

# RESEARCH STATEMENT

Ying Zhang (wingying@umich.edu)

## I. RESEARCH AREA OVERVIEW

Despite the enormous popularity and success of the Internet leading to a wide range of network applications available today, we still cannot fully depend on networks for time-critical wide-area applications, such as remote surgery and first-responder emergency coordination. The main concerns are that due to the best-effort nature of the Internet, potential occurrences of unexpected failures and attacks may cause severe performance degradations. Given the scale, complexity, and openness of the Internet, such events can occur at any time at any network locations. Real-time network monitoring is therefore indispensable for enabling quick identification of any performance problems so that mitigation response can be undertaken to limit the impact of any disruption event.

A large body of research work already exists focusing on Internet Service Provider (ISP) based monitoring. However, end-to-end performance is determined by all networks composing the network path. Network performance data, e.g., latency and loss rate, can be passively collected by monitoring within ISPs to detect violations of Service Level Agreements (SLAs). Such an ISP-centric approach is insufficient as these SLAs usually measure coarse-grained performance of paths limited to within each network. Thus, it cannot accurately represent the end-to-end *user-perceived performance* on the paths which traverse multiple ISPs. Moreover, ISP's data are usually proprietary and inaccessible to end-host users, making it challenging to find ISPs responsible for encountered performance degradations. Therefore, we explore a purely end-system based approach for network monitoring which can more accurately represent application-perceived performance and enable efficient mitigation of failures.

The goal of my research is to **design fundamental building blocks to help enhance the robustness and fairness of the Internet in the presence of unexpected failures and attacks to enable effective prevention and recovery**. The improvement in *robustness* entails enhanced capability to detect, predict and react to failures. The improvement in *fairness* means to provide end users with the necessary information to diagnose their own network problems or uncover any potential traffic discrimination. Unlike existing end-host based approaches to network monitoring, my unique contribution is to build large-scale, accurate and efficient network monitoring systems on both routing and forwarding layers from a purely end-systems' perspective, which enables end-hosts to diagnose and react to performance degradations in real time. As a part of my thesis, I have built a large-scale distributed system running on 300 machines monitoring 19 ISPs simultaneously, which demonstrates the feasibility and value of my research in the real world.

## II. DISSERTATION RESEARCH: IMPROVING THE ROBUSTNESS AND FAIRNESS OF THE INTERNET

Monitoring network disruptions is important given their large impact on end-to-end performance [2], [3], [11]. My dissertation research focuses on methodologies for **scalable and accurate end-host based Internet monitoring**. My work explores three underlying themes essential to any monitoring system design.

- 1) Monitoring locations are of critical importance in monitoring system design.
- 2) Information exchanged across locations and protocol layers can significantly improve scalability and accuracy.
- 3) Exploring predictability from monitored events is essential for designing intelligent mitigation.

I illustrate these three themes from four dimensions under the context of designing a large-scale monitoring system: where to monitor, what to monitor, how to diagnose using monitored observations, and finally how to construct mitigation response based on observed monitoring events.

**1) Where to monitor.** We demonstrate the first theme of the *critical importance of monitor location selection* by examining the constraints imposed by monitors on a diverse set of applications depending on it [1]. Network monitors are systems used to collect various performance data. A variety of networking research, e.g., troubleshooting, modeling, and security analysis, all heavily depend on data obtained from network monitoring, while the data can differ significantly across vantage points. Despite its importance, the monitor selection problem has not been well-studied, particularly in the context of routing failure and attack monitoring.

Today, several public route monitoring systems such as RouteViews and RIPE have been widely used for understanding the Internet routing system by gathering real-time Border Gateway Protocol (BGP) updates from

various networks. Many studies have been relying on such data sources by assuming reasonably good coverage and thus representative visibility into the Internet routing system. We studied the impact of diverse deployment schemes in answering several important research questions such as Internet topology discovery, routing attack detection, and inference of AS-level relationships and paths. My study provides insights on selecting route monitors to improve diversity and coverage.

Besides its impact on coverage and visibility, monitor selection is also important for improving scalability. Probing overhead is a key concern especially for end-hosts due to limited resources. To ensure low overload, I have developed a monitoring system consisting of geographically distributed end-hosts to form a connected graph according to the routing topology. Each end-host is assigned specific tasks of active probing. I have designed a multi-ISP optimization algorithm that automatically generates an optimal task assignment periodically upon any topology changes, based on a carefully designed optimization algorithm to minimize active measurement traffic while maximizing monitoring coverage by monitoring multiple ISPs simultaneously.

*Summary statement:* I have studied monitoring selection and its impact on both coverage and system overhead to demonstrate its importance in designing monitoring systems.

**2a) What to monitor.** I illustrate the second theme, *information exchanged across protocol layers can significantly improve accuracy*, by building a cross-layer monitoring system. Internet monitoring can be done on the *control plane* (routing data) and the *data plane* (packet forwarding). Control-plane monitoring is often performed passively by setting up BGP sessions with routers in various networks. Data-plane monitoring normally takes the form of actively sending probe packets along the network path. Previous works have studied monitoring in either the control plane or the data plane individually. Each approach has its own benefits: the former is less noisy with less overhead while the latter can capture the actual path traversed and performance experienced by the application. I explore the benefits of combining both approaches to improve the accuracy and reduce monitoring overhead through two applications.

One application combines monitoring in both planes to understand the impact of routing dynamics on data plane performance. I developed a measurement framework deployed at each vantage point with access to real-time BGP routing updates [3] for monitoring data plane failure induced by routing changes. Light-weight probing is triggered by locally observed routing updates. The probing target is an identified live IP address within the prefix associated with the routing change. Using this framework, I focused on detecting *data plane failures*, which indicate severe performance degradation on packet forwarding manifested as reachability loss or forwarding loops, which are primarily caused by routing changes. In this work, I identified that 45% of observed updates cause the destination becoming unreachable shortly after respective routing changes. A total of 91% of the unreachable events are short-lived, lasting less than 5 minutes. This observation uncovers the potential significant impact of routing changes on data plane performance.

Another application to illustrate the benefit of cross-layer protocol monitoring is to detect and mitigate network attacks. Most systems designed to detect and mitigate critical security problems like Internet worms and spam ignore control plane signatures, and focus only on Internet data plane information (packet headers and payloads). I have explored how to break down the barrier between the two planes in order to use information and anomalies detected on the data plane to inform the control plane decision support and to use anomalies detected on the control plane to inform data plane filtering [8]. Besides the information exchanged between protocol layers, I have also studied the correlation across spatial aggregation levels, e.g., routers, PoPs, ASes using multi-resolution analysis [11] to better understand the properties of Internet traffic distribution at different layers.

*Summary statement:* I have built systems by correlating information from the control and the data plane to effectively detect performance degradation and security attacks, demonstrating the benefit of information exchange across protocol layers in network monitoring.

**2b) How to diagnose.** Data obtained from network monitoring is subsequently used to detect and diagnose problems. In the following, I illustrate the other angle of the second theme, *correlating information across locations can significantly improve diagnosis accuracy*.

Data from network monitoring are often used to diagnose network failures and performance degradations. However, any end-host based system usually faces challenges in achieving high accuracy due to limited visibility and a lack of ISP proprietary information such as network topology and configurations. I have explored two applications using end-system based monitoring: to accurately *locate the disruptions* and to *detect traffic differentiation*. The

key idea is to use *collaboration* to achieve accuracy.

In the first application, I focused on diagnosing the locations and root causes of network disruptions, e.g., loss of reachability, from monitored data [4]. Correlating the traces from multiple locations, I designed an inference algorithm to identify the minimum set of root causes that can explain the most observed disruptions. I modeled the problem as a bipartite graph to search for plausible explanations of these events using a greedy algorithm. This is based on the intuition that routing events occurring close together are likely caused by only a few causes, which do not create many inconsistencies. My work is the first to enable end-systems to scalably and accurately diagnose causes for routing events associated with large ISPs without requiring access to any proprietary data such as real-time routing feeds from many routers inside an ISP.

Besides failure diagnosis, observations from monitoring can be used to detect traffic differentiation which may result in performance difference, dependent on various factors such as network paths, congestion degrees, or traffic types. Particularly, I am interested in detecting performance differences caused by intentional differentiated traffic treatment by ISPs. This is motivated by the topic of “network neutrality” which has become a critical social and technical problem, as ISPs may discriminate traffic in various ways, e.g., giving peer-to-peer application traffic lower priority, providing differentiating quality of service based on customer identities. I have designed a novel application content-aware probing technique to monitor the service provided for diverse applications and customers [5]. This helps detect traffic discrimination useful for ensuring fairness of the Internet.

*Summary statement:* I have designed and implemented diagnosis algorithms to correlate information from multiple locations to detect routing-induced failures and performance differences.

**3) How to mitigate.** Given the information on detected failure and the diagnosed causes, proactive mitigation is performed to improve the network performance. Next, I illustrate the third theme of *exploring predictability from monitored events for constructing mitigation responses*. The monitored observations and diagnosis results should be used to construct reactive mitigation schemes to reduce failure impact and prevent attacks in the long term. I have studied two mitigation approaches: (1) exploring the predictability of routing induced failures to bypass the disruption, (2) constructing filtering schemes to immediately limit attack impact.

I explored the predictability of failures associated with certain routing changes [3]. Using probability and entropy theory, I developed a methodology to predict the impact of future routing updates based on the identified inherent stability of routing changes. The model is trained based on the analysis of the AS paths in the routing updates, especially in the changing ASes by comparing the new and old AS paths. We have been able to achieve 90% accuracy. The ability to accurately predict routing-induced data plane failures is directly useful for applications such as overlay route selection and backup path selection. As another dimension, I have also studied the predictability of performance changes such as latency caused by routing events [2]. We observed the predictable latency difference associated with certain routing changes. Correlating and predicting latency with corresponding changes can help end-systems construct more intelligent routing to improve application performance.

As another example of exploring predictability in observations from monitoring, I developed a system to secure the local interdomain routing system based on learning from historical information. Since the current routing infrastructure is vulnerable to various types of attacks, I have built a system to help routers detect anomalies and protect routers from attacks by malicious peers, which is one step towards achieving better routing security [6]. My approach is to protect routers by adding a firewall in front of each router to detect and stop any unwanted routing traffic from reaching the local routers. The anomalous updates are detected by finding significant deviations from historical behavior.

*Summary statement:* I have designed algorithms to learn the predictable observations from data obtained from monitoring to construct mitigation for both routing failures and security attacks.

Overall, my dissertation is a first attempt towards monitoring SLA compliance from an end-host user’s perspective. The system carefully selects monitoring locations, performs cross-layer monitoring and cross-location diagnosis. Moreover, I have explored the predictability from monitored data for intelligent mitigation.

### III. OTHER RESEARCH: UNDERSTANDING AND IMPROVING INTERNET SECURITY

While my dissertation mainly focuses on building efficient network monitoring systems, I have a wide range of interests in the area of network security as Internet users and network operators have become more frequently plagued by malicious network activities. I have worked on the detection and prevention of IP prefix hijacking

attacks [9], [10], where malicious routers inject and propagate false routes to the global Internet potentially causing traffic to be redirected to the attacker networks.

Another form of attacks targeting routing protocols is the exploitation of transport layer vulnerabilities. Currently there is usually no protection in the form of prioritized use of router resources for control plane packets. Thus, congestion of other data traffic can adversely affect BGP packets. I have conducted a study on how the recently identified low-rate TCP-targeted DoS attacks disrupt interdomain routing [7]. I systematically examined the severe impact of interdomain routing by attacks exploring the deterministic congestion control behavior in the transport layer on commercial routers. I subsequently proposed prevention strategies by prioritizing routing traffic using existing router support.

#### IV. FUTURE RESEARCH AGENDA

In the course of my research, I have gained several insights into network monitoring and troubleshooting, which can be generally applied in other types of networks. In the near future, I am interested in identifying the new challenges of network monitoring in networks such as large-scale enterprise networks, data centers, and online social networks. The monitoring techniques in the Internet cannot be trivially applied to such systems. There are a number of interesting research questions in these areas.

**Monitoring in enterprise network and data center environment** has to consider more complex common underlying shared components at different levels. Today's data centers contain tens of thousands of computers running various applications ranging from static content-based web services to complex cloud computing. There are three key challenges in monitoring such networks. First, complex sharing structure exists. Machines in the same rack share common power failures. Clusters of machines connecting to the same switches/routers are impacted by the same set of network failures and bandwidth limits. Processes of the same applications running on different machines may share common logical failures, e.g., application level deadlock. Second, the requirements vary for monitoring at different network levels and for different applications. Third, the deployment is difficult in a real production system. Unlike the Internet, active monitoring in the data center environment is usually difficult given cost, reliability, privacy, and security concerns. The design of monitoring systems with little overhead and without information leakage can help overcome deployment obstacles.

**Monitoring in online social networks** has larger challenges as the normal behavior is not well-defined. Learning techniques can be used to detect anomalous behavior. Social network properties such as degree distribution, connectivity, and communication patterns can be explored in the construction of the monitoring system. Overall, I believe my previous work provides a solid background to address these challenging problems.

**Analyzing the economical and technological factors in the Internet.** My research so far has focused on enabling end-users to gain sufficient information about ISP networks instead of treating the network as a black box. The topic of network neutrality is currently highly contentious for today's Internet. Previously, ISP networks have been assumed neutral in carrying traffic without any preferential treatment. I am interested in exploring both the economic and technological factors together with the monitoring results to give reasons about the necessity and benefit of various "net-neutrality" violation actions. For example, end-users and ISPs can be modeled as multiple parties on the market with equivalent information given end-host based monitoring. We can study the profit gain, technology cost, and potential risk and penalty in having traffic differentiation. I would also be interested in new network architecture and protocol designs that could create a more open market model for the Internet.

#### REFERENCES

- [1] Ying Zhang, Zheng Zhang, Z. Morley Mao and Y. Charlie Hu, "On the Impact of Route Monitor Selection", *Proc. ACM SIGCOMM IMC*, Oct. 2007.
- [2] Himabindu Pucha, Ying Zhang, Z. Morley Mao and Y. Charlie Hu, "Understanding network delay changes caused by routing events", *Proc. ACM SIGMETRICS*, June 2007.
- [3] Ying Zhang, Zhuoqing Morley Mao and Jia Wang, "A Framework for Measuring and Predicting the Impact of Routing Changes", *Proc. IEEE INFOCOM*, May 2007.
- [4] Ying Zhang, Z. Morley Mao and Ming Zhang, "Effective Diagnosis of Routing Disruptions from End Systems", *Proc. Symposium on Networked Systems Design and Implementation (NSDI)*, Apr 2008.

- [5] Ying Zhang, Morley Mao and Ming Zhang, “Ascertaining the Reality of Network Neutrality Violation in Backbone ISPs”. *Proc. ACM SIGCOMM HotNets-V*, Oct 2008.
- [6] Ying Zhang, Zhuoqing Morley Mao, and Jia Wang, “A Firewall for Routers: Protecting Against Routing Misbehavior”, *Proc. Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, June 2007.
- [7] Ying Zhang, Z. Morley Mao, and Jia Wang, “Low-Rate TCP-Targeted DoS Attacks Disrupts Internet Routing”, *Proc. Annual Network & Distributed System Security Symposium (NDSS)*, Feb. 2007.
- [8] Ying Zhang, Evan Cooke, and Z. Morley Mao, “Internet-scale Malware Mitigation: Combining Intelligence of the Control and Data Plane”, *Proc. ACM CCS Workshop on Rapid Malcode (WORM)*, Nov. 2006.
- [9] Zheng Zhang, Ying Zhang, Y. Charlie Hu, and Z. Morley Mao, “Practical Defenses Against BGP Prefix Hijacking”, *Proc. 3rd International Conference on emerging Networking EXperiment and Technologies (CoNEXT)*, Dec. 2007.
- [10] Zheng Zhang, Ying Zhang, Y. Charlie Hu, Z. Morley Mao, and Randy Bush, “iSPY: Detecting IP Prefix Hijacking on My Own”, *Proc. ACM SIGCOMM*, Aug. 2008.
- [11] Abhinav Pathak, Himabindu Pucha, Ying Zhang, Y. Charlie Hu, and Z. Morley Mao, “A Measurement Study of Internet Delay Asymmetry”, *Proc. Passive and Active Measurement Conference (PAM)*, Apr. 2008.
- [11] Ying Zhang, Zihui Ge, Suhas Diggavi, Z. Morley Mao, Matthew Roughan, Vinay Vaishampayan, Walter Willinger, and Yin Zhang, “Internet Traffic and Multiresolution Analysis”, *Markov Processes and Related Fields: A Festschrift in Honor of Thomas G. Kurtz*, S. N. Ethier, J. Feng and R. H. Stockbridge (eds.), *IMS Lecture Notes–Monograph Series*, 2007.
- [12] Jian Wu, Ying Zhang, Z. Morley Mao, and Kang G. Shin, “Internet Routing Resilience to Failures: Analysis and Implications”, *Proc. 3rd International Conference on emerging Networking EXperiment and Technologies (CoNEXT)*, Dec. 2007.