Dapper: Data Plane Performance Diagnosis of TCP

Mojgan Ghasemi
Princeton University
mojgan@cs.princeton.edu

Theophilus Benson
Duke University
tbenson@cs.duke.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

ABSTRACT

With more applications moving to the cloud, cloud providers need to diagnose performance problems in a timely manner. Offline processing of logs is slow and inefficient, and instrumenting the end-host network stack would violate the tenants’ rights to manage their own virtual machines (VMs). Instead, our system Dapper analyzes TCP performance in real time *near* the end-hosts (e.g., at the hypervisor, NIC, or top-of-rack switch). Dapper determines whether a connection is limited by the sender (e.g., a slow server competing for shared resources), the network (e.g., congestion), or the receiver (e.g., small receive buffer). Emerging edge devices now offer flexible packet processing at high speed on commodity hardware, making it possible to monitor TCP performance in the data plane, at line rate. We use P4 to prototype Dapper and evaluate our design on real and synthetic traffic. To reduce the data-plane state requirements, we perform lightweight *detection* for all connections, followed by heavier-weight *diagnosis* just for the troubled connections.

Keywords

Performance Diagnosis, Measurement, Network Monitoring

CCS Concepts: Networks → Network monitoring; Network performance analysis; Network measurement; Transport protocols

1. INTRODUCTION

Public clouds need timely and accurate visibility into application performance to ensure high utilization (to reduce cost) and good performance (to satisfy tenants) [23]. Yet,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SOSR’17, April 03–04, 2017, Santa Clara, CA, USA
© 2017 ACM. ISBN 978-1-4503-4947-5/17/04 ...\$15.00
DOI: <http://dx.doi.org/10.1145/3050220.3050228>

Figure 1: Dapper monitors performance at network edge

cloud providers must monitor performance within their own infrastructure, since they cannot modify the end-host network stack without violating tenant control over their own VMs. Studies reveal that more than 99% of traffic in some data centers is TCP traffic [7]. However, TCP measurements collected in the network core would not have an end-host perspective (e.g., of round-trip times) and would require combining data from multiple switches because the two directions of a connection may traverse different paths.

Instead, measurement at the “edge”—in the hypervisor, NIC, or top-of-rack switch, as shown in Figure 1—offers the best viable alternative. The edge (i) sees all of a TCP connection’s packets in both directions, (ii) can closely observe the application’s interactions with the network without the tenant’s cooperation, and (iii) can measure end-to-end metrics from the end-host perspective (e.g., path loss rate) because it is only one hop away from the end-host.

Fortunately, emerging edge devices offer flexible packet processing at line rate, in software switches [30], NICs [36], and hardware switches [13]. There have been a variety of proposals to achieve stateful programmable dataplanes, including POF [33], OpenState [11], Domino [32] and P4 [12]. New capabilities, such as flexible parsing and registers that maintain state open up the possibility of detecting and diagnosing TCP performance problems directly in the data plane. Still, a major challenge for data-plane connection diagnosis is finding a “sweet spot” that balances the need for *fine-grained* metrics for diagnosis, while remaining *lightweight* enough to run across a wide range of devices with limited capabilities.

This paper presents Dapper, a **Data-Plane Performance** diagnosis tool that infers TCP bottlenecks by analyzing packets in real time at the network edge, as shown in Figure 2.

Dapper quantifies the contribution of the sender, network, and receiver to poor end-to-end performance. Table 1 shows examples of problems that can limit a TCP connection’s performance. Identifying the entity responsible for poor performance is often the most time-consuming and expensive part of failure detection and can take from an hour to days in data centers [8]. Once the bottleneck is correctly identified, specialized tools within that component can pinpoint the root cause. To achieve this goal, we need to infer essential TCP metrics. Some of them are easy to infer (e.g., counting number of bytes or packets sent or received), while others are more challenging (e.g., congestion and receive windows).

Dapper analyzes header fields, packet sizes, timing, and the relative spacing of data and ACK packets, to infer the TCP state and the congestion and receive window sizes. A unique challenge in a public IaaS clouds is that different tenants run different versions of TCP, each possibly with tuned parameters. Thus, our techniques must be applicable to a heterogeneous set of TCP connections. The advantage of a data-plane performance diagnosis, apart from line-rate diagnosis, is that the data plane can use this information for quick decision making (e.g., load balancing for network-limited connections). However, a data-plane monitoring tool often has more resource constraints: limited state and limited number of arithmetic and Boolean operations per packet. We discuss design challenges and the necessary steps for Dapper to run in the data plane, including reducing the accuracy of some measurements (e.g., RTT) to lower the amount of per-flow state, as well as two-phase monitoring where we switch from collecting *lightweight* metrics for *all* flows to only collecting *heavyweight* metrics for *troubled* ones.

Roadmap: Section 2 explains the TCP performance bottlenecks we identify and how to infer the metrics necessary to detect them. Section 3 explains how Dapper diagnoses performance problems from the inferred statistics. Section 4 explains how to monitor TCP connections in real time using commodity packet processors programmed using P4 [12, 3]. Section 5 discusses our two-phase monitoring to reduce the memory overhead in the data plane. Section 6 evaluates the overhead and performance of our system. Section 7 discusses the limitations of our approach and implementation. Section 8 discusses related work, and Section 9 concludes the paper.

2. TCP PERFORMANCE MONITORING

A TCP connection may have bottlenecks at the sender, receiver, or the network. In Table 1, we present several examples of performance problems that may arise at each location. With many applications and TCP variants in a public cloud, it is challenging to decide what minimal set of metrics to collect that are both affordable (i.e., does not consume a lot of resources) and meaningful (i.e., helps in diagnosis). In this section, we discuss the metrics we collect to diagnose the performance bottlenecks at each component, and the streaming algorithms we use to infer them.

Figure 2: Dapper’s architecture : (1) data plane monitoring on edge, (2) control plane diagnosis techniques

Location	Performance Problems
Sender	slow data generation rate due to resource constraints, not enough data to send (non-backlogged)
Network	congestion (high loss and latency), routing changes, limited bandwidth
Receiver	delayed ACK, small receive buffer

Table 1: TCP performance problems at each component

We denote a TCP connection by a bi-directional 4-tuple of server and client IP addresses and port numbers. We focus on the performance of the data transmission from the server, where the server sends data and receives ACKs, as shown in Figure 3. Hence, we monitor traffic close to the server, for several reasons. First, being close to the server enables us to imitate the internal state of the server’s congestion control and monitor how quickly the server responds to ACKs. Second, if a connection is end-user facing (as opposed to east-west traffic in the data center), we do not have access to the client’s edge, so would need to monitor near the server. Finally, monitoring at one end reduces the overhead and avoids redundancy as opposed to keeping per-flow state at both ends. Note that for east-west traffic, a cloud provider can enable monitoring at both ends of a connection, which would offer higher accuracy and better visibility into the connection’s state by collecting more statistics on the client-side (e.g., delayed ACKs).

2.1 Inferring Sender Statistics

Performance problems at the sender limit the TCP connection’s performance; for example, an application that is constrained by its host’s resources (e.g., slow disk or limited CPU) may generate data at a lower rate than the network or receiver accept, or it simply may not have more data to send; this application is referred to as *non-backlogged*. To find such problems, we measure the *ground-truth* by counting sent packets and compare it with the connection’s *potential* sending rate, determined by the receive and congestion windows.

If we could directly monitor the TCP send buffer inside the VM, we could easily observe how much data the application writes to the buffer, how quickly the buffer fills, and the maximum buffer size. In the absence of OS and appli-

Figure 3: Tracking options, segment size, and application reaction time for a simplex connection (server-client)

Figure 4: Tracking a TCP connection’s flight size

cation logs from tenant VMs, we rely on two independent metrics to see if the application is sending “too little” or “too late”: (i) *generated segment sizes*, to measure if an application is sending too little, and (ii) *application’s reaction time*, to see if it is taking too long.

Inferring flight size to measure sending rate: Upon transmission of new packets, we update the packet and byte counters for each connection and measure the connection’s flight size to infer its sending rate (i.e., $\frac{\text{flight size}}{\text{RTT}}$). The flight size of a connection is the number of outstanding packets—packets sent but not ACKed yet—as shown in Figure 4, and is inferred via examining the sequence number of outgoing packets and the incoming acknowledgment numbers to track how many segments are still in flight. In this figure, the “send window” represents the available window to the sender, that is, the maximum packets that the sender can send before waiting for a new ACK; this window is limited by both the RWND and CWND. In this example, the application is not backlogged, because it has not fully used the available window. We can see that flight size can at most be equal to the send window.

Note that the flight size of a connection is an important metric in diagnosis because it depends on the congestion window, the receive window, and the application’s own data-generation rate (i.e., is the application backlogged?). We will revisit this metric in Section 2.2.

Extracting MSS from TCP options: The segment sizes in

a connection indicate the amount of data that the sending application generates. We infer MSS by parsing the TCP options exchanged during the three-way handshake, as shown in Figure 3, in the SYN and SYN-ACK packets.

Measuring sender’s reaction time via cross-packet analysis: We define the sender’s reaction time as the time window between the arrival of a new acknowledgment and the transmission of a new segment. The reaction time evaluates the sending application’s data-generation rate, i.e., whether or not it is backlogged. Lower reaction times indicate that the data was already processed and was just awaiting an opportunity to be sent. We measure the reaction time using timestamps of incoming acknowledgments and outgoing packets. We compare the reaction time with an empirically derived threshold, calculated based on the latency between the edge and the VM.

Note that “cross-packet analysis” happens at the edge, thus the reaction time consists of the application’s own data-generation latency plus the communication latency between the application and edge (i.e., latency in the hypervisor and virtual switch). More importantly, the network latency does not influence this metric. Also, notice that the edge acts as *the single point of observation*, seeing both directions of the flow to do cross-packet analysis. These two conditions are not necessarily true on the switches in the core of network to measure the sender’s reaction time since the packets could have been delayed at earlier network hops—as opposed to the application itself—or worse yet, taken different paths.

2.2 Inferring Network Statistics

Network problems cause poor performance in a TCP connection. For example, in a congested network, the increased packet loss and path latency cause the sender’s TCP congestion-control algorithm to decrease the sending rate. The sender’s reaction to congestion varies based on the congestion control in use (e.g., Reno vs Cubic) and the severity of the congestion itself. To determine and quantify the network limitation in the performance of a TCP connection, we need to measure and compare its congestion window—how much the network allows the connection to send—against the receive window—how much the receiver allows the connection to send—and how much data the sender has available to send.

To infer a connection’s congestion window, we “imitate” the VM’s congestion-control algorithm by tracking flight size, packet losses and their kinds, and duplicate ACKs. We track RTT and RTTvar to pinpoint the effect of network congestion, routing changes, and queuing delay on TCP performance.

Inferring loss via retransmission: We calculate a connection’s loss rate by counting the lost and total packets in the flow. In addition, we use a counter to track the duplicate ACKs and use it to infer the kind of loss: fast-retransmission (FR) is triggered after a fixed number of duplicate ACKs are received (normally 3); however, a timeout is triggered when

no packet has arrived for a while (RTO). We track the sequence number of outgoing packets (i.e., packets sent by the server) which helps us find the retransmission of previously seen sequence numbers. Using the duplicate ACK counter, if we see at least three duplicate acknowledgments before a retransmission, we treat it as a fast retransmission, otherwise, we deduce that the loss was recovered by a timeout.

Estimating latency by passive RTT measurements: We use Karn’s algorithm [22] to estimate SRTT, SRTTvar, and RTO using a series of “RTT measurements”. An RTT measurement is the time between when a segment was sent and when its acknowledgment reached the sender. For any connection, upon transmission of a new segment, we create a (time-stamp, sequence number) tuple and maintain it in a queue. Upon arrival of an ACK, we inspect the queue to see if any of the tuples are acknowledged by it. If so, we create an RTT measurement and use it to update the latency statistics via Karn’s algorithm.

Note that a retransmitted packet cannot be used as an RTT measurement, because the corresponding ACK cannot be correctly mapped to a single outgoing time-stamp. Also, if a connection has multiple outstanding packets, the queue will have multiple tuples, i.e., the length of the queue grows with the flight size. Finally, if an incoming ACK acknowledges multiple tuples (i.e., a delayed ACK) we de-queue multiple tuples but only create one RTT measurement, from the most recent tuple, to exclude the effect of delayed ACK on RTT.

Estimating congestion window via flight size and loss: It is challenging to estimate a connection’s congestion window outside the VM’s networking stack due to following reasons: (1) *Many TCP variations:* There are many different TCP congestion control algorithms used today (e.g., Reno, New Reno, Cubic, DCTCP). Some are combined with tuning algorithms (e.g., Cubic combined with HyStart to tune *ssthresh*) or have configurable parameters (e.g., initial window). (2) *Thresholds change:* *ssthresh* is the threshold that separates slow-start (SS) from congestion-avoidance (CA) in the TCP state machine and is initially set to a predefined value. However, the Linux kernel *caches* the *ssthresh* value to use it for similar connections in the future. Also, *ssthresh* changes throughout the life of a connection (e.g., under packet loss). Therefore, if the full history of a connection (and even *past* connections!) is not observed, *ssthresh* is unknown, making it impossible to detect transitions from SS to CA based on the *ssthresh* threshold.

In the presence of these challenges, we rely on these TCP invariants to infer congestion window: (1) The *flight size* of a connection is bounded by congestion window, as shown in Figure 4. We denote this lower-bound estimate of congestion window with *inf_cwnd*. In a loss-free network, *inf_cwnd* is a moving maximum of flight size of the connection. Note that in the absence of loss, if the connection’s flight size decreases, it is either due to the sender producing less data (i.e., not fully utilizing the window) or the receiver’s limited re-

Algorithm 1: Estimating *inf_cwnd*

Input: multiplicative decrease factor (*C*), initial window (*IW*)

Output: *inf_cwnd*

```

1 while  $P \leftarrow \text{capture new packet}$  do
2   if  $P$  is new segment and flight size > inf_cwnd then
3      $\text{inf\_cwnd} \leftarrow \text{flight size}$ 
4   else if  $P$  is retransmitted then
5     if fast retransmit then
6       if first loss in fast recovery then
7          $\text{inf\_cwnd} \leftarrow C \times \text{inf\_cwnd}$ 
8     else if timeout then
9        $\text{inf\_cwnd} \leftarrow IW$ 

```

ceive window; hence, *inf_cwnd* does not decrease. (2) If a packet is *lost*, we adjust *inf_cwnd* based on the nature of loss, a timeout resets it to *IW* and a fast-retransmit causes a multiplicative decrease.

Estimating the congestion window based on these invariants makes *inf_cwnd* “self-adjustable”, working regardless of TCP variant and configuration, and is calculated according to Algorithm 1. Algorithm 1 consists of a while loop that inspects every new packet. If a new segment is transmitted and the connection’s flight size grows beyond *inf_cwnd*, we update *inf_cwnd* to hold the new maximum value of flight size. In case of retransmissions (loss), we decrease *inf_cwnd* by the multiplicative decrease constant, *C*, if loss is recovered by fast recovery. Otherwise, *inf_cwnd* is reset to initial window if recovered by a timeout. In this algorithm, we assume the cloud provider knows the value of *C* and *IW*, since they can be easily inferred, either indirectly via observing how large the first window is and how it changes after a loss, or directly via tools such as Nmap [2].

Note that *inf_cwnd* as estimated by Algorithm 1 does not require full knowledge of connection’s history, thresholds, or the congestion control algorithm, and is only dependent upon measuring the connection’s flight size and loss, thus solves the challenges above without tenant’s cooperation¹.

2.3 Inferring Receiver Statistics

The receiver-side of a TCP connection can limit the flow by decreasing its advertised window (i.e., *RWND*) or slowing down the rate of acknowledgments [40] to control the release of new segments.

Tracking *RWND* per-packet and per-connection : To quantify the receiver limitation in a TCP connection, we track the advertised *RWND* value *per-packet*, reflecting how much buffer is available on the client. We also track the *per-connection* agreed upon window scaling option during the TCP hand-

¹To keep our heuristics general across all TCP variants, we do not rely on selective acknowledgments.

Figure 5: Dapper’s packet-processing logic

shake, as shown in Figure 3, which is used for scaling RWND.

Inferring delayed ACKs via RTT samples: When an incoming ACK acknowledges multiple tuples in the queue, it must be a delayed ACK, as the client is acknowledging multiple segments at once. When we de-queue tuple(s) based on an incoming ACK, we count and average the number of de-queued tuples per ACK to reflect the effect of delayed ACK.

Summary: Figure 5 summarizes how Dapper updates a TCP connection’s performance statistics while processing a new packet. The packets are first hashed on the four-tuple to either initialize a new flow, or read the existing statistics. Then, based on the direction of the packet, the relevant header fields are extracted and used to update the metrics. The blue boxes show the analytics performed to keep per-flow state (e.g., update flight size) and the pink diamonds show the conditions used to decide which state to update.

3. TCP DIAGNOSIS TECHNIQUES

In this section, we describe how Dapper uses the statistics gathered by the streaming algorithms discussed in section 2. The high-level objective of the diagnosis techniques is to troubleshoot a connection’s performance limitation.

Diagnosing Sender Problems: Our goal is to find if the sender-side is not limiting the connection’s performance (backlogged), or limiting the sending rate via not having enough data to send or taking too long to produce it (non-backlogged).

1. Exponential sending rate indicates a backlogged sender: On transmission of new segments, we first examine the connection’s macroscopic behavior, i.e., the sending rate, and check to see if it is growing exponentially to infer if the connection is in slow start. More accurately, we compare the

relationship between ACKs and the data packets to see how many packets the sender transmits after a new ACK. If sending rate grows exponentially, we know the connection is *not sender-limited*. Otherwise, we attempt to understand if the connection is sender-limited by checking the next heuristics.

2. If the sender is backlogged, it will “completely” use the send window: When sending rate does not grow exponentially, the connection could either be in congestion avoidance with a backlogged sender, or it could suffer from a non-backlogged sender, producing less data than CWND. This heuristic checks to see if the connection’s flight size is consistently less than the allowed window to send, determined by the minimum of RWND and CWND, i.e., if $flight_size < \min(inf_cwnd, RWND)$. If so, the connection’s performance is limited because the sending application does not send more, not because it’s not allowed to.

3. Sending less than allowed, or later than allowed, indicates a non-backlogged sender: Here we examine the connection’s microscopic behavior to see if the connection is under-utilizing the network, not concerning the “number” of packets in flight like the previous heuristic, but instead focusing on the “size” and “timing” of packets. More concretely, we check to see if a connection is sending packets that are smaller than MSS, or if the application’s reaction time (i.e., data generation time) is larger than an empirically derived threshold for backlogged applications. If either of these conditions are met, we conclude that the application is non-backlogged, hence the connection is *sender-limited* indicating that the sender is either not generating enough data to fill up a whole packet, or not responding almost immediately when it is allowed to send.

4. How the flight size changes during a loss recovery gives us a clue of how backlogged the sender is: In addition to the heuristics above, during loss recovery a connection reveals some information about its internal state². As a reminder, fast-recovery causes the congestion window to decrease by a multiplicative factor, C .

Consider a connection not limited by receiver, and assume that the sending application’s data generation rate remains unchanged during the network loss. We denote the flight size of the connection before loss by f_1 and after the loss by f_2 . Figure 6 shows three example scenarios, where loss happens at 30ms, prompting CNWD to decrease by half. The $\frac{f_2}{f_1}$ ratio gives us the following insights: if flight size is closely tracking CWND, $\frac{f_2}{f_1} = C$, the sender is backlogged (app 1); if the connection was not fully using the CWND before loss but is backlogged after the loss, $C < \frac{f_2}{f_1} < 1$ (app 2); finally, if the flight size remains unchanged the sender is not backlogged (app 3). Note that in these examples we assume the state of the sender remains unchanged during the loss recovery.

Diagnosing Network Problems: Our goal is to determine

²This heuristic can be treated as a bonus, and the diagnosis algorithm does not rely on seeing a loss.

Figure 6: Flight size before and after loss

if the network is restricting TCP performance, either due to limited bandwidth (congestion window limited), high packet loss rate, or increased latency due to problems such as queuing delay or routing problems.

1. *Small congestion window hurts TCP’s performance:* When the network has performance constraints, for example limited network bandwidth, the congestion window will limit the rate of the connection, that is: $flight\ size \leq inf_cwnd < RWND$. Upon a packet retransmission, if loss causes the inf_cwnd value to drop below the RWND, we deduce the connection is limited by network.

2. *Increased network path latency slows TCP’s rate:* The sending rate of the connection is a function of both the flight size and RTT; the sender can only increase the window after a new ACKs arrives, which usually takes an RTT. To track the impact of network latency on TCP performance, we use the RTT measurements as explained in section 2. The cloud provider can either define an “expected RTT” per TCP connection based on SLAs, or use the minimum RTT sample per-flow as the baseline. To diagnose path latency problems, we compare the RTT values with the expected RTT to detect if a connection’s latency is acceptable.

Diagnosing Receiver Problems: Our goal is to find if the receiver is restricting TCP performance, either by offering a small receive buffer (receive window), or by delaying ACKs.

1. *Small receive window hurts performance:* Upon updating the RWND and inf_cwnd values, this heuristic compares them with the current flight size to see if the connection’s sending rate is receiver-limited, that is: $flight\ size \leq RWND < inf_cwnd$. If so, the connection is diagnosed as receiver-limited.

2. *Delayed acknowledgment hurts performance:* The receiver can limit TCP performance by sending ACKs with a delay; delayed ACKs has been shown to cause issues in datacenters [40]. When the client sends ACKs with a delay, for example, sending acknowledgments for every other packet, the sender’s opportunity to increase its window is halved³.

³Note that the congestion window on sender-side opens up upon

For each connection, we measure the average number of RTT samples freed by each new ACK and if the average is greater than one, we diagnose the connection as receiver-limited due to delayed acknowledgment.

4. DATA-PLANE MONITORING

In this section, we describe how Dapper tracks TCP connections in the data plane and discuss the principles behind our target-independent solution using P4. We outline the P4 features (e.g., metadata) that enable us to monitor TCP connections according to Figure 5 (Section 4.1); then, we discuss the target-specific resource constraints and how to mitigate them (Section 4.2).

4.1 TCP Monitoring Prototype in P4

P4 is a programming language that allows us to express how packets are processed and forwarded in a target-independent program, therefore, our P4 prototype can run in a public cloud on a variety of targets, as long as at least one of the elements at the edge (the switch, the hypervisor, or the NIC) can run P4 programs.

To monitor TCP connections in the data plane in real time, we need to extract and retain packet header information (P4’s flexible parsing), carry information across multiple stages of packet processing (P4’s metadata), and store state across successive packets of the flow (P4’s registers). Furthermore, to realize the logic in Figure 5, we need to perform specific operations on each packet, shown with blue boxes (P4’s tables and actions) and check test conditions based on both the packet headers and the flow state to invoke the relevant tables, shown by pink diamonds (P4’s flow control).

1. Extracting headers and options via flexible parsing: Using header definitions, we identify the relevant header fields in a packet. In our prototype we assume the TCP packets have Ethernet, IPv4, and TCP headers, although this can be easily extended to include other protocols (e.g., IPv6). The following snippet shows some relevant TCP headers. In addition, we provide a parser that extracts headers (e.g., source and destination IP from the IP header).

```
header_type tcp_t {
  fields {
    srcPort : 16;
    dstPort : 16;
    seqNo : 32;
    ackNo : 32;
    ...
  }
}
```

P4 models the parser as a state machine represented by a parse graph. The parsed headers need to be “de-parsed”, i.e., merged back, to a serial stream of bytes before forwarding. TCP options require TLV (Type-Length-Value) parsing. For parsing options, we use “masks” to identify the “type” (e.g., type 2 represent MSS), then a parser is called to extract that receiving each new ACK, as every ACK is a sign that a packet has left the network, hence the network can receive more.

option knowing its “length” (e.g., `parse_mss` for MSS in the snippet below) which returns the control back to the original parser when done. This creates a loop in the parsing graph causing the exact de-parsing behavior to be undefined. To solve it, we impose a fixed order for de-parsing, by using the *pragma* keyword as the following code snippet shows. We will use this extracted MSS value in the subsequent parts.

```
@pragma header_ordering ethernet ipv4 tcp options_mss
options_sack options_ts options_nop options_wscale
options_end

parser parse_tcp_options {
  return select(mymeta.opt_counter, current(0,8)) {
    ...
    0x0002 mask 0x00ff : parse_mss;
    ...
  }
}
```

2. Keeping per-flow state in registers: Registers are stateful memories, which are essential to Dapper because they maintain the per-flow state as it gets updated after processing each packet. Registers consume resources on the target, hence are a major limitation in running our solution on specific targets, therefore we will minimize the required per-flow state to ensure our program runs on commodity hardware in Section 5.

P4 registers can be global, referenced by any table, or static, bound to a single table. The following code shows one of our global registers, MSS, as an array of 16-bit values, and `instance_count` is the number of entries in the flow table. Each packet is hashed to find its flow index in the register array. We will explain our bi-directional flow hashing in more details shortly.

```
register MSS {
  width : 16;
  instance_count : ENTRIES;
}
```

When tracking a TCP connection for diagnosis, some register values depend on the value of other registers; for example, only by comparing a packet’s acknowledgment with previous ACKs of the flow can we detect a duplicate ACK. To update such dependent registers, we have to *read* other register(s), test *conditions*, and finally *update* the target register.

3. Carrying information per-packet via metadata: Metadata is the state associated with each packet, not necessarily derived from the packet headers, and can be used as temporary variables in the program. We use metadata to carry the information belonging to the same packet from one table to the other. The code below shows the most widely used metadata in our program, `flow_map_index`, which is the flow’s index produced by hashing. This metadata carries the index over to the subsequent tables, each using it to index their registers for read/write. Below, we show a code snippet for declaring metadata field named `flow_map_index` where `FLOW_MAP_SIZE` indicates its width in bits. In the next sub-

section, we explain how we use this metadata in hashing.

```
header_type stats_metadata_t {
  fields {
    flow_map_index : FLOW_MAP_SIZE; // flow’s map index
    ...
  }
}
metadata stats_metadata_t stats_metadata;
```

Some metadata has special significance to the operation of the target (i.e., the standard intrinsic metadata). In particular, we use the target’s `ingress_global_timestamp` as the arrival time of the packet, which is necessary to infer latency metrics such as the sender’s reaction time and SRTT.

4. Bi-directional hashing using metadata and registers:

As discussed earlier, our streaming algorithm must see both directions of traffic to capture our cross-packet metrics, e.g., application reaction time. To do this, we need to hash both directions to the same index and process them as one entity. Unfortunately, P4 provides no primitives or methods for hashing both directions to the same index—no symmetric hashes. Although some targets may allow configuring the hash function through run-time APIs, this support may vary across targets [3]. Therefore, we build our own symmetric hash using P4’s default hash algorithm, e.g., `crc32`, by defining two sets of headers to hash on, with one in the reverse order of the other. In other words, one direction is hashed based on (src IP, dst IP, src Port, dst Port) fields, and the reverse direction is hashed on (dst IP, src IP, dst Port, src Port) fields. To keep the direction’s hash function consistent, we use a simple and consistent comparison of the two IPs: if $srcIP > dstIP$, we hash the packet header in the former order, otherwise we hash the packet headers in latter order. This guarantees that each side of packet stream gets consistently hashed by one of these hash functions, but results in the same index value per flow.

5. Realizing operations using actions and tables: To realize the blue boxes in the flowchart of Figure 5, P4 tables and actions are used. A P4 table defines the fields to match on and the action(s) to take, should the entry match. P4 tables allow us to express different sets of match-action rules to apply on packets; for example, the set of actions for an outgoing packet differs from incoming packets. Furthermore, some tables could be dedicated to monitoring while others are dedicated to forwarding packets (e.g., `ipv4_lpm` and `forward`). A fundamental difference between our monitoring tables from regular forwarding tables in P4 is that our monitoring tables have a single static entry that matches on every packet—hence, have no match field. In contrast, `ipv4_lpm` is a forwarding table that uses longest prefix matching to find the next hop. The following code snippet shows two of Dapper’s tables, the lookup table, that hashes every packet to find its flow index, and the `init` table that initializes the flow upon observing its first packet (e.g., saves the extracted MSS value in the MSS register array, at the flow’s index).

```

table lookup{
  actions {
    lookup_flow_map;
  }
}
table init{
  actions {
    init_actions;
  }
}

```

```

action lookup_flow_map() {
  modify_field_with_hash_based_offset(stats_metadata.
    flow_map_index, 0, flow_map_hash, FLOW_MAP_SIZE);
}

action init_actions() {
  register_write(MSS, stats_metadata.flow_map_index,
    options_mss.MSS);
  ...
}

```

Actions in P4 are declared imperatively as functions, inside the tables. Actions can use registers, headers, and metadata to compute values. An example action is `register_write`, which takes a register, and index, and a value as input, and sets the value of the register array at the index accordingly. Actions are shown with red color in our code snippets.

6. Conditions via control-flow: The control flow of a P4 program specifies in what order the tables are to be applied. Inside the control segment, we can “apply” tables and test conditions. The choice of which block to execute may be determined by the actions performed on the packet in earlier stages. The control flow is what enables us to design the pipeline and implement the conditions (pink diamonds) in P4 as the flowchart shows in Figure 5.

In the “widely-supported” P4 specification [3], conditional operations are restricted to the control segments of program; that is, we cannot have `if-else` statements inside a table’s logic. Fortunately, P4 offers metadata, which can be used as temporary variables in the program. The metadata gives us an opportunity to read the current value of *conditional registers* inside an earlier table—the “loader”—in the pipeline, store their values in the metadata, test the conditions in the control section, and apply the appropriate set of tables conditionally. Note that we need the “loader” table because of the current restrictions in P4 that allows conditions only in the control segment.

```

control ingress {
  if ( ipv4.protocol == TCP_PROTO ) {
    if( ipv4.srcAddr > ipv4.dstAddr ) {
      apply(lookup);
    }else{
      apply(lookup_reverse);
    }
  }
  if ( (tcp.syn == 1) and (tcp.ack == 0) )//first pkt
    apply(init);
  else
    apply(loader);
  if (ipv4.srcAddr == stats_metadata.senderIP){
    if( tcp.seqNo > stats_metadata.seqNo ){
      apply(flow_sent);
      if(stats_metadata.sample_rtt_seq == 0)
        apply(sample_rtt_sent);//“temp” has the new
        flightsize
    }
  }
}

```

```

    if(stats_metadata.temp > stats_metadata.mincwnd)
      apply(increase_cwnd);
  }else{
    if(stats_metadata.dupack == DUP_ACK_CNT_RETX)
      apply(flow_retx_3dupack);
    else
      apply(flow_retx_timeout);
  }
}
else if(ipv4.dstAddr == stats_metadata.senderIP ) {
  if( tcp.ackNo > stats_metadata.ackNo ){
    apply(flow_rcvd);//new ack
    if( tcp.ackNo >= stats_metadata.sample_rtt_seq and
      stats_metadata.sample_rtt_seq>0){
      if(stats_metadata.rtt_samples ==0)
        apply(first_rtt_sample);
      else
        apply(sample_rtt_rcvd);
    }
  }else
    apply(flow_dupack);//duplicate ack
}
}
apply(ipv4_lpm);
apply(forward);
}

```

4.2 Hardware Resource Constraints

In this section, we explain our design choices to monitor connections in P4. These choices stem from a variety of restrictions, in particular, the limited resources on hardware switches, missing features in the P4 spec, and the diversity of hardware targets, which would require us to design for the least common denominator among the supported features.

1. Handling hash collisions: We use 32 bits for hashing in our prototype; regardless, collisions are often a concern in hash tables. In our software implementation, we handle collisions in the hash table by “hash-chaining”: we store the four tuple key of the connection and create a linked-list of flows in the same index with different keys. However, since memory in hardware is limited, we decide to not store a connection’s tuple. Still, if collisions go undetected they may pollute the accuracy of collected statistics. Hence, it is useful to assess if the accuracy of a connection’s statistics has been compromised. Therefore, we perform basic checks on the packet’s sequence number versus the flow’s previously sent sequence numbers and available windows (i.e., does the sequence number fall within the acceptable window?). This comparison requires additional tables or registers per-flow, but can store the result of conditions in a Boolean variable, named “sanity check”. The sanity check can be queried from the data plane along with the connection metrics to indicate whether collected statistics are reliable for diagnosis.

2. Keeping one RTT sample at a time: As discussed in Section 3, to accurately track the connection’s RTT, we maintain a queue of tuples based on the outgoing packets, where each tuple is (sequence number, time-stamp). The received ACKs are compared to the tuples of the queue to make an RTT measurement. Since the queue of tuples grows with the flow’s flight size, it increases the amount of state per-flow in our P4 program. Because the hardware resources on a switch

are limited, we limit the number of outstanding RTT tuples in our P4 program to one at a time, per-flow: we only sample an outgoing packet for RTT if the flow’s queue is empty.

3. Multiple accesses per register array: Our program accesses some registers from multiple points to use them in test conditions; e.g. duplicate ACK count register is read in the loader table and used for identifying the kind of loss, and upon a new duplicate ACK another table updates it. Currently, conditional operations in P4 are restricted to the control segments. Thus we cannot avoid accessing some registers in multiple tables. Unfortunately, accessing a register from multiple tables limits the processing rate. However, the ternary operator ($?:$) will be supported in the next P4 version [1], allowing us to perform simple conditional assignments, eliminating the need for global registers, and permitting our solution to run at line rate.

4. Relying on control plane to scale RWND: In TCP, the advertised RWND should be shifted by the window scale (as negotiated in handshake) to calculate the actual receive window. However, most P4 targets can only “shift” by a “constant” value [3]. So, we instead record both values and allow the control plane to query both and perform the shift.

5. Foregoing RTTvar: Calculating RTTvar involves capturing the absolute difference of the smoothed moving average RTT (SRTT) and the current RTT sample (RTT). This difference can be captured via an abs operator or by introducing a new comparison test, i.e., a new pipeline stage. Unfortunately, the P4 specification does not support the abs operator and adding a new stage impacts the processing rate of our implementation to gain a single metric. By default, Dapper does not include this stage but can be enabled optionally.

5. TWO-PHASE TCP MONITORING

Our goal in this section is to lower the cost of monitoring and diagnosing TCP connections. We present a two-phase monitoring technique to decrease the amount of state required. The first phase monitors all connections continuously but only collects low-overhead metrics, enough to *detect* but not *diagnose* performance problems. When a connection meets the “badness” criterion, heavier-weight monitoring is enabled to diagnose the poor performance.

Phase 1: Lightweight detection: In the first phase, we collect lightweight metrics that suffice to “detect” the existence of performance problems, based on a badness criterion. The statistics used in the first phase must be: 1. *lightweight* to maintain, ensuring that the continuous monitoring of connections in the data plane is cheap, and 2. *general* enough to capture the badness of the flow, regardless of the component limiting the performance. We use the *average rate of the flow* as an indicator of how well it’s performing.

To maintain the average rate of flow, we keep three registers: 1. *init_time*, the time-stamp when monitoring be-

gan, 2. *bytes_sent*, total bytes sent so far, and 3. *update_time*, the time-stamp when the flow was updated last. The control plane can query these states per-flow and find the connections that look troubled based on a cloud-wide tenant-specific threshold⁴. Upon observing low rate (or when tenant complains), the cloud provider can turn on the heavy-weight monitoring mode to do diagnosis.

Phase 2: Diagnosis of troubled connections: The second phase is “diagnosis”, where we collect heavyweight metrics for a troubled connection to shed light on the component that is hindering the flow’s performance. These metrics include the complete set of TCP statistics (discussed in Section 2) and our diagnosis techniques (discussed in section 3). This phase consumes more state on the switch, but note that it is only turned on after a problem is detected, hence the switch state consumption overall decreases. The two-phase monitoring can be thought of as a “long and narrow” table (all flows, few metrics), followed by a “short and fat” table (few flows, many metrics).

There are some challenges with inferring TCP metrics midstream: First, the TCP constants, in particular MSS and window scale are exchanged once, during the handshake. Second, the flow counters (e.g., packets sent, ACKed, or in flight) are unknown. This results in errors in the inferred value of flight size, which is also used to estimate *inf_cwnd*.

Our solution to these challenges is two fold: 1. Parse and keep the TCP options during the handshake for all the flows in the first phase, in case they are needed later. This approach provides high accuracy but requires more data-plane state. 2. Infer them midstream, only when necessary. Of course, inferring constants midstream reduces the memory overhead at the expense of accuracy. To infer MSS from midstream, we need to keep track of the largest segment sizes seen so far. To infer the window scaling option midstream we track the flight size and the unscaled RWND as advertised in received packets. We use the TCP invariant that “the flight size of a connection is limited by the receive window”, hence we can estimate the *lower-bound* window scaling factor:

$$\begin{aligned} \text{flight size} &\leq \text{RWND} \cdot 2^{\text{scale}} \\ \lceil \log_2 \frac{\text{flight size}}{\text{RWND}} \rceil &\leq \text{scale} \end{aligned} \quad (1)$$

Finally, to infer *inf_cwnd*, we rely on flight size, which itself is inferred from tracking sequence numbers of outgoing packets and incoming ACKs midstream. We will evaluate the accuracy of these metrics to show how close to actual values the accuracy of midstream inferred metrics get.

6. EVALUATION

In this section we evaluate the accuracy of our heuristics, and the overhead of our P4 prototype for hardware switches, and our C prototype for hypervisors (Section 6.2). Then

⁴If the tenant’s traffic is not crossing the WAN.

Figure 7: First 5 seconds of a flow shows how closely our heuristic tracks CWND, from monitoring starting points of (a) beginning, (b) 1 sec, (c) 3 sec

we use the software prototype to validate the accuracy of our diagnosis algorithm using synthetic traffic (Section 6.3). Next, we showcase Dapper in the wild by analyzing CAIDA packet traces (Section 6.4). Finally, we demonstrate the trade-offs in accuracy and overhead in the P4 design.

6.1 Accuracy of Heuristics

We begin the evaluation by comparing our heuristics against the ground truth data collected from within the VMs (using Linux kernel’s `tcp_info` structure). In particular, Figure 7 demonstrates how well the CWND is inferred for the first 5 seconds of a flow. The connection belongs to a simple server-client application where the server is transmitting a large file over a 1G link, both machines run TCP Cubic. We make three observations based on this figure: (1) Dapper’s inferred CWND closely follows Cubic’s. This suggests it is possible to accurately infer the congestion window without tenant’s cooperation. (2) If monitoring begins at later times (e.g., after 1 sec in (b) or 3 sec in (c)) Dapper infers CWND accurately, and the inferred CWND converges quickly. (3) We show how losses offer extra information for tracking CWND. Note, that while the figure only presents the first 5 second, our observations generalize to whole flow and to other flows in our dataset.

6.2 CPU and Memory Overhead

We evaluate Dapper’s overhead along two dimensions: memory utilization on hardware switches, and CPU Utilization on hypervisor or vswitch; this is mainly because switches operate at line rate but are constrained for memory, while software solutions are not often constrained by memory, but by CPU utilization. Dapper’s software prototype is implemented in C and uses `libpcap` to capture packets.

Memory in hardware: In single-phase mode, our P4 prototype keeps 67 bytes of state for each connection (i.e., 16 four-byte registers to keep the flow state, a two-byte register to track MSS, and a one-byte register for scale). In addition, 40 bytes of metadata are used to carry a packet’s information

Figure 8: Required state on switch in single vs two-phase monitoring, with 10% and 20% troubled connections

Figure 9: Expected collision rate vs number of flows (k) for different table sizes (N).

across tables. In a typical data center, a host can have 10K connections [40], which results in 670 KB of state required to track their state. In the two-phase monitoring prototype, as described in section 5, we use the average rate of flows as the badness factor in the first stage, which would impose about 16 bytes per flow in the first stage. Figure 8 shows the amount of state needed for single-phase monitoring, as opposed to two-phase monitoring with 10% and 20% troubled connections.

As more connections are monitored, the expected rate of collision in the hash table increases. Figure 9 shows the expected collision rate per number of flows (k), for varying sizes of tables (N). Assuming the N hash values are equally possible, the probability that a flow A shares the same index with flow B is $\frac{1}{N}$. So, the probability that the other k-1 flows will not share the same index is $(1 - \frac{1}{N})^{k-1}$, resulting in the expected likelihood of collisions of $1 - (1 - \frac{1}{N})^{k-1}$. Thus, to track 10K connections with a collision rate of less than 4%, we need the table size to be at least 262,144 (2^{18}), which results in less than 18 MB space on the switch.

CPU in software: We measure the CPU overhead by connecting two servers to a single switch and starting parallel

Figure 10: CPU per aggregate bandwidth processed

TCP flows between them and measuring CPU using `top`. The server machines have Xeon e3-1630 V3, 4 core, 3.4 GHZ processors. All flows are initially established (i.e., completed the TCP handshake) and have the average rate of 1 Mbps. Figure 10 shows the total CPU consumption (including live packet capturing via `libpcap` and copying packets to user space) versus the aggregate bandwidth processed. Dapper’s CPU consumption is close to approaches that use near-real time polling frequency (e.g., [40] at 50ms frequency). Batch-processing can lower the CPU overhead. In addition, the overhead in system calls and memory copies caused by `libpcap` for packet capturing can be reduced if Dapper uses fast packet IO frameworks such as `Netmap` [29].

In Figure 11, we quantify the CPU processing requirements for different packet types. The y-axis shows how many CPU cycles it takes after a packet is fully captured to update the flow state (excluding the packet capturing process). The first packet of a flow usually takes longer to process, because Dapper must allocate and initialize flow state and parse packet options that require extra CPU cycles. Further, the first outgoing packet is used for creating the first tuple for RTT measurement. The variations in the cycles are caused by several reasons: First, our software prototype’s flowtable is a hash table with hash-chaining, thus, in case of a collision, the flow statistics are maintained in a linked-list; incurring the extra overhead of linked-list traversal for collisions. Second, for retransmitted packets or duplicate ACKs, the flow state needs further updates. Finally, packets used for RTT measurement incur extra processing.

6.3 Diagnosis Accuracy

We measure the accuracy of our diagnosis method by systematically creating TCP connections with known problems and comparing the diagnosis results with the ground truth. We create a server-client connection with the client requesting a 1MB file over a link with 1Mbps bandwidth and 50ms round-trip time. We then create the following problems in the connection and repeat the experiments 100 times:

1. *Sender-Limited*: We emulate a resource bottleneck for the server, e.g., slow disk or busy CPU, by making the server

Figure 11: CPU cycles to update the table based on each type of packet, excluding packet capturing and handling

Figure 12: Accuracy vs severeness of problem

wait for T seconds before transmitting each data packet. Higher T indicates more severe problems. We also emulate non-backlogged servers by limiting the transmitted segment sizes to less than an MSS.

2. *Receiver-Limited*: We create receiver-limited connections by changing socket options (using Linux’s `setsockopt`) to limit the client’s receive buffer size, `socket.SO_RCVBUF`.

3. *Network-Limited*: We use the Gilbert-Elliott model [20] to emulate micro-bursts during network congestion: a connection can be in either a *good* (no network congestion) or *bad* (network congestion) state. To emulate the bad state, we generate bursty losses at a rate of 1% to 10% for 2 seconds. We assume that losses in the good state are negligible.

For each problem shown in Table 2, we measure “sensitivity” and “accuracy”: the *true positive rate* (TPR) is the fraction of correctly diagnosed tests, showing the sensitivity to each problem, and *diagnosis accuracy* is the fraction of time Dapper correctly classified the connection to be limited by that problem. The initial results are promising, Dapper achieves an average accuracy of 94%; Dapper’s accuracy is less than 100% because the ability to detect a problem is proportional to its severity. Figure 12 shows an example problem where the severity changes from low (1% loss) to high (10% loss), increasing the average accuracy of diagnosis.

Table 2: Dapper’s Diagnosis Sensitivity and Accuracy

Problem	TPR	Avg Accuracy
Sender-Limited	98%	95%
Receiver-Limited	96%	94%
Network-Limited	94%	93%
Sender-Network-Limited	100%	95%
Receiver-Network-Limited	100%	93%

Figure 13: Diagnosis Results for CAIDA traces

6.4 Analyzing CAIDA Traces

In the absence of traces from a public cloud, we use CAIDA traces collected on Equinix-Chicago [5] to showcase Dapper, by assuming that the applications that produced this traffic have now moved to the cloud. We preprocess the trace by eliminating connections with less than 10 packets in either direction. A total of 244,185 flows are considered in this study. After preprocessing, the traces are treated as a live stream of packets, with no replay. To account for the fact that the packet traces are collected from within the network, rather than the edge, we turn off inferences based on application reaction time because it cannot be reliably captured in the core of the network.

Figure 13 shows the CDF of the fraction of time each type of performance problems limits a connection. Since the bottlenecks in a connection may change over time, the CDF shows the normalized duration of bottlenecks. We observe that: 1. 99% of flows spend less than 1% of their life in the no-limit state; this means that 99% of flows have at least one bottleneck in more than 99% of their lifetime. 2. Although sender and receiver problems have similar rates, we did not use the application reaction time to detect the sender-limited connections, thus we expect the actual rate of sender-limited connections to be higher. 3. Many connections are bottlenecked by two factors simultaneously (e.g., sender-network or receiver-network). 4. About 90% of the connections spend some time in the network-limited state, with almost half of them being network-limited 50% of the time.⁵

6.5 Trade-offs in Accuracy and Overhead

⁵Note that these results are from wide-area network and the characteristics of data-center connections are different from WAN.

Figure 14: Error in inferring SRTT with queue size of 1. As more packets are exchanged, the error decreases.

Using the CAIDA trace, we evaluate the impact of our space optimizations on the accuracy of our measurement. To do this, we operate our software prototype in two modes: first, trades-off space for accuracy (emulate P4) and second, measures with high accuracy and precision (ground-truth).

Limited queue size: Our measurement shows that running Dapper with an unbounded queue increases memory usage by about 9%. In a network with higher bandwidth capacity, (i.e., high bandwidth delay product), more memory will be required. In Figure 14, we examine the error in SRTT when the queue size is bounded to one. We observe that when the queue size is bounded, Dapper requires more samples to reduce error. Note these traces were collected in WAN, not a datacenter, hence the high RTTs.

Two-phase monitoring: Recall, two-phase monitoring offers a trade-off between accuracy of our heuristics and their memory overhead. We filter the connections to consider only those with known options, and use that as the ground-truth to compare against the midstream inferred constants. Figure 15 shows the error in inferring MSS and window scaling options according to Section 5. While the error in MSS decreases as more packets are inspected, the inferred window scale value still differs from the ground-truth by about 20%. This is because most of these connections are not RWND-limited, hence the flight size values used in estimating the scale (Equation 5) do not approach the upper-bound imposed by the receiver.

7. DISCUSSION

In this section, we discuss the implications of our diagnosis granularity and Dapper’s hardware implementation.

(1) *Diagnosis Granularity:* Dapper identifies the component causing the performance problem in a connection (network, sender, or receiver), but does not find the root cause within that component. As part of future work, we plan to explore domain specific inference algorithms for each entity.

(2) *Hardware capabilities:* A major issue in our P4 implementation arises from accessing the same registers at mul-

Figure 15: Error in inferring options (MSS and wscale) midstream, the error rate decreases faster for MSS.

multiple stages, which hurts the line-rate performance at the switch. Fortunately, as discussed in Section 4, future versions of P4 will support the ternary operator, allowing us to effectively restrict access to registers to one stage.

(3) *Hardware capacity*: Switches have a limited number of registers and as discussed in Section 6, this capacity limitation increases the collision rate and ultimately reduces the monitoring accuracy of Dapper. To address this limitation, we propose two solutions: (i) Dapper’s software equivalent (which uses `libpcap` and `BPF`) can be used in the hypervisor in public clouds. The software environment will have much more memory for storing, however, the measurement process will consume the end-host’s CPU instead of the edge switch. Some of the CPU cost could be mitigated via batch-processing [26] and fast packet I/O such as `netmap` [29]. (ii) Sampling and triggering, as discussed in Section 5, can help us cope with the limited capacity.

Monitoring in the network removes packet capturing and processing overhead from the end hosts, and enables the network to adapt in real time to the diagnosis results. Implementing Dapper in the dataplane via P4 introduces limitations that a purely software solution may avoid. Yet, P4 embodies some of the capabilities at the edge (e.g. stateful programming) and provides a realistic environment. We believe that the modest overheads are an acceptable price to pay for the flexibility of automatically embedding Dapper in a wide range of existing edge architectures.

8. RELATED WORK

We have categorized the existing related work to Dapper:

Offline packet trace analysis: Several tools analyze packet traces to find performance limitation [41, 10, 4] for a known TCP variant. Some [39] store packet headers that facilitate diagnosis via running queries on headers. However, *offline* analysis makes this category unsuitable for real-time diagnosis and introduces large data-collection overhead.

Measurement in the core: Some tools use coarse-grained metrics collected on switches, e.g., 5-minute SNMP coun-

ters [38]. Such metrics are not sufficiently fine-grained to diagnose the sources of poor performance [9]; others [25] focus on coordination between network switches to ensure the full packet stream is analyzed. Instead, we rely on the “edge” for better visibility and a simpler solution.

Instrumenting the network stack: Existing end-host techniques [34, 35, 6, 40], while appropriate for private clouds, are too invasive for a public IaaS cloud because they would run in the tenant’s VM. Trumpet [26], similar to Dapper, tracks packets at the hypervisor with the goal of providing a network-wide event monitoring system, but would still rely on the network operators to define events (via filters and predicates) for diagnosis.

Tomography: Network tomography infers link-level properties (e.g., loss, delay) from end-to-end measurements [15, 27, 14, 28, 31, 17] and may use linear [16, 19] or Boolean algebra [19, 37, 24] to find congested links. Our work differs from this body of work since we rely on *direct and continuous measurement* of performance instead. Tomography can be complementary to Dapper, e.g., to locate the congested link(s) among network-limited connections.

Machine learning: Recent work [8] makes the case for using machine learning (ML) in connection’s performance diagnosis. Our solution differs in several aspects: first, instead of relying on ML, we leverage the fact that TCP is a well studied protocol and derive techniques based on common characteristics of TCP and its interaction with applications. Second, our solution is agnostic to application and TCP variant, while the ML would require new training for each variant.

Congestion control from the “edge”: To enable new congestion control algorithms in multi-tenant clouds, [18] offers virtualized congestion control in the hypervisor. Similarly, [21] enforces congestion control in the vswitch without cooperation from the tenant VM.

9. CONCLUSION

Dapper helps cloud providers diagnose TCP connections at line rate without instrumenting the tenant VMs. We use TCP domain knowledge to draw inferences about the internal state of the connections to find if a connection is limited by the sender, the network, or the receiver.

10. ACKNOWLEDGMENTS

This work was supported by NSF CNS-1162112, NSF CNS-1409426, and a gift from Intel. We thank our shepherd, Luigi Rizzo, and the anonymous reviewers for their thoughtful feedback; Changhoon Kim for helpful discussions; Georgios Nikolaidis and Rob Harrison for feedback on earlier drafts.

11. REFERENCES

- [1] Changhoon Kim, Co-chair of P4 Language Design Working Group, personal communication.
- [2] Nmap. <http://www.nmap.org>.

- [3] The P4 language specification, version 1.0.2. <http://p4.org/wp-content/uploads/2015/04/p4-latest.pdf>.
- [4] tcptrace. <http://tcptrace.org/>.
- [5] The CAIDA UCSD 2013 internet traces - 2013/05/29. <http://data.caida.org/datasets/passive-2013>.
- [6] The Web10G project. <http://www.web10g.org>.
- [7] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan. Data Center TCP (DCTCP). In *Proc. ACM SIGCOMM*, 2010.
- [8] B. Arzani, S. Ciraci, B. T. Loo, A. Schuster, and G. Outhred. Taking the blame game out of data centers operations with netpoirrot. In *Proc. ACM SIGCOMM*, 2016.
- [9] P. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. A. Maltz, and M. Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *Proc. ACM SIGCOMM*, 2007.
- [10] P. Barford and M. Crovella. Critical path analysis of TCP transactions. In *Proc. ACM SIGCOMM*, 2000.
- [11] G. Bianchi, M. Bonola, A. Capone, and C. Cascone. Openstate: Programming platform-independent stateful openflow applications inside the switch. *ACM SIGCOMM CCR*, Apr. 2014.
- [12] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker. P4: Programming protocol-independent packet processors. *ACM SIGCOMM CCR*, 2014.
- [13] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz. Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN. In *Proc. ACM SIGCOMM*, 2013.
- [14] R. Caceres, N. Duffield, S. Moon, and D. Towsley. Inference of internal loss rates in the mbone. In *Proc. GLOBECOM*, 1999.
- [15] R. Caceres, N. G. Duffield, J. Horowitz, and D. F. Towsley. Multicast-based inference of network-internal loss characteristics. *IEEE Transactions on Information Theory*, 1999.
- [16] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu. Network tomography: recent developments. *Statistical Science*, 2004.
- [17] M. Coates, R. Castro, R. Nowak, M. Gadhiok, R. King, and Y. Tsang. Maximum likelihood network topology identification from edge-based unicast measurements. In *Proc. ACM SIGMETRICS*, 2002.
- [18] B. Cronkite-Ratcliff, A. Bergman, S. Vargaftik, M. Ravi, N. McKeown, I. Abraham, and I. Keslassy. Virtualized congestion control. In *Proc. ACM SIGCOMM*, 2016.
- [19] N. Duffield. Network tomography of binary network performance characteristics. *IEEE Trans. on Information Theory*, 2006.
- [20] G. Hasslinger and O. Hohlfeld. The gilbert-elliott model for packet loss in real time services on the internet. In *Measuring, Modelling and Evaluation of Computer and Communication Systems (MMB)*, March 2008.
- [21] K. He, E. Rozner, K. Agarwal, Y. J. Gu, W. Felter, J. Carter, and A. Akella. AC/DC TCP: Virtual congestion control enforcement for datacenter networks. In *Proc. ACM SIGCOMM*, 2016.
- [22] P. Karn and C. Partridge. Improving round-trip time estimates in reliable transport protocols. In *Proc. of the ACM Workshop on Frontiers in Computer Communications Technology*, SIGCOMM, 1988.
- [23] C. Kim, P. Bhide, E. Doe, H. Holbrook, A. Ghanwani, D. Daly, M. Hira, and B. Davie. In-band network telemetry (int). <http://p4.org/wp-content/uploads/fixed/INT/INT-current-spec.pdf>.
- [24] R. Kompella, J. Yates, A. Greenberg, and A. Snoeren. Detection and localization of network black holes. In *IEEE INFOCOM*, 2007.
- [25] X. Liu, M. Shirazipour, M. Yu, and Y. Zhang. Mozart: Temporal coordination of measurement. In *Proc. ACM SOSR*, 2016.
- [26] M. Moshref, M. Yu, R. Govindan, and A. Vahdat. Trumpet: Timely and precise triggers in data centers. In *Proc. of ACM SIGCOMM*, 2016.
- [27] H. X. Nguyen and P. Thiran. Binary versus analogue path monitoring in IP networks. In *Proc. PAM*, 2005.
- [28] H. X. Nguyen and P. Thiran. Network loss inference with second order statistics of end-to-end flows. In *Proc. ACM IMC*, 2007.
- [29] L. Rizzo. netmap: A novel framework for fast packet i/o. In *USENIX ATC*, 2012.
- [30] M. Shahbaz, S. Choi, B. Pfaff, C. Kim, N. Feamster, N. McKeown, and J. Rexford. Pisces: A programmable, protocol-independent software switch. In *Proc. ACM SIGCOMM*, 2016.
- [31] Y. Shavitt, X. Sun, A. Wool, and B. Yener. Computing the unmeasured: an algebraic approach to internet mapping. In *Proc. IEEE INFOCOM*, 2001.
- [32] A. Sivaraman, A. Cheung, M. Budiu, C. Kim, M. Alizadeh, H. Balakrishnan, G. Varghese, N. McKeown, and S. Licking. Packet transactions: High-level programming for line-rate switches. In *Proc. of ACM SIGCOMM*, 2016.
- [33] H. Song. Protocol-oblivious forwarding: Unleash the power of sdn through a future-proof forwarding plane. In *Proc. of ACM SIGCOMM HotSDN*, 2013.
- [34] P. Sun, M. Yu, M. J. Freedman, and J. Rexford. Identifying performance bottlenecks in CDNs through TCP-level monitoring. In *ACM SIGCOMM Workshop on Measurements Up the Stack*, 2011.
- [35] P. Sun, M. Yu, M. J. Freedman, J. Rexford, and D. Walker. HONE: Joint host-network traffic management in software-defined networks. *Journal of Network and Service Management*, 2015.
- [36] T. Tofigh and N. Viljoen. Dynamic analytics for programmable NIC's utilizing p4 - identification and custom tagging of elastic telecoms traffic. <http://p4.org/wp-content/uploads/2016/06/P4-Poster-Netronome-ATT.pdf>.
- [37] W. Wei, B. Wang, D. Towsley, and J. Kurose. Model-based identification of dominant congested links. In *Proc. of ACM IMC*, 2003.
- [38] L. Wu, D. Turner, C.-C. Chen, D. A. Maltz, X. Yang, L. Yuan, and M. Zhang. Netpilot: Automating datacenter network failure mitigation. *ACM SIGCOMM CCR*, 2012.
- [39] W. Wu, G. Wang, A. Akella, and A. Shaikh. Virtual network diagnosis as a service. In *Symposium on Cloud Computing*, 2013.
- [40] M. Yu, A. Greenberg, D. Maltz, J. Rexford, L. Yuan, S. Kandula, and C. Kim. Profiling network performance for multi-tier data center applications. In *Proc. NSDI*, 2011.
- [41] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. On the characteristics and origins of Internet flow rates. In *Proc. ACM SIGCOMM*, 2002.