Fault Diagnosis of Path-Vector Routing Through Accumulation of Topological Connectivity Information

Dan Pei, Mohit Lad, Dan Massey, and Lixia Zhang*

1. POSTER ABSTRACT

This paper presents a new approach to diagnosing routing faults (topology changes) in a path-vector routing protocol such as the Internet's BGP routing protocol. The *route diagnosis* problem takes a sequence of updates observed from a node as input, and identifies, as precisely as possible, the topology changes that triggered the updates, and that affected the new paths received by the node. Route diagnosis in BGP is of great interest to both operators and researchers. However, a path vector algorithm by itself fundamentally lacks sufficient information for effective route diagnosis. As a first step towards automatic BGP route diagnosis, we consider a simplified model of BGP where each AS is an atomic node, and nodes are connected by atomic logical links. Given these simplifying assumptions, we propose the following approaches to enhance path vector protocols to achieve effective route diagnosis.

Built-in Diagnosis and Root Causes. Existing path-vector protocols signal the *effects* of the topology changes through the announcement of a new best path. However, a wide range of different topology changes can all produce the same set of new best path updates. In other words, many different events can all result in the same signal (i.e. the set of updates). To achieve accurate diagnosis some (minimal) set of additional information must be added into the routing protocol to facilitate diagnosis.

Recent work including [1] has proposed adding a *root cause* notification attribute (RCN) in order to improve BGP convergence time. The RCN associates a sequence number with a link (or node). A node that detects a link failure will send out a routing update together with the link name and sequence number. The link name and sequence number are referred to as the *root cause* for this update. Any updates that are triggered by the same failure will propagate this root cause information. Analysis and simulations have shown that, by using RCN, routers can detect and avoid obsolete path and achieve substantial reduction in routing convergence time. This root cause attribute directly signals the triggering reason of each update, and we adopt it in our approach. The root cause explains why a previous best has gone, but does not necessarily explain why a node has chosen its new best path among many alternate paths. We propose the *Queued Root Cause Notification* approach to address this issue. In the Queued RCN approach, those root causes received by a node but didn't change this node's path will be stored in a queue because it explains why the alternate paths are gone. When this node's path does change later, the queued root causes are attached to the new path updates and propagated.

Fully Utilizing Topology Information. In our approach, a router can be attached an auxiliary device, called *monitoring station*, which conducts route diagnosis by passively collecting and analyzing the routing updates received at the router. During routing convergence periods, routers may explore a number of transient paths (both valid and invalid), revealing links and nodes that would not have been seen when the network is stable. Instead of discarding this path information after the network stabilizes, our approach uses the information to build a partial view of the topological connectivity. The root cause information provides us explicit information regarding the failure (or recovery) of some links. This combined information allows us to build an estimate of the network topology graph G.

The root cause information attached to updates directly signals the topology changes that have triggered the incoming updates. The network graph G plus the queued root causes help explain why nodes have chosen the new paths among alternate paths. One side benefit is that one can perform "path labeling" based on the topology graph and the root cause information. For example, if a path includes a failed link, one can classify such a path as "obsolete"; if a path appears to be the best available one according to the graph, it can be classified as "eventual". All this information can be combined to provide an operator or researcher with the information necessary to debug/deal with route changes.

Our experiments using real BGP updates show that the topology graph can be implemented with a reasonable overhead. Our proof-of-concept simulation also demonstrates the effectiveness of our approach. More detailed measure of diagnosis effectiveness is ongoing. We are also investigating the attack/misconfiguration detection based on the topology and root cause information. In summary, the accumulated topology with queued root causes open a wide door to the route diagnosis of path-vector protocols.

*Dan Pei, Mohit Lad, and Lixia Zhang are with UCLA. Email:{peidan,mohit,lixia}.cs.ucla.edu. Dan Massey is with Colorado State Univ. E-Mail: massey@cs.colostate.edu

2. REFERENCES

[1] D. Pei, M. Azuma, N. Nguyen, J. Chen, D. Massey, and L. Zhang, "BGP-RCN: Improving BGP Convergence Through Root Cause Notification," Tech. Rep. TR-030047, UCLA CSD, October 2003.