Impact of Routing Protocol and Routing Policies on Internet Resilience

Feng Wang and Lixin Gao
Department of Electrical and Computer Engineering, University of Massachusetts
Amherst, MA 01002
{fewang, lgao}@ecs.umass.edu

Studies have shown that the Internet has experienced the widespread reachability problem [1] [2]. Most of them are transient failures. In Figure 1, we show failure durations at four ASs in August 2003. The data was collected from Oregon Route-View [3]. From this figure, we find that more than half of failures last for less than 90 seconds. In this paper, we study the potential causes of transient failures.

One potential cause of transient failures is that Border Gateway Protocol (BGP) is a path vector routing protocol, in which each AS advertises only its best path to its neighbors. When an AS's neighbor uses the best route sent by the AS, the AS cannot get the best route from the neighbor due to *loop avoidance property* of BGP protocol. For example, Figure 2(a) represents a simple example, in which each node represents an AS, while the edge represents the connectivity between a pair of ASs. In Figure 2(a), AS2's neighbor AS1 selects the path (1 2 0) via AS2 as the best path to the destination d. As a result, AS2 does not know the path (1 0) at AS1. This property can lead to transient failures. For example, if the link between AS0 and AS2 fails, both AS2 and AS3 will experience transient failures. Such failures will last until AS2 and AS3 learn the path (2 1 0) from AS1 and (3 2 1 0) from AS2 respectively. This example shows that the loop avoidance property of BGP protocol could be one potential cause of transient failures.

Another potential cause of transient failures is that BGP is a policy-based routing protocol, in which each AS determines whether to announce a route to its neighbor according to routing policies. The decision is usually determined by AS commercial agreements. An AS typically does not carry provider to provider or peer to peer traffic, which is known as no-valley policy [4] [5]. For example, Figure 2(b) represents an AS graph, in which the edge between a pair of ASs represents their AS relationship. AS0 announces prefix d to both providers, AS1 and AS4. AS4 chooses the path (4 1 0) as the best path. The path (4 0) in AS4 is invisible to AS1 because of loop avoidance property of BGP protocol. Furthermore, AS4 cannot propagate the path (4 1 0) to its peers AS2 and AS3 due to the no-valley policy. As a result, AS2 and AS3 have only one path to d. The no-valley policy can make certain paths in an AS invisible to its neighbors. If the link between AS1 and AS0 fails, AS1, AS2 and AS3 will experience a transient failure until they learn the alternate path from AS4. So routing policies could be a potential cause of transient failures as well.

Additionally, the examples in Figure 2 show that transient failures can be propagated to other ASs. For example, in Figure 2(a), the transient failure occurred at AS2 can cause AS3 to experience a transient failure. However, the failure duration at AS3 could be longer than that at AS2. The reason is that AS3 is further away from AS1 than AS2, and AS1 provides the alternate path. Therefore, some ASs could experience longer transient failures than others.

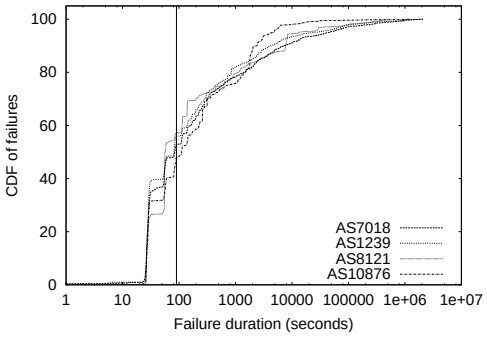

Fig. 1. Distribution of failure durations.

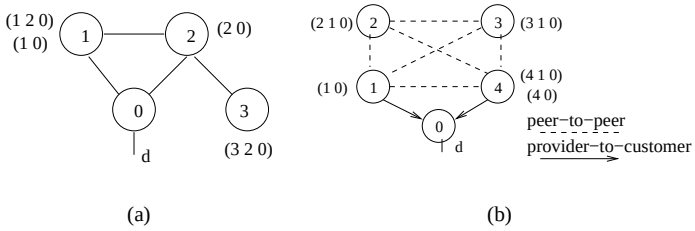

Fig. 2. Examples of transient failures. The paths beside each node indicate available paths to d, and the first one is the best path.

Our study shows that BGP routing protocol and routing policies play a critical role in ensuring the robustness of the Internet. BGP routing protocol and routing policies could cause certain paths invisible to some ASs. They could be the potential causes of transient failures. Furthermore, widespread anomalies such as worm attacks can exacerbate the extent of transient failures. During worm attacks, a large number of BGP sessions in the Internet are reset, which can result in a significant number of routing dynamics. Those dynamics will invoke even more transient failures. Our work can lead to the guidelines for setting routing policies to reduce transient failures in the Internet.

REFERENCES

[1] C. Labovitz, G. Malan, F.Jahanian. Internet Routing Instability. Proc. of SIGCOMM, September 1997.
[2] C. Labovitz, A. Ahuja, F. Jahanian. Experimental study of Internet stability and wide-area network failures. Proc. of Fault Tolerant Computing Symposium, June 1999.
[3] http://www.oregon.com
[4] L. Gao, On Inferring autonomous system relationship in the Internet. IEEE/ACM Transactions on Networking, December 2001.
[5] L. Gao, J. Rexford, Stable Internet routing without global coordination. ACM SIGMETRICS, 2000.