Abstract

Software Defined Networking (SDN) is a new field in the networking area and is growing rapidly. Over the past few years, SDN has gained significant attention from both industry and academia. To accelerate adoption of SDN, in 2011, Open Network Foundation (ONF) was established. It is an industrial driven organization. Around 2012, Stanford and UC Berkeley established a new research center, Open Networking Research Center (ONRC), especially to design, develop, and deploy new frameworks for SDN. SDN promises to provide enhanced configuration, flexible control over network devices, and encourage innovations by decoupling the control plane from the data plane. SDN network is controlled by a central entity called SDN controller. SDN controller requires network statistics to perform various network management tasks such as routing, QoS assurance, load balancing, congestion detection, etc. The main goal of this thesis is to propose new and efficient methods to collect network-wide statistics. To achieve this goal, we have worked on three aspects, (1) collection of consistent statistics, (2) measurement of queueing delay, and (3) determining the polling frequency for network statistics collection.

In an SDN network statistics can be collected by querying the underlying switches to capture their states. States of the switches are causally related i.e., an attribute in one switch is causally affected by the same attribute of another switch. Such a state needs to be measured preserving this causal relation. In certain applications such as load balancing, congestion prediction, trace recording, applying updates consistently on all switches, dynamic visualization of network traffic patterns, a measurement of the switches' states that preserves this causal order is expected to yield accurate results. Prevalent network monitoring tools focus on per device or per port statistics collection.

These statistics, when viewed across the switches, are likely to be inconsistent if a specific order is not enforced while collecting them. Collecting globally consistent statistics requires coordination among all the underlying network switches. Due to the asynchronous nature of the network, the statistics collected through polling the switches have inconsistencies and are not suitable to capture a consistent global state of the network. Globally consistent state detection problem is well studied for asynchronous systems. However, to the best of our knowledge, the current SDN standards, such as OpenFlow, does not provide any support to collect globally consistent statistics. In this thesis, we propose a novel method, "OpenSnap", which extends the capability of OpenFlow to collect globally consistent statistics in OpenFlow based SDN networks. OpenSnap works well for the networks which have switches connected through FIFO (First-In-First-Out) channels. If switches are connected through Non-FIFO channels, OpenSnap requires multiple rounds to collect globally consistent statistics. Also, it has to restart the statistics collection process again in case of an interruption. As an improvement upon OpenSnap, we propose an efficient and robust method, GlobeSnap, to collect globally consistent statistics in OpenFlow based SDN networks with both FIFO and Non-FIFO channels. Unlike, OpenSnap, it does not require any control packet to travel through the data path for statistics collection. Instead, it uses the network traffic itself for the statistics collection. Thus, reduces the monitoring overhead. We have evaluated the proposed solutions in terms of percentage of consistency achieved. OpenSnap provides 100% consistency in a network with FIFO channels and 40.33% in a network with Non-FIFO channels. GlobeSnap provides 100% consistency in both the network with FIFO channels and Non-FIFO channels. Whereas the existing methods achieve a maximum of 59.89% consistency.

Over the last few years, there has been an enormous rise in the number of devices connected to the Internet. In today's world, application requirements change frequently. Ensuring Quality of Service (QoS) to the end-users has become a big challenge for the network operators and service providers. To assure QoS, the network operator has to update the policies and reconfigure each device on the fly, and SDN supports this functionality. To define or update the policies for QoS, the controller needs to measure the QoS parameters such as bandwidth, delay, jitter, etc. In this thesis, we focus on delay measurements. Delay measurement methods can be broadly classified under two categories, active delay measurement, and passive delay measurement. Most of the existing works in the literature to measure delay come under active delay measurement category. The active delay measurement methods are based around calculating the transit time of a specially constructed control packet (probe packet) that travels through the data links. These approaches are not efficient, as the probe packet injected into the data plane incurs considerable overhead. Additionally, a separate probe packet is required to measure each queue's delay if more than one queue is present at the egress port of a switch. Thus, these approaches are not scalable. We propose a novel passive delay estimation method, qMon, to monitor queueing delay in OpenFlow networks. Existing passive delay estimation method works for TCP flows only. qMon is the first of its kind to provide delay measurements in SDN networks irrespective of the type of traffic it carries. qMon leverages the OpenFlow protocol to obtain queue statistics from OpenFlow switches at regular intervals, which are further employed to estimate the mean queueing delay for each interval. The proposed method is scalable w.r.t number of switches and has low monitoring overhead, as it requires sending a single queue-statistics request message to each of the switches to monitor queueing delay at all the links in the network. In an Out-of-band controller configuration, qMon does not have to send probe packets via the data plane to measure the delay. Thus, it does not have any data plane footprints. For an In-band controller configuration, qMon has very low data plane footprint compared to the existing methods. As it requires one queue statistics request message per switch to measure the queueing delay of all the queues. Whereas, the existing methods require one probe packet per queue on each switch. We performed experiments for Poisson traffic and bursty traffic. The results show a high correlation between the qMon and ping RTT values for Poisson traffic and bursty traffic with large ON intervals.

Collection of consistent statistics and delay measurement requires polling of underlying network switches by the controller. So, it is very important to determine an optimal polling rate which minimizes the monitoring overhead without compromising with the accuracy of the collected statistics. There exist two different approaches to poll the network statistics, Pull-based and Push-based. Although, OpenFlow switches provide flow statistics to the controller when the flow expires/time-out. However, these statistics provide sparse information about the network in case there are a large number of elephant flows. Also, the existing push-based approaches require additional software and hardware support. Whereas, in pull-based approaches, the controller polls the underlying switches for the statistics. The issue with pull-based statistics collection methods is that the accuracy of the collected statistics depends on the polling frequency/rate. Polling at a higher rate may provide better accuracy but incurs high overhead. Whereas, a lower polling rate can lead to inaccuracy in collected statistics. Thus, an efficient method is required to determine an optimal polling frequency. An optimal polling rate ensures high accuracy in collected statistics with least overhead. The existing methods use the change in the arrival rate of traffic as a metric to determine the polling rate. These approaches incur high overhead, especially when the change in the collected statistics is linear. Taking this as a motivation, in this thesis, we propose a new pull-based dynamic polling method using curvature based sampling. Unlike existing methods, we tune the parameters used in the proposed method to achieve optimal results. We also provide a cost function to rank different pull-based statistics collection methods. Our method provides high accuracy with low polling overhead when compared with the stat-of-the art. The experiments show that our method performs better than the state-of-the-art methods by at least 23% in terms of accuracy and cost.

Though SDN provides many advantages over traditional networks, many challenges still plague SDN, and transition to SDN can be disruptive and painful. One feasible solution is increment deployment of SDN in traditional networks. This gave birth to a new paradigm called Hybrid SDN. Hybrid SDN, is an architecture where both traditional and SDN paradigm exists and operates in a single network. To leverage benefits of SDN in a Hybrid SDN network, the traffic has to go through the SDN switches. This is ensured through waypoint enforcement. Waypoint enforcement is defined as constraining a packet to take a path which has at least one SDN switch on its way to the destination. Existing methods in the literature achieve only partial waypoint enforcement. In this thesis, we propose a novel framework to achieve full waypoint enforcement. The proposed framework uses virtual IP addresses to divert the entire network's traffic towards the SDN switches. Assuming the network has sufficient capacity, the proposed framework achieves 100% waypoint enforcement.

In this thesis we have proposed methods for consistent statistics collection, delay measurement, and optimal polling frequency determination for SDN networks, and waypoint enforcement in Hybrid SDN network. These solutions have the potential to enhance the performance of SDN networks.