

Characterization of BGP Recovery Time under Large-Scale Failures

A. Sahoo
Univ. of California, Davis

K. Kant
Intel Corporation OR

P. Mohapatra
Univ. of California, Davis

Abstract—Border gateway protocol (BGP) is the standard routing protocol between various autonomous systems (AS) in the Internet. In the event of a failure, BGP may repeatedly withdraw some routes & advertise new ones until a stable state is reached. It is known that the corresponding *recovery time* could stretch into hundreds of seconds or more for isolated Internet outages and lead to high packet drop rates. In this paper we characterize BGP recovery time under large-scale failure scenarios, perhaps those caused by disastrous natural or man-made events. We show that the recovery time depends on a variety of topological parameters and can be substantial for massive failures. The study provides guidelines on reducing the impact of BGP convergence delay on the Internet.

Keywords: Autonomous system (AS), Border gateway protocol (BGP), Massive failures, Recovery time.

I. INTRODUCTION

BGP (Border Gateway Protocol) is the predominant protocol used for inter-domain routing in the Internet. BGP belongs to the class of *path vector* routing protocols wherein each node maintains multiple ordered paths to reach each destination. One of these paths is chosen at any time according to some given policy. When this primary path fails, BGP withdraws this path and selects the next best backup route. The new route is advertised to its neighbors. However there is no guarantee that the backup route is still valid. In case the backup route has failed, it will be withdrawn only after a withdrawal is sent by the neighbor which advertised it, and another backup route is chosen. Thus because in the absence of information about the validity of a route that can cause BGP to go through a number of backup routes before selecting a valid one. This can result in a considerable delay before the cycle of withdraws/advertisements ends and all BGP nodes have a valid path the destination.

Internet routing sports other classes of routing protocols as well, such as the *distance vector* and *link state* protocols. Distance vector protocols advertise the cost of the best path for a destination, to their neighbors whereas link state protocols flood the entire network with information about the cost to reach their immediate neighbors. Distance vector algorithms usually suffer from the *count-to-infinity* problem in that the nodes may continuously increase their cost to reach a failed node, leading to lost connectivity and loops. The flooding approach of link state protocols makes them generally inappropriate for inter-AS use. Distance vector algorithms are generally used within an autonomous system (AS). In contrast, inter-AS routing primarily uses BGP because of its better scalability, flexibility and configurability. In particular, the

scalability of BGP has been a critical factor in the explosive growth of the Internet over the last decade.

Numerous studies [7], [8], [4], [5], [12], [15] have been carried out to study the fault tolerance and recovery characteristics of BGP. In particular, it was shown by Labovitz et al. [8] that the convergence delay for isolated route withdrawals can take > 3 min in 30% of the cases. It is also shown in [19] that during recovery, packet loss rate can increase by 30x and packet delay by 4x. There has been some effort to create analytical models for the BGP convergence delay. These studies have identified factors that affect the recovery time and also computed lower and upper bounds. However it is still difficult to estimate the convergence delays for a fault in an arbitrary network. The problem is complicated further if we consider multiple failures and there has not been any work that studies this behavior.

The primary reason why large scale failures in the Internet have not been studied is their low probability of occurrence. But it is easy to see that large scale failures can cause a significant disruption to the Internet routing infrastructure, not only in the affected ASes but also in the rest of the Internet. Recent events have shown that communication networks are needed the most at the time of crisis, and that increases the importance of a short convergence delay. Therefore it is vital that we have a good understanding of the behavior of the Internet connectivity after a large scale failure. Large scale failures can occur because of a number of reasons such as malicious attacks on the Internet infrastructure, earthquakes, major power outages, massive hurricanes, etc. A large scale failure would typically take out numerous routers belonging to multiple Autonomous systems (ASes). This will obviously cause the routes to the failed routers to be withdrawn (Tdown [8]). Another consequence will be the withdrawal of routes passing through the affected routers and advertisement of backup routes in place of the failed routes (Tlong [8]).

It is more likely that there is a contiguous area of complete failure, scenarios where the affected routers are sparsely distributed over a large area can also be envisioned. In either case, scenarios where only the links (but not the routers) fail are not very likely in a massive failure situation, and are not considered here. The aim of this work is to study the recovery characteristics of BGP networks after multiple BGP router failures and to identify the factors that affect the convergence process.

A. Related Work

There has been a fair amount of work on the analysis of BGP convergence properties. However, most publications have examined only simple networks or a specific set of sources and destinations. Although many parameters affecting the convergence time have been identified, it is still not possible to estimate the convergence time for a set of simultaneous failures in an arbitrary network. In this section we talk about the important papers published in this area and the conclusions therein.

One of the first papers that tried to analyze the BGP convergence delay was by Craig Labovitz et al [8] where the authors injected faults in the Internet and measured the convergence times from 5 different ASes. The authors observed different convergence times for the same event from different ASes. The authors also computed the lower and upper bounds for the convergence time for a completely connected graph. In a follow-up paper Labovitz et al.[9] concluded that the convergence time for a route is proportional to the length of the longest possible backup path from the source to the destination. Obradovic [12] also found the convergence time to be dependent on the path length for path vector protocols that use shortest path first strategies.

Griffin and Premore [5] studied the effect of BGP's MRAI (*minimum route advertisement interval*) timer on the convergence time after a fault in simple BGP networks. They found that as the value of the MRAI timer is increased, the convergence time first goes down and then increases. The number of update messages however, stabilizes after decreasing initially. The authors concluded that there is an optimum value for the MRAI timer for each source destination pair and a fixed value for the MRAI timer is not optimal in terms of the convergence time.

The rest of the paper is organized as follows. Section I-A briefly discusses the related work. Section II discusses our study methodology and the assumed network characteristics. Section III discusses the results.

II. STUDY APPROACH

We used a number of synthesized topologies for our studies and varied their parameters to analyze the effect of these parameters on the recovery times. The basic topology generator tool used was BRITE [11] and BGP simulations were carried out using the SSFNet *citessfnet* simulator.

A. Topology Generation

BRITE can generate topologies with a configurable number of ASes and with multiple routers in each AS. BRITE supports a number of AS topology generation schemes such as Waxman [17], Albert-Barabasi [1], and GLP [2]. In the Waxman scheme, the probability of two ASes being connected is proportional to the negative exponential function of distance between the two ASes. The Albert-Barabasi and GLP models use preferential connectivity and incremental growth for edge creation. In these schemes the probability of connecting to a node is proportional to the degree of that node. Both these schemes try to generate a power-law degree distribution. We

modified BRITE to allow more flexible degree distribution so that it is possible to assess the impact of degree on recovery time in more controlled settings as well (e.g., uniform degree, mixture of high and low degree, etc.). We also modified the code to generate variable number of routers for the ASes. The number of routers for each AS was selected from a heavy tailed distribution.

Geographical placement is essential for studying large scale failures since such failures are mostly expected to be geographically contiguous (e.g., earthquake zone). However, directly using the geography of actual Internet is not only difficult (precise identification & location of routers is a hard problem) but also considerably limits the scenarios that can be studied. Instead, we placed all ASes and their routers on 1000x1000 grid, where the physical size of the grid is chosen so that the average router density (#routers per square mile) is roughly the same for the US based Internet. For reasons of simulation complexity, we do not, however, consider physical area of the size of the entire US.

Studies of real internet have found that the geographical extent of an AS is strongly correlated to the AS size (i.e., number of routers in the AS) [10]. Here we assume a perfect correlation and make the geographical area of an AS (in terms of number of squares occupied on the grid) proportional to its size (number of routers). In particular, the routers of the largest AS are distributed over the entire grid. For smaller ASes, the area is reduced proportionately, and its extent (also a square) is placed randomly on the grid subject to the limitation that it doesn't go beyond the overall grid. The routers of an AS are distributed randomly over the geographical area assigned to it. In most cases, we used a uniform distribution; however, we also studied the impact of clustered router placement on recovery time.

Internet studies also show that larger ASes are better connected [16]. This is handled as follows: We first create a sequence of AS degree values according to the selected AS degree distribution and sort them. Similarly, the AS list is also sorted according to the number of routers in the ASes. The inter-AS degree of an AS in the AS list is then set to the value at the corresponding location in the inter-AS degree list. This creates a perfect correlation between AS sizes and degree. Again, although a perfect correlation is unlikely in practice, it is a reasonable approximation for our purposes.

Normally we did not take geographical location into account when creating inter-AS edges, but we did run a few cases where we used a Waxman (distance-based) connectivity function. The ASes are linked together using a pseudo-preferential connectivity in the sense that one of the ends of an edge is selected randomly but the other end is selected according to the degree of the node. Once an inter-AS edge has been created, we randomly select a router from one of the ASes and connected it to the closest router in the other AS. We used the default Waxman scheme for creating the inter-AS edges. However we observed that distance based connections inside the ASes did not have any significant impact on the convergence delays.

With the above changes, our modified BRITE can generate networks with arbitrary degree distribution for the constituent

ASes. This allowed us to experiment with degree distributions with different decay characteristics, distributions extracted from real networks besides the uniform and constant degree distributions.

B. BGP Simulation

We used the SSFNet simulator for our experiments because it has been used extensively in the research community for large scale BGP simulations and BRITE can export topologies in the format supported by SSFNet. The interdomain routing protocol used was of course BGP and OSPFv2 was the intradomain routing protocol. We used the standard values for the MRAI timers, 30 seconds for eBGP and 0 seconds for iBGP. All the timers are jittered as specified in RFC 1771 [13] and are reduced by up to 25%. We used a mesh of iBGP peering instead of route reflection [6] inside the ASes as the number of routers in the ASes is not very large. For the experiments, we simultaneously failed a group of routers and then analyzed the update messages generated by BGP. In most of the cases we failed the routers in a circular region around the center of the map. A failure will lead to both Tdown [8] and Tlong [8] events. A number of routes will be withdrawn and a subset of these routes will be replaced by other routes. Some of these newly advertised routes might be invalid and would hence be replaced later routes. The total convergence delay depends on the number of Tdown and Tlong events, We observed the recovery time and messages sent out per unit time, for failures of different magnitudes.

III. EXPERIMENTAL RESULTS

In studying the impact of massive BGP router failures on recovery time, the following parameters are the most relevant:

- 1) Extent of failure, expressed as percentage of routers failed.
- 2) Impact of BGP router degree distribution (average degree & its variability).
- 3) Shape of the failure area – contiguous (“area failure”) vs. scattered. Since past work has examined scattered failures, we concentrate mainly on area failures.
- 4) Impact of MRAI timer setting.
- 5) Impact of distance distribution between routers.
- 6) Impact of non-uniformity in geographic location of routers.

In the following we study these aspects. All the results in the following assumed 200 ASes, unless stated otherwise. Each AS had between 1-100 routers drawn from a heavy tailed distribution with an average of about 5.

Our initial experiments indicated a considerable variability and complexity in BGP recovery time behavior. Consequently, in the following we vary only one parameter at a time and also consider several simple topologies in addition to those modeled after the real topologies.

A. Degree Distribution

In studying the impact of degree distribution, we first examined how the recovery time with a “realistic” degree

distribution would compare against one with a constant or uniform degree with the same average value. The average measured inter-AS degree from the Internet AS-level topology is about 8.0 [18]. However, the Internet has over 22000 ASes and the maximum inter-AS degree is in the thousands. For our 200 AS network we decided to restrict the maximum degree to 50 and used the degree distribution in the range 1-50. This gave us a degree distribution which decays as a power law with an exponent of about -1.9. The average degree was very close to 4. We found that, a topology with a constant inter-AS degree equal to 4 yielded a recovery time 5-6 times as high as the realistic case! This prompted us to examine the recovery time as a function of the degree distribution. We started off with topologies in which all the ASes have a constant inter-AS degree. To avoid contamination of results due to other factors, routers were located uniformly in this case and distance wasn’t considered while creating the inter-AS edges. Fig. 1 shows the recovery time as a function of extent of failure. (The graphs still contain some statistical variability and thus minor irregularities should not be considered significant.)

In all cases, the recover time increases initially with the failure area to some maximum value and then slowly rolls off. It is seen that a higher degree consistently increases the recovery time. The sharpness of initial increase also increases with the degree. This happens because, the number of possible backup paths goes up as the degree is increased. The recovery time rises initially because, a larger failure translates into more failed routes and more failed backup routes. However as the number of failed nodes continues to grow, the residual network gets smaller and a larger proportion of backup routes are invalidated quickly. This causes the decline in the convergence delay. It must be noted however that the number of failed routes must keep increasing with the size of the failure. We are only looking at the BGP convergence delay here. Common sense dictates that larger failures are less probable than smaller ones. Thus for a network like the Internet, only the left end of the graph might be of any realistic interest.

Two, rather surprising, results from these graphs are as follows. First, for a reasonable connectivity (recall that “realistic” average connectivity is 4), the recovery time shoots up close its maximum value at a much smaller failure percentage (10%) than one would have suspected. The practical implication of this that we don’t need truly large failures to experience a very high recovery time. The second surprising result is that average degree is not at all a reliable indication of the recovery time. This is illustrated in Fig 2, which compares the convergence delays for the “realistic” topology mentioned earlier, a topology constant inter-AS degree of 4.0 and a third topology (referred to as the 70-30 case henceforth) where 70% of the nodes have a low connectivity (1-3) and the other 30% have a single higher connectivity value (=9) such that the average degree is 4.0. It is seen that variable connectivity helps bring down the maximum recovery time considerably.

The effect of variability in the degree can also be seen in Fig 3. Here we compare the convergence delays for topologies with constant inter-AS, with topologies that have a uniform degree distribution but with the same average degree as the constant case. For the uniform case, the inter-AS degree is

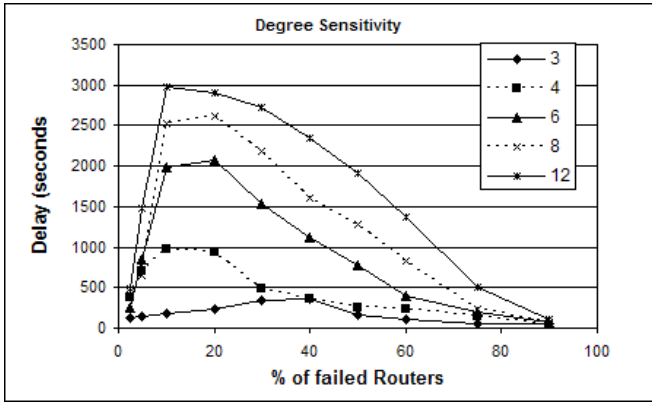


Fig. 1. Recovery time for constant degree networks

uniformly distributed in the range $[1..2x-1]$, where x is the desired average degree. Again we see that the convergence delays for the uniform case are significantly lower than constant case.

We have seen that the 70-30 and the uniform distributions yield lower recovery times than the constant connectivity case, and the convergence delays for the power law degree distribution are lower still. The reason for this behavior is that the overall recovery time is a result of two opposing factors wrt degree:

- A Number of routes: Higher degree nodes create many more routes, which means that during a failure, the number of withdrawn routes as well backup routes is higher.
- B Route Lengths: As the degree of a node increases it can, in general, reach other nodes in fewer hops. This helps in a quicker propagation of updates known to this node to other nodes. In other words, high degree nodes can act as “short circuits” and actually help lower the recovery time.

We find that, the effect (A) is generally much stronger than (B). As a result, a uniform increase in the degree of most nodes results in higher recovery times as shown in Fig 1. The effect (B) can be seen in Fig 2 where the convergence delay for the 70-30 case is less than the topology with constant inter-AS degree. Thus, the presence of a small percentage of high degree nodes can provide the beneficial short circuit effect and lower the recovery time. This can be seen more clearly in Fig 4 which shows the maximum recovery time as a function of the fraction of nodes that have a low degree. Recall that in Fig 2 we showed a situation where 70% of nodes have degree in the range 1-3 and others have a fixed degree (=9). In Fig 4, we use a similar idea except that percentage of nodes with low degree is varied. In all cases, we still maintain the same average degree. This means that as the fraction of high degree nodes decreases, their degree goes up. Fig. 4 shows the curve for average degrees of both 4.0 and 8.0. It is seen that the curves are almost monotonically decreasing (but for statistical fluctuations). This reinforces the idea that a small number of well connected nodes among a large number of poorly connected nodes forms the ideal situation for low recovery time.

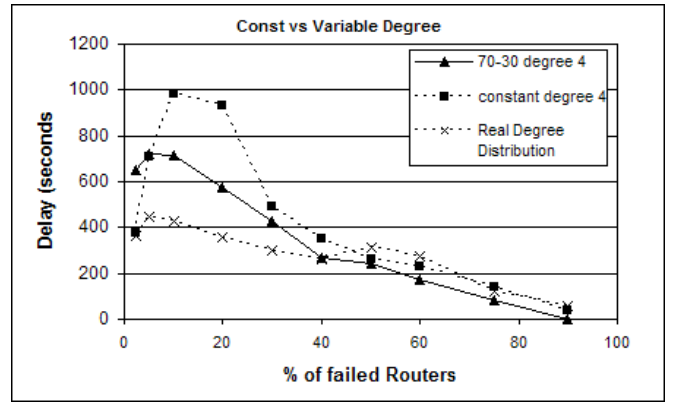


Fig. 2. Recovery time for different degree distributions

The arguments above still fail to explain why a distribution (e.g., power law) should yield lower recovery time than the fixed low-high mixture of degrees. This result follows by applying the above arguments recursively. Given the constant high degree regime, we can lower the recovery time by again splitting it into parts: a large subset with lower than average degree, and a smaller subset with higher degree. Note that a recursive high-low degree partitioning is akin to cascade multifractal construction and in the limit yields the log-normal distribution.

One issue not addressed above is the behavior of the convergence delay as a function of average degree (with the degree distribution kept the same). This is shown more clearly in Fig. 5 where we show the convergence delay of 10% failure for topologies with constant inter-AS degree. It is seen that the curve shows a diminishing return behavior, which may appear counter to the explanation of effect (A) above. The explanation for damping is that once the node degrees are already high, a node can reach others in very few hops. Thus, although the number of possible paths increases, the shortest paths grow even shorter which can limit the time needed for the failure notification to reach a source node.

B. Distance-based Connections

As stated earlier, in reality, routers connect preferentially to other routers that are nearby [10]. For small ASes, a similar property should hold with respect to AS-AS connectivity. For large, ASes, the concept of “nearby AS” may not be very meaningful; however, for uniformity, we conducted experiments with distance based inter-AS connectivity where the inter-AS distance was defined to be the distance between the “center” of the respective ASes. In our model the largest ASes cover almost the entire area of the map and hence their “location” will always be close to the center of the map. However, the heavy tailed distribution, which is used to generate the number of routers for each AS, ensures that the number of large ASes is small and hence the location is much more meaningful for the rest of the ASes. We used the Waxman connectivity scheme for creating the inter-AS edges. The probability that two ASes are connected was proportional to $e^{-d/\beta M}$ where d is the distance between the “locations” of

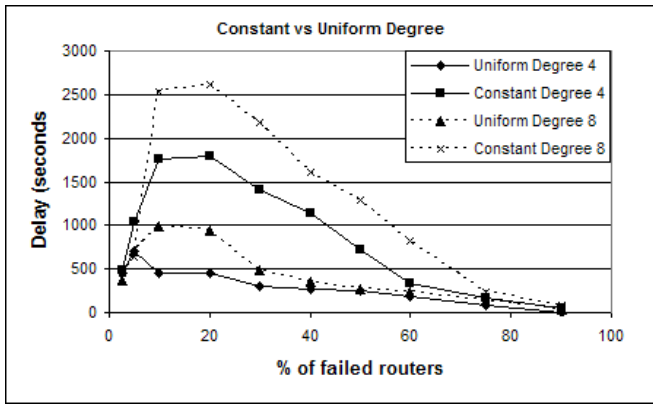


Fig. 3. Recovery time for Constant vs. Uniform degree

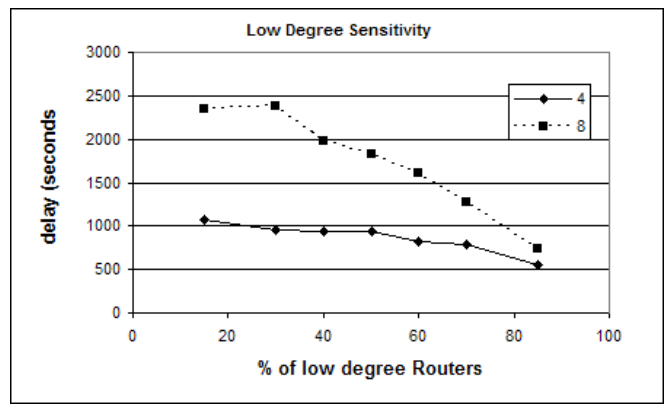


Fig. 4. Recovery time vs. percentage of low degree nodes

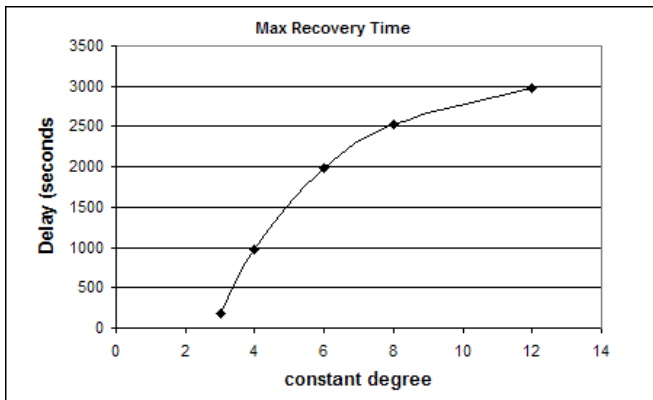


Fig. 5. Max Recovery time vs. constant degree for 10% failure

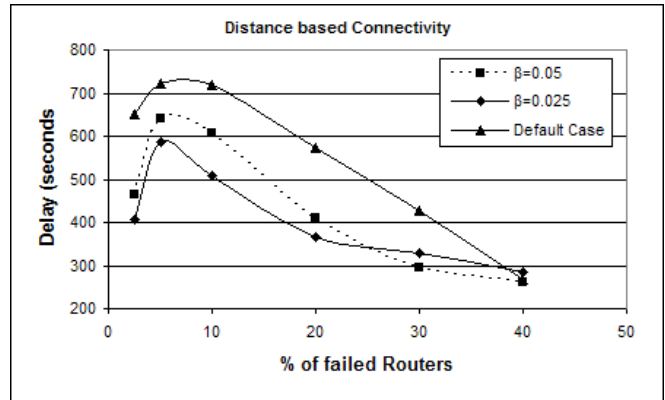


Fig. 6. Recovery time vs. distance based connectivity

the two ASes, and M is the maximum possible distance and β is a dimensionless parameter. For our experiments we varied the values of β and observed the variation in the convergence delay.

For all the cases we used the same degree distribution: 70% low degree(1-3) nodes and 30% high degree nodes. The overall inter-AS degree for the topologies was equal to 4. Fig 6 shows the results. It is clear that the convergence delay goes down as the decay rate β is decreased, i.e. as the probability of connecting to closer ASes is increased. The reason for the behavior is simple. A decrease in β leads to more links between geographically proximate ASes, and this means that these ASes now have less links connecting them to the rest of the network. The failure of a bunch of ASes in a contiguous area has less effect on the rest of the network, and hence the convergence delays go down.

C. Variable MRAI

The MRAI timer limits the number of updates sent from one AS to another during BGP convergence. The default value of the MRAI timer for external BGP connections is 30 seconds. A higher MRAI timer reduces the number of messages sent out after a failure since the information that becomes available later may make the earlier update superfluous. However, a larger MRAI also prevents quick dissemination of route changes to other nodes. Griffin and Premore citebgp-cnv shows

that there is an optimal value of the MRAI for the routes originating from each network. These results coupled with the probability of self synchronization in the network, leads us to speculate that a variable MRAI may work better than a fixed one. Fig 7 compares the effect of fixed MRAI against the one where the MRAI is uniformly distributed around the average value. In particular, we chose MRAI to be uniformly distributed in the range 2-58 secs. We can see that despite the jittered timers the convergence delay for variable MRAI is less than that for the fixed MRAI scenario. We used a 70-30 degree distribution with an average degree of 4 for both sets of runs.

We also show the number of messages generated during the convergence process in Fig 8. Each point represents the number of *update* messages generated during the preceding 10 second interval. We plot the time elapsed after the failure on the x-axis. Here we compare two specific cases with about the same convergence delay: one using a fixed MRAI timer and the other using a variable MRAI timer. We see that the number of messages generated for variable MRAI case is lower than that of the fixed MRAI case. We observe a similar behavior for other cases also. Thus we see that a variable MRAI is helpful in reducing both the convergence delay and the number of generated messages.

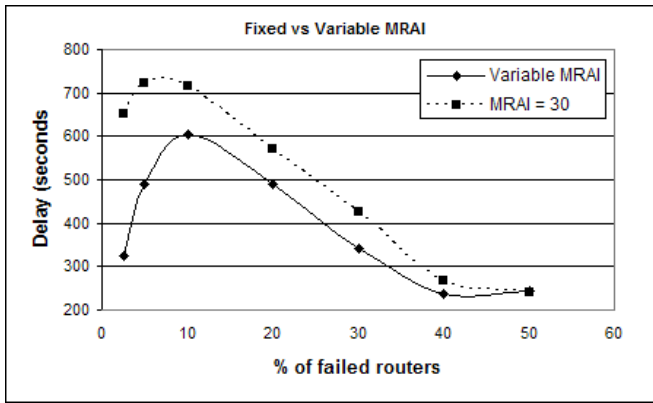


Fig. 7. Recovery time vs. MRAI distribution

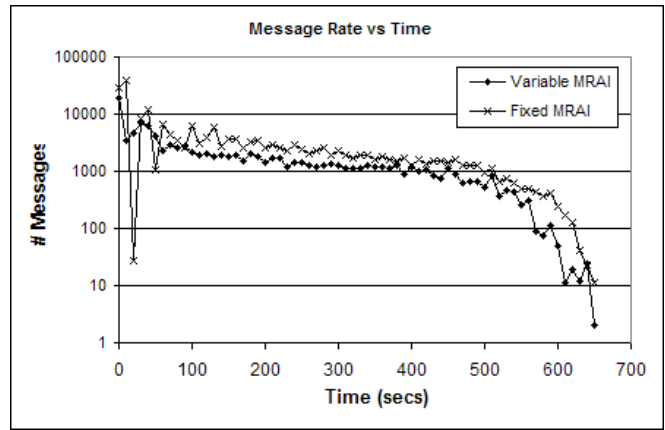


Fig. 8. Message count during convergence

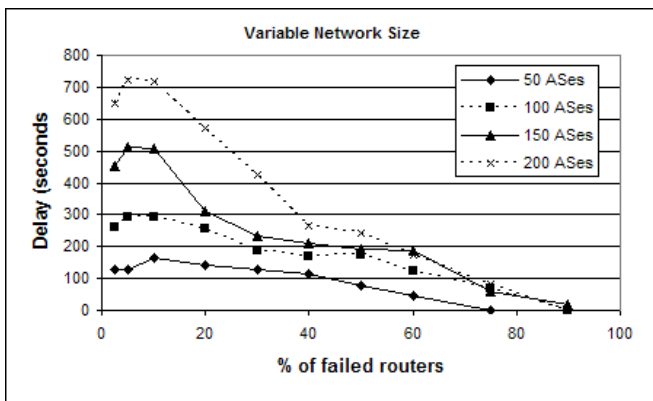


Fig. 9. Recovery time vs. AS size




Fig. 10. Effect of distributed failure

D. Network Size

In Fig 9 we show the effect of the size of the network on the convergence delay. As expected, we see that the convergence delay increases with the number of ASes in the network. That is because the number of routes goes up with the size. The interesting thing to note here is that the convergence delays go up even if the number of failed routers stays the same as the number of ASes grows. Thus for large networks, even moderate sized area failures could result in long recovery times. Given the continued growth of the Internet, we expect that BGP recovery times will continue to grow for expected failure sizes. This clearly points to the need for stop-gap mechanisms that can avoid substantial packet losses or route resolution errors during the recovery process.

E. Other Observations

In all the results that we have discussed till now, we considered a contiguous area of failure. However there can be scenarios in which the failed routers are sparsely distributed over a large area. Possible reasons could be a worm attack on the world wide web, an attack on routers sharing the same vulnerable software, etc. So we experimented with a few topologies in which the failed nodes were randomly distributed over the map. The results are shown in Fig. 10. We see that

the maximum convergence delays for the distributed failure are greater than that for the contiguous failure case. That is because, in a distributed failure that number of routers, for which an edge is failed, is much higher. In a contiguous a number of the failed edges are between the failed routers and hence do not have any effect on the convergence process. Thus the overall effect of a distributed failure is greater.

By default, the router placement in our experiments was uniform over the entire grid. We examined a few cases in which the distribution of the routers was non- uniform. For this, the entire grid was divided up into 5x5 blocks, and within each block a consistent non-uniform placement pattern was used. This pattern made the routers most likely to be located near the center and with decreasing probability towards the edges. No distance based connections were used in this case. It was found that non-uniform placement did not change the convergence time in any significant way.

IV. CONCLUSIONS

In this paper we studied the recovery time of BGP for large-scale failure scenarios. The study sheds light on how inter-domain routing of Internet will behave under natural or man-made disaster scenarios. It is found that the recovery time increases initially as the area of failure grows to about 10% and then rolls off. Furthermore, even with fixed number

of failed routers, the recovery time increases as the number of ASes increases. As indicated earlier, this has significant implications for the availability of the Internet. The paper also points to a number of other interesting aspects about BGP recovery time. In particular, degree distribution has a much stronger influence on the recovery time than distance based connectivity or clustered location of routers. Also, a heavy tailed distribution of connectivity actually helps in bringing down the recovery time. It is also found that recovery messages rate shows a sharp jump initially, which can be reduced by a more variable MRAI timer value.

The future work includes a more thorough study of BGP recovery and messaging and mechanisms to a) reduce the recovery time, and b) to reduce the impact of recovery process on packet loss and delays.

REFERENCES

- [1] A.L. Barabasi and R. Albert, "Emergence of Scaling in Random Networks," *Science*, pp 509-512, Oct 1999.
- [2] T. Bu and D. Towsley, "On Distinguishing between Internet Power Law Topology Generators," in *Proc. Infocom 2002*, June 23–27, New York.
- [3] "The Border Gateway Protocol", Cisco Systems, at <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito.doc/bgp.htm>.
- [4] Dan Pei, B. Zhang, et al., "An Analysis of Path-Vector Routing Protocol Convergence Algorithms," *Computer Networks*, 2005.
- [5] T.G. Griffin, B.J. Premore, "An experimental analysis of BGP convergence time," 2001 International Conference on Network Protocols ICNP, pp. 53-61, 2001.
- [6] Bassam Halabi, *Internet Routing Architectures*, Cisco Press, 1997.
- [7] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability," *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 515–528, 1998.
- [8] Labovitz, C., Ahuja, et al., "Delayed internet routing convergence," In *Proc. of SIGCOMM (Aug 2000)*, ACM, pp. 175-187.
- [9] C. Labovitz, A. Ahuja, et al., "The Impact of Internet Policy and Topology on Delayed Routing Convergence," *INFOCOM*, Anchorage, Alaska, April 2001.
- [10] A. Lakhina, J.W. Byers, et al., "On the Geographic Location of Internet Resources," *IEEE Journal on Selected Areas in Communications* 21 (2003) 934–948.
- [11] A. Medina, A. Lakhina, et al., "Brite: Universal topology generation from a user's perspective," In *Proc. of MASCOTS*, October 2001.
- [12] D. Obradovic, "Real-time Model and Convergence Time of BGP," in *Proc. Of IEEE Infocom*, 2002.
- [13] Y. Rekhter and T. Li, "Border Gateway Protocol 4", RFC 1771, SRI Network Information Center, July 1995.
- [14] "SSFNet: Scalable Simulation Framework" *Network Models*. <http://www.ssfnet.org/>.
- [15] G. Siganos, "Analyzing BGP Policies: Methodology and Tool," *IEEE INFOCOM 2002*. <http://citeseer.ist.psu.edu/652452.html>
- [16] H. Tangmunarunkit, J. Doyle, et al, "Does Size Determine Degree in AS Topology?," *ACM Computer Communication Review*, Oct 2001.
- [17] B. Waxman, "Routing of Multipoint Connections," *IEEE J. Select. Areas Commun. SAC-6(9)*: 1617-1622, Dec 1988.
- [18] B. Zhang, R. Liu, et al., "Collecting the Internet AS-level Topology," *ACM SIGCOMM Computer Communication Review (CCR)*, special issue on Internet Vital Statistics, January, 2005.
- [19] X. Zhao, D. Pei, D. Massey, and L. Zhang, "A Study on Routing Behavior of Latin America Networks," *IFIP/ACM SIGCOMM Latin America Networking Conference*, La Paz, Bolivia, Oct 2003.