

Programmable Measurement Architecture for Cloud Operations

We believe programmable network measurement is an important piece for cloud operations. Network measurement is at least half of network operations. Most network operations such as traffic accounting, traffic engineering, load balancing, and performance diagnosis, rely on accurate and timely measurement of time-varying traffic across the entire network. Many companies (e.g., Google, Microsoft, Facebook) require huge capital investments to build larger networks with higher link speeds; yet provide surprisingly little visibility into the network and traffic. These companies still rely on obsolete measurement tools on isolated devices, which are either too coarse-grained to capture important events (e.g., NetFlow, sFlow), or too heavyweight to run all the time (e.g., tcpdump). Operators have to take great efforts to manually join a large amount of data from many devices and identify useful information. Such measure practice becomes increasingly difficult when these companies drive their networks to higher utilization (even up to 100%), which requires finer-grained flow information, more timely report of traffic changes, and automatic performance diagnosis.

There are three key problems of today's measurement support: (1) Network device vendors often treat measurement as a second-class citizen, devoting most of the resources to meet the networking control functions (forwarding, firewalls, load balancing, etc.), leaving limited resources for supporting measurement. (2) Operators have limited control over what (not) to measure and when to measure. As a result, the already limited measurement resources are sometimes wasted to measure flows that operators do not care, leaving even fewer resources for measuring important flows. (3) These solutions are device oriented instead of network-wide. Operators have to dive into the limited measurement support at multiple devices, taking great efforts to understand these measurement results offline, and find it challenging to answer their network-wide queries.

We claim that today's cloud lack a programmable measurement framework. Inspired by software-defined networking, the programmable measurement framework allows operators to flexibly program their *network-wide* measurement queries in a controller. To answer these queries, the framework automatically configures and coordinates new measurement primitives at different places across the network stack. The framework also allows operators to dynamically change their queries and automatically reconfigure the primitives accordingly to handle network dynamics.

This programmable measurement framework is potentially useful for a wide variety of cloud applications: **For cloud operators**, the framework allows them to monitor in real time the health status of the cloud, and perform daily network management operations such as traffic engineering and load balancing, and identify potential attacks and malicious traffic (e.g., DDoS, port scan). **For application developers**, our measurement framework could be useful for them to diagnose performance problems (e.g., localize the devices/application modules that cause long delay for requests) and track unexpected behaviors (e.g., identify root causes of packet losses).

There are four key components in the framework that are important for the cloud infrastructure today: (1) We need novel measurement algorithms and designs throughout the network stack (VMs, hypervisors, switches, and packet sniffers). We need to redesign the *measurement primitives* at these devices to make them both *generic* in supporting diverse measurement requirements, and *efficient* in packet processing performance with limited resources and capabilities. (2) We need new *declarative measurement abstractions* for operators to clearly express what to measure

at the network level and their accuracy/timeliness requirements. We also need *a runtime system* that automatically matches the measurement abstractions with primitives at devices, dynamically allocates resources across tasks and handles network dynamics such as mobile hosts and routing changes. (3) Our framework will improve the interactions between measurement and control by optimizing measurement queries for specific control functions, and redesigning control solutions based on measurement results.

The proposed work will fundamentally change network measurement and management practice in enterprise and data centers. The research will lead to new designs of network devices, by bringing in new measurement algorithms, making the devices more programmable and efficient for measurement. The research will facilitate close interactions between theory, programming languages, and networking. Compared to companies with large data centers (e.g., Google, Microsoft), which have little flexibility in changing their current measurement infrastructure, we believe that the cloud infrastructure for researchers should have better flexibility in supporting the programmable architecture and can serve as the testbed for future measurement architecture in the production clouds.