes to reach each destination network. From its CMP_PVT, each router selects its best route to reach each destination, stores all these best routes in its own PVAT and advertises them to all its neighbors. In the absence of policies, each BGP router chooses its best route to each destination as the one that has the minimum number of hops in its AS-PATH vectors.
2. In order to keep itself prepared to deal with any link or router failure that may occur any time in its immediate neighborhood, each BGP router maintains two neighbor-related tables. The first one is a Neighbor Particulars Table (NPT) in which are stored particulars like AS#, IP address, prefixed advertised, etc. The second one is a Special Neighbor Status Table (SNST) in which are stored information about each neighboring router like whether it is a SCN or MCN and who are its CNs and QNs, etc. Each BGP router builds up its NPT and SNST from its CMP-PVT.

5 Modifications Incorporated in BGP

In this section, we shall outline the steps followed by the MBGP router R_x to achieve greatly reduced path exploration following the detection of failure of any of the five components listed in section 4.

5.1 Failure of a BGP Router

1. If a neighboring BGP router R_y fails, R_x does not receive the KEEPALIVE message from R_y before the hold timer for R_y times out and hence it detects the failure of either R_y or the link R_xR_y.
2. R_x now checks whether its TCP connection with R_y has been broken or reset by attempting to open a TCP connection with R_y afresh. Obviously, the attempt fails in this case.
3. R_x then consults its SNST to know whether R_y is its SCN or MCN. If R_y is found to be a SCN, then R_x simply announces a "permanent withdrawal" of the router R_y to all BGP routers in the Internet by broadcasting (through flooding), a PERMANENT SCN-ROUTER WITHDRAWAL (R_y) message over the entire Internet. On receipt of this message, all BGP routers just delete all routes advertised by R_y from their respective CMP_PVTs and PVTs; no alternative routes need be discovered.
4. In step 3 above, if R_y is found to be a MCN, then R_x needs to ascertain whether R_y itself or the connecting link R_xR_y has failed. Towards this, R_x checks its SNST to know whether it has one or more CNs or QNs or both for R_y. If yes, then R_x sends one or more ROUTER-FAIL CHECK messages to R_y via these CNs and QNs. However, since no ROUTER-OK (R_y) response comes back, R_x learns that the router R_y has failed. Then R_x first broadcasts a PERMANENT MCN-ROUTER WITHDRAWAL (R_y) message to declare the MCN R_y an LD. This results in (i) permanent removal of all routes stored in all BGP routers in the Internet that were originated by R_y and (ii) temporary withdrawal of those routes which only passed by R_y.
5. Immediately thereafter, R_x (as well as other routers which had temporarily withdrawn all routes which passes by R_y) tries to discover alternative routes and advertises them to their neighbors. This initiates some amount of path exploration, but best alternative routes are soon found out for the temporary withdrawn destinations. The following points may be noted in the present context.
   (a) In case R_x finds from its SNST that no CN or QN exists for R_y, it searches its CMP_PVT for knowing if any neighbor of it had advertised to it any path originating from R_y and then sends a ROUTER-FAIL CHECK message to R_y via this path. Actually, the SNST only provides some shorter paths and that too readily.
   (b) Sending multiple ROUTER-FAIL CHECK messages via possibly independent paths, if available, increases the reliability of the router/link failure checking process.
   (c) Each neighbor of R_x, after receiving the updated (new) routes from R_x, are free to accept or ignore them in case they themselves have better (shortr) routes stored in their CMP_PVTs.

5.2 Failure of an e-BGP Link

Let us assume that in step 4 in section 5.1, R_x receives a ROUTER-OK response from R_y against ROUTER-FAIL CHECK probe message sent by it. As a result, R_x learns that the link R_xR_y has failed and, consequently, R_y is no longer its neighbor. So, R_x

first removes the entry of R_y from its NPT and SNST and then broadcasts a PERMANENT LINK WITHDRAWAL (link-id) message over the entire Internet. Next R_x removes the set of routes, that were advertised by R_y as well as the subset of these routes that R_x had thereafter propagated to its other neighbors, from its CMP_PVT and its PVT, respectively. Then R_x discovers alternative routes to those destinations(avoiding the failed link) and advertises them. Thus each BGP router in the Internet receives the PERMANENT LINK WITHDRAWAL (link_id) message from R_x, immediately followed by the BGP UPDATE message(s) sent by R_x. In between the two messages, each BGP router discovers and uses alternative routes, although these routes may soon be replaced by better routes.

5.3 Failure of Components Within an AS

From the simplified model of the Internet shown in Figure 2, it is evident that, within an AS, three types of components may fail, namely, a gateway router, the link connecting it to the BGP router and, finally, an internal network or the link connecting it to the gateway router. The BGP router can detect the failure of the gateway router or the connecting link by the non-receipt of the DV table from the gateway router and the failure of the network from the content of the DV table received from the gateway router. In case of any failure within its AS, the BGP router thus simply uses the concept of a SCN and broadcasts a PERMANENT NETWORK WITHDRAWAL (network prefix) message.

5.4 New Messages Use the Existing BGP Message Format

BGP uses only 4 types of messages, namely, OPEN, UPDATE, NOTIFICATION and KEEPALIVE [1]. All BGP messages have a common 19-byte header followed by separate or special format for each message type, with the exception of KEEPALIVE which is just the 19-byte header containing no information. The header has a 16-byte MARKER field, a 2-byte LENGTH field and a 1-byte TYPE field. Presently, only 4 values, viz, 1,2,3 and 4 have been assigned to the TYPE field to identify the OPEN, UPDATE, NOTIFICATION and KEEPALIVE messages, respectively. Thus, it is possible to use the TYPE field in the BGP header to create the new routing control messages needed by MBGP. MBGP needs 6 additional messages, namely, PERMANENT SCN-ROUTER WITHDRAWAL (router-id), PERMANENT MCN-ROUTER WITHDRAWAL (router-id), PERMANENT LINK WITHDRAWAL (link-id),PERMANENT NETWORK WITHDRAWAL (network prefix(es)), ROUTER-FAIL CHECK (router-id) and ROUTER OK (router-id). The value of the TYPE field and the format of the respective attributes may be assigned following the convention used in the design of the BGP message format.

6 Simulation Procedure and Results

Though RFC 4271 [1] describes the BGP in details, it does not contain much idea about its implementation. As a result, most router vendors have come up with their own

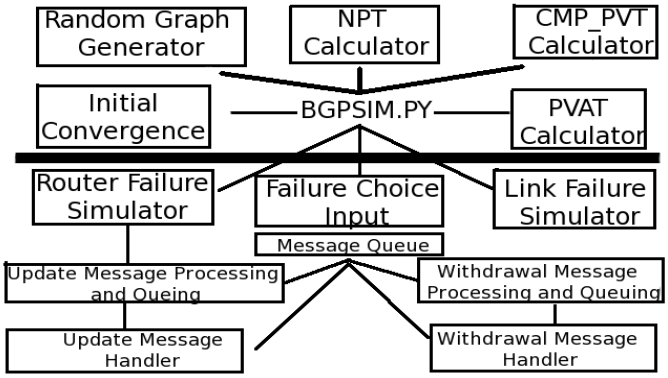

Fig. 3. Simulation Model of BGP 4.0

implementations details of which are, unfortunately, not available in the public domain. So, for simulating the process of convergence in BGP and MBGP, we have employed the simplified model of BGP described in section IV and simulated it, leaving out the internal networks within the ASes, using the simulator module shown in figure 3.

The main routine or the simulator program bgpsim.py in Fig 3, in association with the different subroutines, first builds the various tables from a given graph of ASes (in the absence of internal networks, an AS is reduced to just a BGP or MBGP router) which is fed as its input, either manually or from a random graph generation subroutine. It exchanges the initial messages, on behalf of the nodes, till the BGP converges into a steady state, thereby simulating the nodes in the graph booting up and exchanging messages till the network stabilizes. The simulator then injects a random failure of a link or router in accordance with the user's choice and simulates and prints the exchange of messages between the ASes till the network stabilizes again.The different subroutines that have been used in the simulation of the BGP are shown in Fig 3. Similarity in the basic design and the method of simulation has allowed reuse of codes and flowcharts of BGP while simulating MBGP. Only two new subroutines, namely, Router and Link Failure Simulator and the Permanent Withdrawal Message Handler, needed to be written for the MBGP.

The process of convergence in BGP and MBGP has been studied on four different network graphs with two router failures and two link failures, all chosen randomly, being successively injected in each network. The number of messages that were exchanged during the process of convergence were counted in each case to obtain an idea of how fast the MBGP converges relative to the BGP, following identical failures injected on identical network graphs. The four network graphs, along with the four failed components in each graph are shown in Fig. 4 through Fig. 7, each figure being accompanied by the number of messages exchanged till the BGP and the MBGP converge. The results clearly demonstrate that the MBGP has a significantly reduced path exploration compared to BGP and, as a consequence, it converges much faster.

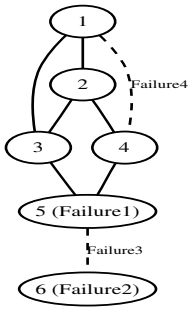

Failure #	BGP msgs	MBGP msgs	% impr
1	90	33	63 %
2	63	18	71 %
3	68	18	74 %
4	12	9	25 %

Fig. 4. Network graph 1 and results for the failures shown.

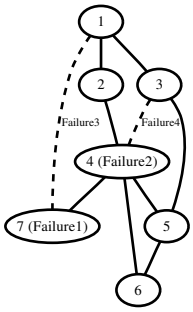

Failure #	BGP msgs	MBGP msgs	% impr
1	165	27	84 %
2	90	72	20 %
3	17	10	41 %
4	13	10	30 %

Fig. 5. Network graph 2 and results for the failures shown.

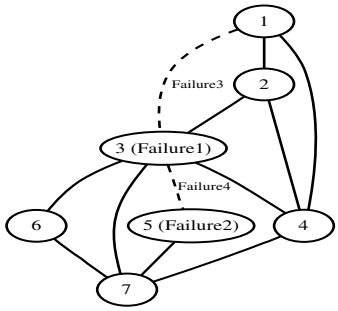

Failure #	BGP msgs	MBGP msgs	% impr
1	325	114	65 %
2	662	29	95.6 %
3	28	9	68 %

Fig. 6. Network graph 3 and results for the failures shown.

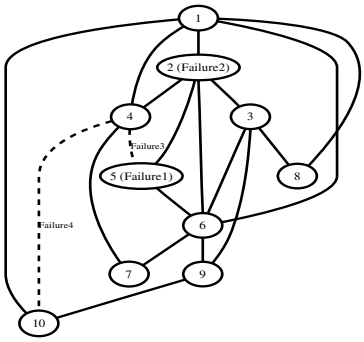

Failure #	BGP msgs	MBGP msgs	% impr
1	3256	41	98.7 %
2	63	18	71 %
3	68	18	74 %
4	12	9	25 %

Fig. 7. Network graph 4 and results for the failures shown.

7 Concluding Remarks

Some modifications have been proposed to BGP, in the form of a Modified BGP (MBGP), to enable it to converge much faster following a link, router or internal network failure. This will reduce packet losses, router congestion, increased packet delay and other deleterious effects that occur in the Internet during the delayed convergence of BGP. The MBGP adopts the overall strategy of "distributed fault detection, fault notification and rediscovery-cum-readvertisement of alternative valid routes". A novel concept of special neighbors in conjunction with a new query-response message pair, enable each MBGP router to detect a component failure and identify the failed component locally, quickly, reliably and with negligible overhead. All MBGP routers in the Internet are immediately notified by the fault-detecting router, through flooded broadcast of one of four new permanent withdrawal messages, about the identity of the failed component. Two of these messages which broadcast the permanent withdrawal of a singly connected router and of a singly connected network, result in immediate MBGP convergence with no path exploration. The other two messages, which broadcast the permanent withdrawal of a link and of a multiconnected router, start a path exploration which, however, dies down quickly because no invalid routes are ever generated in MBGP and very few downstream routers actually switch routes and fewer still do it multiple times. Simulation studies have demonstrated a significantly faster convergence of MBGP over BGP. With its distributed fault identification and concomitant fast converging capabilities, the MBGP as well as similar future algorithms will have the potentiality to make the global Internet "the largest self-regulating engineering system in the world".

References

1. RFC4271. www.ietf.org/rfc/rfc4271.txt
2. C. Labovitz, A. Ahuja, A. Bose and F. Jahanian: Delayed Internet Routing Convergence. SIGCOMM 2000
3. C. Labovitz, R. Wattenhofer, S. Venkatachary, and A. Ahuja: The Impact of Internet Policy and Topology On Delayed Routing Convergence. INFOCOM, April 2001
4. Craig Labovitz, Abha Ahuja, Abhijit Bose, Farnam Jahanian: Delayed Internet Routing Convergence. IEEE/ACM Trans. Network. pp 293-306, 2001
5. T. Griffin and B. Presmore: An Experimental Analysis of BGP Convergence Time. IEEE ICNP, November 2001
6. Z. M. Mao, R. Govindan, G. Varghese, and R. Katz: Route Flap Damping Exacerbates Internet Routing Convergence. ACM SIGCOMM, 2002
7. T. Griffin and F. Shepherd: The Stable Path Problem and the Interdomain Routing. IEEE ICNP, November 2001
8. Dan Pei, Xiaoliang Zhao, Lan Wang, Massey, D.Mankin, A. Su, S.F. Lixia Zhang: Improving Bgp Convergence Through Consistency Assertions. INFOCOM 2002
9. A. Bremler-Barr, Y. Afek and S. Schwarz: Improved BGP Convergence via Ghost Flushing. IEEE INFOCOM, 2003
10. D. Pei, M. Azuma, D. Massey, and L. Zhang: BGP-RCN: Improving BGP Convergence Through Root Cause Notification. Computer Networks, Volume 48, Issue 2, pp. 175-194, 2005
11. Hongwei Zhang, Anish Arora, Zhijun Liu: A Stability-Oriented Approach to Improving BGP Convergence. IEEE International Symposium on Reliable Distributed Systems, 2004
12. Chandrashekar, Z. Duan, Z.-L. Zhang, and J. Krasky: Limiting Path Exploration in Bgp. INFCOM, Miami, USA, 2005
13. Anja Feldmann, Olaf Maennel, Z. Morley Mao, Arthur Berger, Bruce Maggs: Locating Internet Routing Instabilities. ACM SIGCOMM Computer Communication Review archive Volume 34, Issue 4, October 2004
14. M. Lad, Akash Nanavati, D. Massey, L. Zhang: An Algorithmic Approach to Identifying Link Failures. IEEE PISPDC, pp. 25-34, 2004
15. R. Teixeira and J. Rexford: A Measurement Framework for Pin-pointing Routing Changes. IEEE SIGCOMM '04 workshop, 2004
16. S. K. Ray, S. K. Paira and S. K. Sen: Modified Distance Vector Routing Avoids the Count-to-Infinity Problem. Proc. Intl. Conf. on Commun. devices and Intell. Systems (CODIS 2004) held in Calcutta during Jan 8-10, 2004, pp. 31-34
17. Santanu Kr Sen: An Improved Network Routing Scheme Based on Distance Vector Routing. Ph.D (Engg.) Thesis of Jadavpur University, January 2008
18. S. K. Ray, P. Ghosh and S. Sen: Modified BGP has Faster Convergence. proc. NCC 2006 held at I.I.T Delhi during Jan 27-29, 2006, pp. 404-408
19. Susmit Shannigrahi: Reducing path Exploration in BGP. M.C.S.E Thesis of Jadavpur University, June 2009