Routing Research Issues
(Position Statement for WIRED 2003)

Z. Morley Mao

How to debug the routing system?

Problem: Today, network operators have very limited tools to debug routing problems. Only primitive tools such as traceroute and ping are commonly used to identify existing routing behavior. There is very little visibility into the routing behavior of other ISPs' networks from a given ISP's perspective, making it even more difficult to identify the culprit of any routing anomalies. This also means that it is difficult to predict the impact any routing policy change has on the global routing behavior. Oftentimes, routing problems are noticed only after a customer complains about reachability or severe degradation of performance. There is lack of proactive, automated analysis of routing problems that detect routing problems at early stages. As certain routing problems initially may not be very obvious and result in suboptimal and unintended routes. Diagnosing Internet routing problems often requires analysis of data from multiple vantage points.

Proposed solutions: (1) Build routing assertions, so that nothing fails silently. When network operator configures a network, it is important to create a set of assertions, equivalent to integrity constraints in database or assertions in software programs. This generates the expected behavior of the routing protocols in terms of which routes are allowed, the resulting attributes of the routes, etc. These constraints can be checked dynamically by a route monitor.

(2) cooperation among networks Each network builds a measurement repository to collect data from multiple locations. It builds a profile of the expected routing behavior to quickly identify any deviations using statistic techniques. Cooperation across networks is absolutely necessary to diagnose global Internet routing problems. It is a challenge to provide summaries of measurement data at sufficiently detailed level to be useful but without revealing sensitive information about internals of ISP's networks. A complementary approach is to allow special distributed queries of the detailed network data from multiple vantage points without direct access to the data.

(3) scalable distributed measurement interpretation and measurement calibrations Routing measurement (e.g., BGP) can result in significant data volume and it may be infeasible to perform real-time or online interpretation of such measurement data by combining all the data from multiple locations in distinct networks at a centralized location. Distributed algorithms are useful to interpret measurement results locally and then aggregate them intelligently to identify routing anomalies. Interpreting measurement can be challenging as there is a lack of global knowledge of topologies and policies which can arbitrarily translate a given measurement input signal to observed output signals. We propose the use of calibration points to help identify expected or normal routing behavior and correlate the output with the input. Calibration points are well-controlled active measurement probes with known measurement input. The BGP Beacons work is one such example of an attempt to understand the patterns of output for a known input routing change.

(4) Internet-wide emulation for network configurations The impact of a single routing configuration change caused by a policy change for example could be global; thus, it is important to emulate the behavior in advance to study its impact. It is useful to abstract the routing behavior in a single network at a higher level to study the perturbation on the global routing system. Currently, the routing configuration is done at a device level. Higher-level programming support is needed to provide semantically more meaningful configuration of networks. Predicting the output of a routing configuration implicitly assumes that routing is deterministic. However, nondeterministic routing may be more stable by preferring routes that have been in the routing tables the longest. Such tradeoffs are important to study.

(5) Understanding the interaction of multiple routing protocols and implementation variants Internet routing consists of multiple protocols, e.g., interdomain, intradomain routing protocols, and MPLS label distribution protocol. All these protocols interact to achieve end-to-end routing behavior from an application's point of view. It is critical to understand their dependency on each other. For instance, in BGP/MPLS IP VPNs, the label distribution protocol is needed to set up label switched paths across the network and if that is unsuccessful, BGP cannot find a route. There is similar dependence of BGP on OSPF or IS-IS. Implementation variants among router vendors determine routing

dynamics which is poorly understood. The interaction among the variants may result in unexpected behavior and needs to be studied.

(6) Understanding routing ”politics” When a customer complains about routing problems either in terms of reachability or poor performance, it typically is in the context of some applications. Network operators install route filters in the routers to determine which routes to accept in calculating the best path to forward traffic. Packet filters at the routers are much more flexible in the sense that they determine which packets are accepted for forwarding based on attributes of the packets, e.g., port numbers, protocol types. Given a route in one’s routing table received by one’s upstream provider, there is no guarantee that all application traffic can reach the destination due to the presence of packet filters. Some networks, for instance, perform port-based filtering to protect against known worm traffic. When debugging routing problems, one needs to view from application’s perspective to understand which type of application traffic is correctly forwarded.

How to improve the application performance?

Problem: Today, the Internet has no performance guarantees for real-time or delay-sensitive applications, such as VoIP, gaming, especially if traffic goes across multiple networks. To obtain flexible routing in terms of control over cost and performance of network paths, end users resort to either multihoming to multiple networks or overlay routing. However, studies have shown that there may be potential adverse interaction between application routing and traffic engineering at the IP layer. Multihoming, similarly, is not a perfect solution as it does not directly translate to paths with performance guarantees, has little impact on how incoming traffic reaches the customers, and may further amplify the amount of routing traffic during convergence.

Proposed solution: Application is the king: correlate routing with forwarding plane, evaluate and improve in the context of application performance metrics: delay, loss rate, and jitter.

When studying routing protocol performance, researchers often use convergence delay as a universal metric. However it does not translate directly to metrics applications care about, e.g., delay, loss rate, and jitter. Understanding the stability of such measurements as a function of the network topology and time provides a way for overlay routing algorithms to intelligently route around network problems. Application performance measurements also expose the detailed interaction between the dynamics of forwarding plane and control plane.

How to protect the routing system?

Problem: There has been relatively little studies on protecting the Internet routing infrastructure against attacks. Vulnerabilities in router architectures are relatively unknown and have not been widely exploited. The routing system can also be indirectly affected due to enormous traffic volume. Recently, there has been a large number of worms exploiting end host OS vulnerability. Significant attack traffic volume causes router sessions to time out. Session resets result in exchange of entire routing tables and disruption of routing. Cascaded failures can occur if the session reset traffic subsequently cause router overload and other peering sessions to be affected.

Proposed solution: (1) Understanding vendor implementation of routing protocols Through detailed black-box testing and support from vendors, one can better understand the obscure, undocumented behavior of routers that are not documented in RFCs and their implication on router security.

(2) Understanding vulnerability points on the Internet Network topology and policy information are more widely known through various Internet mapping effort. Such mapping efforts help us discover vulnerability points by analyzing failure scenarios.

(3) Higher priority for routing traffic The delay and loss of routing traffic, especially keepalive HELLO messages, can cause sessions to reset. This can occur when there is significant data traffic. Increasing the queuing and processing priority of routing packets in the routers is one possibility to reduce the impact of bandwidth attacks on the routing system.

(4) Automated dynamic installation of packet and route filters The attack against windowsupate.com was prevented just in time by invalidating the relevant DNS entry in the DNS system, which takes at least 24 hours to propagate any change globally. To react to any attacks in real time, there needs to be a faster and automated way. One possibility is to dynamically install relevant packet and route filters across a selected set of networks to eliminate/reduce the impact of the attacks. Routers have limited memory for such filters and the order of the filters determine the actual routes or packets permitted. We need to study efficient algorithms to compute such filters on the fly.