

COP²: Continuously Observing Protocol Performance

Olivier Tilmans*

Université catholique de Louvain
olivier.tilmans@uclouvain.be

Olivier Bonaventure

Université catholique de Louvain
olivier.bonaventure@uclouvain.be

Abstract

As enterprises move to a cloud-first approach, their network becomes crucial to their daily operations and has to be continuously monitored. Although passive monitoring can be convenient from a deployment viewpoint, inferring the state of each connection can cause them to miss important information (e.g., starvation). Furthermore, the increasing usage of fully encrypted protocols (e.g., QUIC encrypts headers), possibly over multiple paths (e.g., MPTCP), keeps diminishing the applicability of such techniques to future networks.

We propose a new monitoring framework, Flowcorder, which leverages information already maintained by the end-hosts and records Key Performance Indicators (KPIs) from their transport protocols. More specifically, we present a generic approach which inserts lightweight eBPF probes at runtime in the protocol implementations. These probes extract KPIs from the per-connection states, and eventually export them over IPFIX for analysis.

We present an application of this technique to the Linux kernel TCP stack and demonstrate its generality by extending it to support MPTCP. Our performance evaluation confirms that its overhead is negligible. Finally, we present live measurements collected with Flowcorder in a campus network, highlighting some insights provided by our framework.

1 Introduction

Network performance depends on a variety of factors such as link delays and bandwidth, router buffers, routing or transport protocols. Some of these are controlled by the network operators, others by the end-hosts. To detect potential issues, and ensure their proper operations, most network operators monitor a wide range of statistics on the health of their networks, which can be classified in three categories. First, health metrics capture the status of network elements. Most networks record those using SNMP, polling their devices every few minutes to collect various statistics (e.g., link load, CPU usage, size of forwarding tables). Operators often also collect statistics about the traffic itself, usually using NetFlow/IPFIX [38, 67, 75]. These provide more detailed information about the flows crossing the network (e.g., layer-4 5-tuples, volumes in bytes and packet), and enable various management applications [52] (e.g., identifying major source/destination pairs [83], heavy-hitters [31], or detecting DDoS attacks [69, 79]). Finally, operators monitor key performance metrics which are important for many end-to-end

applications, such as delays, packet losses, and retransmissions. On one hand, active measurements techniques [16, 53] collect these metrics by generating test traffic (e.g., pings). On the other hand, passive measurements [28, 46] infer these performance metrics by analyzing the packets that traverse the network (e.g., using network taps which maintain per-flow states to accurately measure Round-Trip-Times (RTT), retransmissions, packet losses and duplications [54]).

Although widely deployed, passive monitoring suffers from several important limitations. First, as link speeds increase, it becomes more and more difficult to maintain the per-flow state that is required to collect detailed performance metrics [76]. Second, as multipath protocol deployment increases (e.g., MPTCP [29] is used in iPhones [3] and for other services [9]), passive monitors only see a subset of the packets belonging to a connection. This compromises their ability to operate properly [61]. Finally, the most important threat against the passive collection of network performance metrics is the deployment of encrypted protocols, such as QUIC [50]. QUIC replaces the HTTP/TLS/TCP stack with a simpler protocol that runs over UDP. Google estimates [50] that QUIC already represents more than 7% of the total Internet traffic. Recent measurements indicate that content providers have started to deploy QUIC massively [66]. The IETF is currently finalizing a standardized version of QUIC [42].

From a performance monitoring viewpoint, an important feature of QUIC is that all the payload and most of the header of the packets are encrypted. This prevents the middlebox ossification problems that affect protocols such as TCP [39, 60], but it also greatly decreases the ability for network operators to monitor network performance. This prompted some of them to ask to modify QUIC to be able to extract performance information from its headers [70]. The IETF answered those operational concerns by reserving one bit in the QUIC header (the spin-bit [74]), exposing limited delay information. Multipath extensions to QUIC have already been proposed [20, 80].

To keep collecting end-to-end performance metrics of their users flows, enterprise network operators need a different approach than passive monitoring to be future proof.

Problem statement *How can we support the legitimate need of fine grained performance information from enterprise network operators in presence of encrypted, multipath protocols?*

Key challenges Designing a monitoring framework that answers this question raises at least four challenges. First, this framework must accurately depict the performance experienced by the end-hosts. This limits the applicability of active measurements, as this might hide issues specific to

*O. Tilmans is supported by a grant from F.R.S.-FNRS FRIA

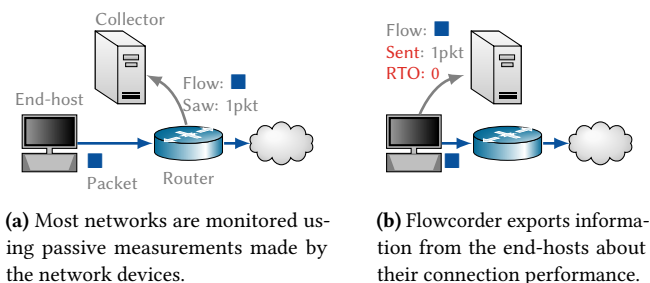


Figure 1. Flowcorder enhances network monitoring with fine-grained measurements about connections.

the used protocol (e.g., TCP RTO). Second, it must support multipath protocols, and thus monitor the performance of all paths used by a given connection. This limits the possibility of using passive monitoring since this would require coordination among the monitors located on different paths. Third, supporting encrypted protocols prohibits such framework from analysing packet headers or contents and prohibits the utilisation of “transparent” proxies. Finally, it should operate with a low overhead, limiting the generated statistics to the minimum to establish a baseline for normal operation, while also enabling to quickly capture and detect performance issues.

Flowcorder We introduce Flowcorder, a novel enterprise network monitoring framework which addresses the above challenges. The key insight behind Flowcorder is to leverage the per-connection information that is already maintained by the end-hosts themselves.

Instrumenting the transport stacks of the end-hosts enables Flowcorder to compute Key performance Indicators (KPIs) for each connection. By capturing such KPIs at specific moments of the connection life-cycle, Flowcorder can then build *performance profiles* of connections. Finally, Flowcorder aggregates those profiles and exports them over IPFIX, integrating with existing monitoring infrastructure and enabling analyzes across hosts, protocols, remote services and/or ISPs.

Contributions Our main contributions are:

- A novel enterprise monitoring framework to monitor the network performance experienced by the end-hosts (§2).
- A generic approach to export performance profiles of connections by transparently extracting KPIs from existing protocol implementations (§3).
- An application of the approach to realize an event-based instrumentation of the Linux kernel TCP stack (§4), a demonstration of the generality of the approach by extending it to support MPTCP (§4.3), and an evaluation showing its low-overhead (§5).
- A case study highlighting insights provided by Flowcorder when deployed in a campus network (§6).

2 Flowcorder

Many networks monitor their traffic using in-network appliances that inspect packets crossing them, and eventually export statistics to measurement collectors using a protocol such as IPFIX (Fig. 1a). While sufficient to track traffic demands, or collect rough traffic statistics through passive inference of the connection states, these techniques hardly scale if the operators requires fine-grained performance measurements on a per-connection basis. Flowcorder instead pushes the monitoring processes directly on the end-hosts (Fig. 1b). By monitoring the per-connection states, Flowcorder can then record the performance of the connections, as experienced by the end-users, and then export those over IPFIX to complement existing measurement infrastructure. The rest of this section illustrates the different building blocks making up Flowcorder, visible on Fig. 2. More specifically, we consider a network administrator who wants to use Flowcorder to answer the following a high-level question: “Which provider performs the best to connect to a remote storage service accessed over TCP?”

Computing performance profiles. The first step to answer this high-level question is to identify KPIs (§3.1) that enable to characterize the performance of the instrumented protocol. Such KPIs should contain general statistics about the connection, as well as metrics indicating possible performance issues, specific to the protocol.

For example, high-level KPIs to answer our illustrative question could be: (i) the number of bytes transferred and assumed to be lost; (ii) the amount of reordering [5, 8, 45] that occurred in the network; and (iii) signs of bufferbloat, such as the number bytes received multiple times, thus signaling a retransmission timeout on the sender, or times where the connection stalled and was blocked from sending pending data for several RTTs (TCP RTO).

Continuously streaming the collected KPIs is inefficient as, beside wasting resources, it might hide the key performance outliers in the noise generated by the huge number of smaller variations. Instead, Flowcorder exports the KPIs of a connection only at specific moments in the connection life-cycle (§3.3). In-between these exports, the KPIs are buffered in a lightweight aggregation daemon, local to the end-host. Once the decision to export the measurement is made, this aggregation daemon computes a *performance profile* of the connection: statistics computed over KPIs (e.g., moving averages, counter increase) during well-defined moments of the connection life-cycle. The performance profile is then serialized as an IPFIX record and added in a pending IPFIX message buffer. As we want to minimize the processing load on the collector and take advantage of the features provided by IPFIX, the message is only exported once its size reaches the local MTU.

In our example, a connection towards the remote storage service that would experience one retransmission timeout

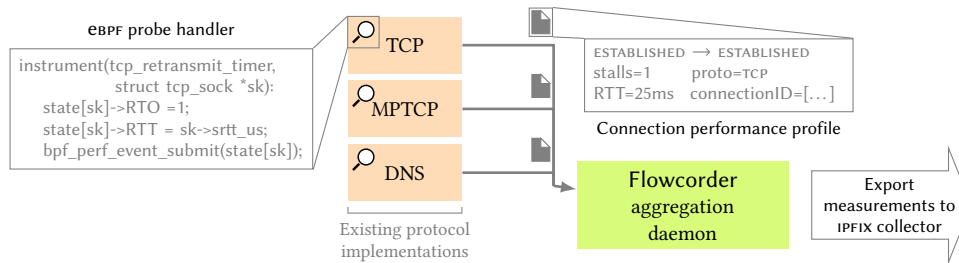


Figure 2. Flowcorder enables to evaluate network performance from generic Key Performance Indicators collected on the end-hosts for every connection.

in its entire life-cycle would generate four performance profiles: (i) one describing the connection establishment; (ii) one describing the performance of the data transfer (e.g., average RTT, byte counters, number of RTO experienced) up to the RTO; (iii) one describing the performance while the connection is considered as lossy; and (iv) a final one describing the performance since the end of the lossy state and how the connection ended (e.g., did it abruptly end with a TCP RST?).

Collecting KPIs. Under the hood, Flowcorder instruments existing transport protocol implementations on the end-hosts. Many methods exist to collect such statistics, such as extracting them from a general purpose loggers [55, 63] or polling [14]. Instead, Flowcorder uses an event-based method. More specifically, Flowcorder inserts eBPF probes at specific code paths in the transport protocol implementations (§3.2). When the end-host stack reaches one of these probes, the probe handler is executed, computes KPIs of the connection, exports them in an asynchronous channel to the aggregation daemon, and then resumes the normal execution of the protocol implementation. Beside minimizing the instrumentation overhead (§5), this approach is also extremely flexible as it does not require any support from the implementation (e.g., MIBs), and is thus not restricted to a predefined set of metrics, computed in an opaque manner.

In the example of Fig. 2, we see that one such probe has been setup to intercept the expiration of the TCP retransmission timer. If any connection experiences a RTO, this handler then increases the KPI counting RTO’s and updates the connection’s RTT estimated by TCP, then exports it for processing in user-space.

Analyzing performance profiles. Flowcorder produces measurements that can be collected, parsed and analyzed by any IPFIX collector supporting custom Information Elements [10]. Performance profiles are independent views of the performance of a connection during a given window of time, and one can be analyzed separately from the others belonging to the same connection. These performance profiles thus enable the network operator to build several views of the network according to key metrics using simple database queries, and to analyze them (§6).

Table 1. Key Performance Indicators can answer most questions about transport protocol performance

KPI	Description
Σ^* Sent	Data [†] sent towards the remote host
Σ Received	Data received and processed by the end host
Σ Lost	Data assumed to be lost in the network
Σ Errors	Data received corrupted
\mathcal{A}^{\ddagger} RTT	Mean Round-Trip-Time and variance (i.e., jitter)
Σ Duplicates	Received data already acknowledged
Σ OFO	Data received out-of-order
\mathcal{A} OFO-dist	Distance of out-of-order data from the expected one
Σ Stalls	Count when the connection delays the sending of any pending data during several RTTs

* Σ denotes a counter over a time window
[†]Most KPIs can be duplicated to track byte-counts and packets ('data')
[‡] \mathcal{A} denotes an average and a variance over a time window

For example, to answer his question, our network administrator could compute generic statistics such as mean, variance and median of all performance profiles contained in a given time window, aggregated by provider, and run hypothesis tests. These results could also be split based on the IP version, or compared against the general trend to access all other remote services. Finally, beside numerical tests, one can also generate time series and plot them in monitoring dashboards.

3 Recording protocol performance

Flowcorder records *performance profiles* of connections directly on end-hosts, and exports them to a collector for further analysis. Achieving this requires addressing three issues: (i) What should a performance profile contain to describe a connection and indicate performance issues (§3.1)?; (ii) How can we collect these key metrics from the protocol implementations?; and (iii) When should these profiles be computed to maximize the accuracy of the measurements while minimizing the overhead of Flowcorder (§3.3)?

3.1 Characterizing protocol performance

Connection-oriented transport protocols such as TCP maintain state and usually expose some debugging information (e.g. `struct tcp_info` [49] on Linux or macOS). However, recording the entire state for each established connection is impractical. Most of this information is very specific to the protocol implementation and does not always relate to connection performance. For example, one can find the distance (in terms of TCP segments) between the last out-of-order segment and the expected sequence number or the value of the slow-start threshold in the `struct tcp_info`, both of which give almost no insight to qualify the connection performance. Finally, while Flowcorder aims to collect fine-grained measurements about protocol performance as experienced by the end-hosts, recording every single data point would be counter-productive, as the more critical observations will end up buried in a huge pile of data.

Instead, we characterize protocol performance by recording the evolution of Key Performance Indicators (KPIs) during a connection. Example KPIs are listed in Table 1. Recording Sent and Received bytes quantifies the volume transported on a connection, while tracking the number of segments quantifies the packet rate (e.g., an interactive SSH session produces many small TCP segments). Recording Lost segments or segments with a checksum error (Errors), enables to qualify the path used by the connection. Tracking the evolution of the RTT (and thus implicitly its jitter) can be used to estimate whether congestion is building up in the network (and is the main source of information of some congestion control algorithms such as BBR [13]). Similarly, recording the reception of segments containing already acknowledged data is an indication that the remote host mistakenly assumed their loss, which could be a sign of a possible bufferbloat. Measuring the amount of packet reordering is also useful, especially in the context of transport protocols, as its occurrence often limits the maximum achievable throughput. Finally, recording when a connection is prevented from making progress is a strong signal that something bad happened in the network (e.g., triggering a TCP RTO).

From these KPIs, network administrators can then answer complex high-level questions characterising the performance of the network, such as: (i) what is the best response time that can be expected when connecting to a remote server?; (ii) Is the connection suitable for bulk transfers?; or (iii) Is the network congested?

3.2 Collecting KPIs from implementations

Recording the evolution of the KPIs of a connection on the end-hosts requires to extract them directly from the protocol implementation. Achieving this is usually possible using poll-based techniques. For example, SNMP can be used to query the TCP Management Information Base (MIB) [64]. Some OS'es also define APIs to retrieve information [2, 49], or log events to a centralized journal [55] which can then be monitored.

These techniques however come with two limitations. First, the information they give is limited to the explicitly defined metrics. For example, counting TCP out-of-order packets, as well as characterizing their out-of-order distance is impossible on Linux with the existing API. Counting received duplicates is not feasible either. Second, by requiring the monitoring tool to poll them, getting more accurate information about performance changes imposes a polling frequency and thus a high resource usage on the end-hosts. For example, characterizing the connection establishment times requires to precisely track the first few packets of a connection, which could be exchanged within a few milliseconds.

To address these issues, Flowcorder bypasses these traditional techniques, and directly instruments the protocol implementation at runtime.

Dynamic tracing using eBPF. Flowcorder leverages the existing dynamic tracing tools such as *kernel probes* [35], or DTrace [6]. These enable to insert lightweight probes at runtime at arbitrary locations in either kernel (e.g., to instrument the TCP kernel implementation §4) or user-space code (e.g., to instrument DNS resolution routines, for which we present collected measurements in §6), typically around function calls. Conceptually similar to breakpoints and debugging watches, these probes automatically call user-defined handlers before and after executing the probed instruction. These handlers have complete access to the memory, as well as to the content of the CPU registers. More recently, the Linux kernel added code to define such handlers using extended Berkeley Packet Filters (eBPF) [43].

eBPF code is pre-loaded in the kernel using the `bpf()` system call. This eBPF code is executed in an in-kernel virtual machine that mimics a RISC 64-bits CPU architecture, with 11 registers and a 512 bytes stack. This code can be interpreted, but many architectures include a JIT that compiles the eBPF bytecode. Before accepting to load an eBPF code, a verifier ensures safety guarantees such as proof of termination (e.g., by limiting the overall number of instructions and disallowing non-unrollable loops) and checks memory-access. eBPF code executed within the kernel can asynchronously communicate with user-space processes using perf events (FIFO queues). Additionally, eBPF programs can define maps, which let them maintain state in-between executions. When an eBPF probe handler is executed, it receives an instance of the `struct pt_regs`, which describes the content of the CPU registers when the probe was hit, including the value of the stack pointer. This enables the eBPF handler to inspect the function arguments, or to explore the memory of the instrumented code. These capabilities make eBPF a target of choice to write probe handlers, as they guarantee that the handlers will not cause crashes nor hang the instrumented code, while also enabling it to compute complex statistics and easily report them to user-space.




Figure 3. Transport protocol flows can be abstracted in a general FSM, where state-transitions act as signal to compute performance profiles.

This approach has at least five advantages. First, by leveraging state transitions that are internal to the implementation, it ensures an accurate translation to KPIs. For example, by recording retransmission timer expirations, it easily distinguishes between a connection that had no data to send for a while and a connection that was stalled and had to wait a complete RTO before sending anything else. Second, it seamlessly adapts to settings local to the host — for example, the TCP duplicate ACK threshold, or the support for SACK on a per connection basis — that alter the behavior of the transport protocol. As such, it accurately captures the performance experienced by all instrumented end-hosts. Third, as it implements a push-based model where the transport stack itself calls Flowcorder, it minimizes the overhead on the end-hosts. Indeed, as the probe locations guarantee that all KPI changes will be detected, this avoids the need for constant, high-frequency, polling of the state-variables. Fourth, as it enables to both read per-connection states and to compute arbitrary statistics that can be stored in maps (thus defining custom ancillary state), this approach is highly flexible, as it does not rely on specific support from the protocol implementation. Finally, it could also be applied to encrypted transport protocols such as QUIC since it does not use the packet data but instead the state-variables of the protocol implementation.

3.3 Creating performance profiles

To use dynamic tracing and eBPF handlers to instrument a particular transport protocol, one needs to pick probe insertion locations to catch updates to the state of a connection. While a straw-man approach would pick the main functions involved in every send and receive operation, and continuously stream the connection KPIs after each sent and received packet, this would impose a high overhead without necessarily providing useful measurements. Indeed, once the probes are inserted, their handlers are executed for every connection hitting that code path. Instead, we aim at recording the evolution of KPIs between key events in the connection life-cycle. To this end, we place probes at locations that are seldom reached, yet catch all important events affecting the connection, and record statistics describing the evolution of

the KPIs between two events. We call such set of statistics the *performance profile* of a connection.

A first set of events are defined by the protocol specifications. Such specification is usually composed of two different parts. The first is the syntax of the protocol messages, which can be expressed informally with packet descriptions or more formally by using a grammar (e.g., ABNF [19], ASN.1 [44]). The second part of the specification describes how and when these messages are sent and processed. Most Internet protocols specifications use Finite State Machines (FSM) to represent the interactions among the communicating hosts. Although implementations are usually not directly derived from their specification (e.g. for performance reasons or ease of maintenance), most implementations also include the key states and transitions of the protocol specifications. For example, most TCP implementations include the SYN_RCVD, SYN_SENT and ESTABLISHED state of the TCP specifications [62]. While state transitions signal that a connection is making progress, not all of them provide similar information (e.g., transitions into the TCP TIMEWAIT state give no information on the connection besides that "it is about to close"). Ultimately, these FSM describe the life-cycle of a connection. They can thus be abstracted by mapping their state and transitions to the three key phases in a connection life-cycle: (i) the connection establishment; (ii) the exchange of data; and (iii) the connection tear-down. These three stages enable us to define the abstract FSM visible on Fig. 3. When the state of a connection in this simplified FSM changes, it is a signal that Flowcorder needs to create a *performance profile* for the connection. Performance profiles should thus also contain the start and end states corresponding to their transition, enabling to compare the performance of connections for similar transitions (e.g., characterize the connection establishment delay).

A second set of events that requires Flowcorder to generate a performance profile are the functions in the protocol implementation that indicate that an unexpected event occurred (e.g., a retransmission timeout). We model this by a looping transition in the ESTABLISHED state in Fig. 3.

Finally, a third set of probe locations is defined by KPIs that are not computed by default by the protocol implementation. For example, metrics related to reordering for the TCP instrumentation. Tracking these KPIs then implies to create an ancillary state for the connection (e.g., using an eBPF map), and updating it as the connection advances.

Once exported by the eBPF handlers, these performance profiles will eventually be received by an user-space aggregation daemon. This daemon then serialises these profiles to an IPFIX record, adding in the process information to identify both the connection (e.g., the TCP 5-tuple) as well as the network path used (e.g., the egress interface and source address). This record is then eventually exported to the collector.

4 Instrumenting TCP with eBPF

To demonstrate the applicability of our approach, we have applied it to the TCP implementation of the Linux kernel. This is a high-performance and widely used TCP implementation that has been tuned over more than a decade. We first introduce the KPIs building up the performance profiles of TCP connections (§4.1). Then, we describe the various eBPF handlers that are used, and illustrate their interactions (§4.2). Finally, we present how we have extended this instrumentation to support MPTCP (§4.3), showing the genericity and the flexibility of our approach.

4.1 Selecting KPIs

Instrumenting the Linux kernel TCP stack requires to map the chosen KPIs to TCP state variables. A TCP connection is represented in the kernel using the struct `tcp_sock`. As-is, this structure already contains most of the KPIs presented in Table 1. For example, `bytes_received` tracks the received bytes; `srtt_us` is a moving average of the estimated TCP RTT. Computing the statistics to create a performance profile from these state variables thus requires the eBPF handler to: (i) retrieve the address of the connection state from the parameters of the instrumented functions; (ii) copy the relevant state variables from the kernel memory to the eBPF stack; and (iii) compute the statistics on the evolution of the KPIs that these variables represent.

Unfortunately, not all KPIs from Table 1 are directly available in the TCP implementation. More specifically, four KPIs are missing. First, the number of duplicate incast bytes (`Duplicates`) is never recorded. If a connection receives a segment already (partially) acknowledged, the implementation ignores its payload. Second, the number of retransmission timeouts (`Stalls`) is not recorded. Similarly, the number of bytes and packets that arrived out of order (OFO) is not tracked. Finally, the existing reordering connection state variable is not sufficient to represent the distance between out-of-order packets (`OFO-dist`). Indeed, while it does express an out-of-order distance, it does so in terms of number of MSS-sized segments, and represents only the value computed for the last packet. Furthermore, it is clamped by a `sysctl` value.

Recording such "custom" KPIs thus requires to create an eBPF map alongside the probe handlers. This map can then be used to contain the ancillary state for each monitored connection (i.e., map a connection state to a data structure containing the value of the KPIs not provided by the protocol implementation). Managing this map has two implications. First, new entries must be added for any connection that will be monitored. This is especially important for connections initiated by the end-host itself. Indeed, if the TCP SYN they send is lost, the retransmission timer will expire, and the count of connection stalls will need to be increased. This does not apply for inbound connection requests, as creating state before their acceptance by user-space application would provide a Denial-of-Service attack vector. Similarly,

this ancillary state must be purged when the connection is over. The second implication of managing such ancillary state is that it imposes to insert eBPF code at every location where one of its value needs to be updated. Fortunately, as the missing KPIs represent very specific behaviors, these only require to instrument two extra locations (see §4.2).

4.2 Defining eBPF probes

Table 2 lists the functions of the Linux kernel where we insert our probes as well as their handler(s). These functions were chosen to minimize the overhead induced by the probes, i.e., they are never executed in the context of the TCP "fast-path" processing. They fall into two categories. First, we instrument the functions that correspond to state changes in the TCP FSM (i.e., from `tcp_v6_connect` to `tcp_set_state`). These indicate changes in the connection life-cycle and thus mandate to compute KPIs. Second, we instrument functions that denote events which require us to update our ancillary connection state. More specifically, `tcp_retransmit_timer` let us track expirations of the retransmission timer. If a connection experiences a RTO, and its write queue is not empty or the user-space is blocked on a syscall, then it means that the connection has stalled. `tcp_fast_retrans_alert` may signal that a connection has recovered from a RTO (i.e., that the network is stable again) and moved back in the established state. `tcp_validate_incoming`'s instrumentation is split into two handlers. First, it detects whether an incoming segment has already (partially) been acknowledged. Such a segment is an explicit signal that the other host experienced a retransmission timeout. Second, if the function accepts the received segment, this means that it is an out-of-order segment, and the handler updates the statistics tracking the reordering. Furthermore, as both `tcp_retransmit_timer` and `tcp_fast_retrans_alert` indicate that a significant performance event has occurred (a succession of losses in the network, and then a recovery), their handler also export KPIs. This eventually creates performance profiles looping on the ESTABLISHED state, enabling to describe the performance of the connection before, during, and after such transient events (e.g., a flash crowd causing congestion).

Collecting KPIs for a new outbound connection. We now illustrate how Flowcorder exports KPIs describing the establishment of a new outbound connection. In the example shown in Fig. 4, an application creates a regular TCP socket. Then, it tries to establish a TCP connection with the `connect()` system call. This system call is processed by the kernel, and eventually reaches the `tcp_v4_connect()` function, for which Flowcorder had registered a probe. This probe is executed before the instrumented function. It registers basic information about this connection establishment, such as its destination address and the time at which it started. Then, the kernel executes the `tcp_connect()` function, eventually sending a TCP SYN segment. When the function exits, the post handler is executed and immediately returns as

Probe location	Pre	Post	Handler description
tcp_v[46]_connect	✓	✓	Register a new connection attempt and initialize its ancillary state; export KPIs to an error state if the function returns an error which indicates a cancellation of the connection.
tcp_finish_connect	✓	□	Exports KPIs indicating the establishment of a new outbound connection.
inet_csk_accept	□	✓	Exports KPIs for a new inbound connection accepted by user-space.
tcp_set_state	✓	□	If a connection moves to TCP_CLOSE, compute its final state and exports its KPIs.
tcp_retransmit_timer	✓	□	Export KPIs if the connection has stalled and enters a lossy state once established.
tcp_fastretrans_alert	□	✓	If the connection congestion control state moves back to TCP_CA_OPEN (e.g., has recovered from an RTO), exports KPIs to mark the end of the lossy state.
tcp_validate_incoming	✓	✓	Detect incast duplicates; update the reordering KPIs if the packet enters the ofo_queue.

Table 2. A few probes in the Linux TCP implementation act as events to detect many performance changes.

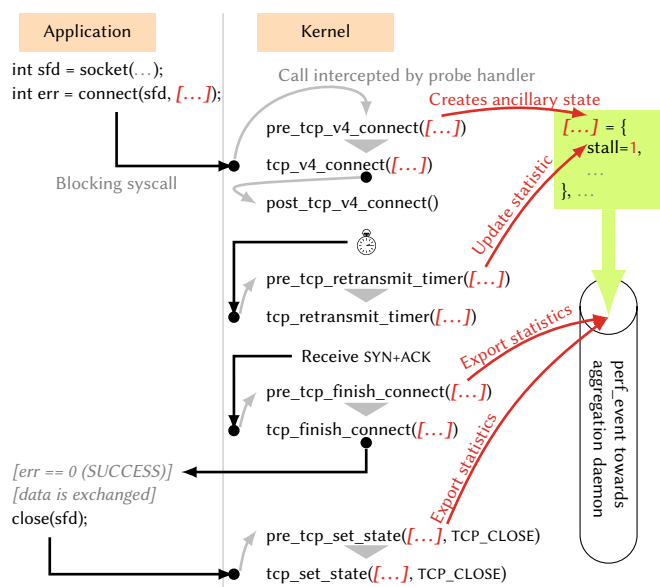


Figure 4. Abstract time-sequence diagram of the generated performance profiles of a TCP connection which loses its initial SYN, exchanges data, then closes. With a few kernel probes, our eBPF handlers trace the entire connection life-cycle and report it to an user-space daemon.

the connection was successfully initiated and the kernel switches to other tasks. Unfortunately, this initial SYN does not reach the destination. After some time, the retransmission timer expires. This causes the kernel to execute the `tcp_retransmit_timer()` function. Again, Flowcorder intercepts that call using a probe, which increments the number of stalls. The kernel then sends a second TCP SYN.

When receiving the corresponding SYN+ACK, the kernel reaches `tcp_finish_connect()`. As its corresponding eBPF handler is awoken, Flowcorder marks the connection as established, computes its KPIs and sends them to the user-space

aggregation daemon using a `perf_event`. This daemon asynchronously fetches and analyzes the KPIs, builds the performance profile of this new connection and adds it in its IPFIX pending message buffer to send it later to the collector. In parallel, the `tcp_finish_connect()` kernel function completes and wakes up the application which can use the connection.

If the network then behaves perfectly (e.g., no reordering, and no losses), the probes placed in the kernel are never reached thus never executed for that connection. Finally, when the application closes its socket, the kernel eventually calls `tcp_set_state` to move the underlying connection to the TCP_CLOSE state. Flowcorder intercepts this call, computes the final set of KPIs for this connection, and exports a performance profile covering the entire connection and reaching a final state describing how the connection ended (e.g., FINISHED if both TCP FIN's were received and acknowledged).

4.3 Supporting MPTCP

MPTCP is a new TCP extension which enables to operate a single TCP connection over multiple paths [29]. Two main implementations of this protocol exists: the reference one in the Linux kernel [59] and one deployed by Apple on iOS [3]. We now demonstrate the genericity of Flowcorder, by enabling it to record performance profiles of MPTCP connections.

To instrument MPTCP, a few architectural details have to be taken into account. Despite being a relatively complex implementation (~18kLOC), it is heavily tied to the existing TCP implementation. At its heart, a MPTCP connection operating over two paths is composed in the kernel of two TCP connections, and of one meta-socket. This meta socket is the one exposed to user-space. It hijacks the socket API used by TCP (i.e., user-space programs use MPTCP by default). Sending data using MPTCP requires to break the bytestream received on the meta-socket into chunks with a MPTCP sequence number (DSS), and then to send those over one of the subflows. The receiver's meta socket then reads the receive queues of its subflows, and reassembles the original bytestream thanks to the DSS.

Instrumenting this implementation poses three challenges: (i) differentiating between a new MPTCP connection and regular TCP one can only be done once the SYN+ACK has been received, since MPTCP connection will contain a dedicated option (MP_CAPABLE); (ii) MPTCP subflows will trigger the same eBPF probes as regular TCP connections; (iii) new subflows can be created directly by the meta-socket.

KPIs specific to MPTCP. As MPTCP subflows operate as regular TCP connections, we use the same set of KPIs as in §4.1 with one addition. When a retransmission timeout occurs on a subflow, its unacknowledged segments are retransmitted both on the subflow itself, as well as on another (they is *reinject*ed on another subflow). We record the number of reinjections done by a subflow in a new KPI present in the ancillary state of the MPTCP subflows. Additionally, the meta-socket provides a bytestream service pretending to be TCP. As such, it supports most of the KPIs supported by TCP, with four tweaks. First, as it gets its segments from underlying TCP connections, it cannot receive corrupted segments and has no concept of latency, removing those KPIs. Second, segments arriving out-of-order on the meta-socket no longer indicate reordering happening in the network. Indeed, such reordering is hidden by the subflows. Instead, reordering on the meta-socket is instead tied to the relative performance difference between the subflows¹. Third, duplicate incast segments now indicate reinjections. Finally, retransmission timeouts at the meta-socket level indicate that the connection is suffering from head-of-line blocking (e.g., a lossy subflow prevents all others from making progress). As one of the more common causes of such a behaviour are too small receive buffers, this defines a new KPI specific to the meta-socket.

eBPF probes handlers. All probes defined in §4.2 also record the performance of MPTCP subflows as-is. In addition to them, we update the ancillary state tracking reinjection across subflows by instrumenting `__mptcp_reinject_data`. Recording the performance of the meta-socket also requires the addition of probes to record the expiration of its retransmission timer (`mptcp_meta_retransmit_timer`). New subflows initiated by the instrumented host are automatically handled by the probes handling the creation of TCP connections. Detecting the creation of new subflows initiated by the remote host requires instrumenting `MPTCP_CHECK_REQ_CHILD`.

5 Evaluation

In this section, we begin by evaluating the overhead of Flowcorder when instrumenting the Linux TCP stack. We first run micro-benchmarks to estimate the overhead of Flowcorder in function of on the characteristics of the underlying

¹Consider two successive segments A and B, such that A comes first in the MPTCP bytestream. If B arrives before A on the receiver’s meta-socket, it then follows that: (i) A and B were sent over different subflows, as subflows guarantee in-order delivery; and (ii) the subflow of B was “better”, e.g., had a lower latency, and/or less losses.

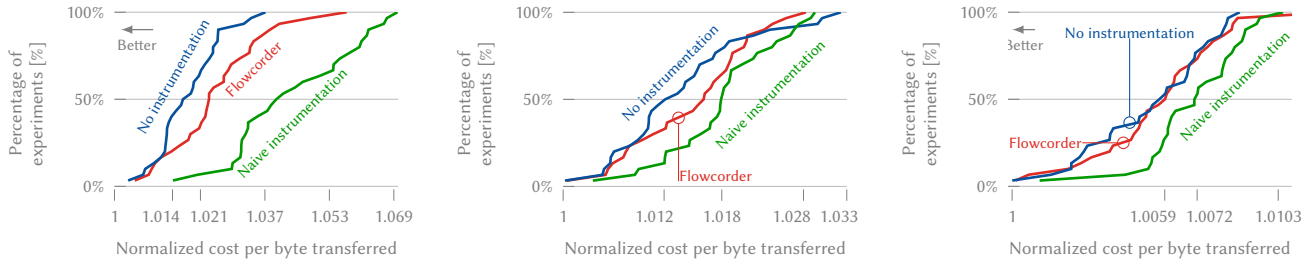
network (§5.1). Then, we evaluate the application-visible performance impact of instrumenting the TCP stack (§5.2). Both sets of experiments confirm that using Flowcorder induces close to no performance overhead on the end-hosts.

Finally, we conclude the section by presenting how to verify that the performance profiles produced by Flowcorder are accurate, especially after kernel upgrades containing potential changes in the instrumented protocol implementation (§5.3). We confirm that Flowcorder supports multiple versions of Linux (v4.5 to v4.18) without any modification.

5.1 Instrumentation overhead

To estimate the overhead induced by the monitoring daemons as well as the kernel probes injected in the TCP stack by Flowcorder, we use a simple benchmark between two servers (each with 8-cores CPUs at 2.5Ghz and 8G of RAM) and connected using 10G interfaces. We use `NTTTC` [56] to initiate multiple parallel TCP connections from one server to the other (between 8 and 100), effectively saturating the 10G link. For each experiment, we record how many bytes were successfully transferred, and use `perf` [48] to record the number of CPU instructions that were executed during each experiment, as reported by the hardware counters. Each experiment ran for 60 seconds, in order to average out measurement errors. To evaluate all instrumented code paths, we also vary the RTT applied over the link (from a few hundred μ s to 100ms), its jitter (10% of the RTT), and its loss rate (from 0 to 1% of random losses). We performed 100 experiments per combination of RTT and loss rate.

To provide quantitative baselines, we repeated each benchmark three times: (i) without any instrumentation; (ii) with Flowcorder running on a server; and (iii) with a naive eBPF instrumentation. This naive instrumentation consists of a version of Flowcorder where an eBPF probe updates KPIs at each incoming segment once the connection reached the `ESTABLISHED` state, i.e., it instruments `tcp_rcv_state_process` in place of `tcp_validate_incoming` to detect out-of-order segments, or incast duplicates. We define the instrumentation overhead as the average number of instructions executed on the servers, divided by the number of bytes successfully transferred. On one hand, this metric let us easily quantify the overhead induced by Flowcorder as it directly gives the amount of extra work carried by a server to execute the eBPF probes. On the other hand, we can compare the gains brought by carefully selecting the probe locations by comparing the overhead of the two different instrumentations. Moreover, the probe induced by the naive implementation is executed for every incoming segment but rarely does any significant work as few segments cause KPI changes (i.e., it often results in a no-op). As such, it implicitly estimates the intrinsic overhead of placing a probe in the TCP “hot” path (i.e., the cost of the software interrupt and the preparation of the eBPF stack). Using the number of executed CPU instructions as metric has at least four advantages: (i) it is independent of



(a) Flowcorder induces a small overhead when used over a link with no loss and sub-ms RTT. (b) Introducing a 10ms RTT and 0.1% loss rate reduces the overhead by an order of magnitude. (c) 30ms RTT and 1% losses cause overhead of Flowcorder to become negligible.

Figure 5. Analyzing the number of instructions executed to saturate a 10G link shows that the overhead induced by the kernel probes is negligible, especially when the link exhibits losses or reordering.

the precise duration of the experiment (i.e., coarse-grained timers have no incidence on the results); (ii) it isolates the results from the transient states of TCP congestion control; (iii) it is independent of the CPU frequency, which is adjusted dynamically by the CPU; and (iv) it captures both the load induced by the kernel probes and the load induced by the user-space daemons aggregating KPIs and exporting IPFIX records. We show a summary of the results in Fig. 5, which plots the cumulative distribution of the fraction of experiments according to their normalized cost (i.e., we normalize all costs by the lowest one).

When operating over a perfect link (Fig. 5a), we see that Flowcorder increases by less than 1% the number of instructions executed during a test. As the experiments had almost no delay and no losses, this gives a baseline as how expensive it is to run Flowcorder, when all connections are processed in the kernel fast path (i.e., the path leveraging as many optimizations as possible, such as hardware offload or skb coalescing, which decreases the overall CPU cost of the connection) thus triggering as few events as possible. This contrasts with the naive instrumentation which has an overhead of more than 2%. When adding some delay (10ms of RTT, and 1ms of jitter), and a small random loss probability of 0.1%, we see in Fig. 5b that the per-byte instruction overhead decreases quite substantially to approximately 0.3%. Indeed, as segments start to arrive out-of-order, or are lost, the TCP stack begins to process them in the slow path, which is much more expensive CPU-wise than the load induced by Flowcorder. This impact is even more visible as we reach a RTT of 30ms±3ms, with a loss rate of 0.5% (Fig. 5c) where the overhead induced by Flowcorder is almost 0.

This indicates that the relative cost of using Flowcorder decreases when the network quality worsens, thus when Flowcorder starts to actually produce performance profiles. The handling of lost or out-of-order segments has a much larger impact on the performance than the kernel probes inserted by Flowcorder and associated monitoring daemons.

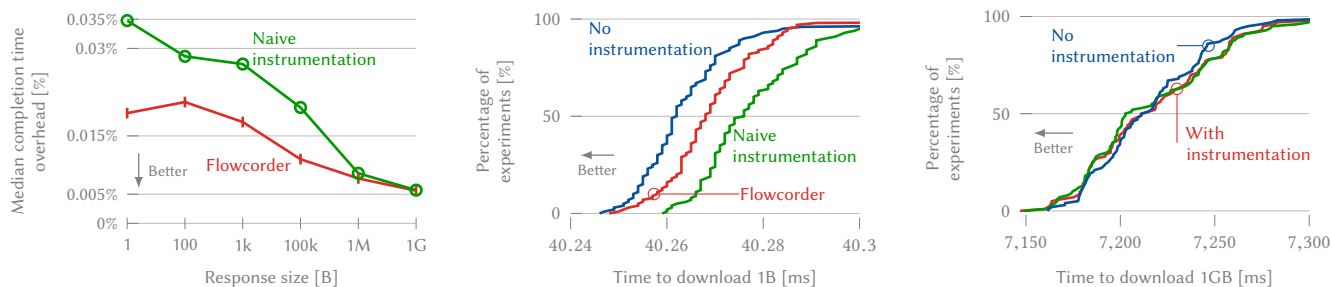
The decrease in the number of instructions per byte transferred between Fig. 5a and Fig. 5b is expected, as increasing the RTT by several orders of magnitude increases the idle periods of connections as they wait for ACKs.

We performed the same experiments when instrumenting the MPTCP implementation (§4.3) and observed similar overhead figures, although there were almost no differences between the two instrumentations as MPTCP disables the kernel TCP fast path processing. Finally, we stress that Flowcorder’s memory overhead is limited by design, as it only has to allocate memory for the ancillary state (bounded by default to about 600kb, i.e., 3000 TCP flows), as well as a python VM holding an MTU-sized IPFIX buffer.

5.2 Impact on application performance

The previous section showed that Flowcorder was inducing some overhead on the instrumented end hosts. In this section, we evaluate whether this overhead can cause application-visible performance degradations. To this end, we configure one host to run a HTTP server. We then record the time to perform an HTTP GET to download a file of a given size from the server. As we saw earlier (§5.1), the overhead of Flowcorder is maximum in a perfect network. As such, we directly connect both the client and the server, configure their interfaces to induce a 20ms RTT, and enable Ethernet flow-control to prevent packet losses. We simulate the client requests using ApacheBench [30], with a variable number of parallel connections (up to 100). Each experiment is repeated 2000 times (i.e., we open a total of 2000 connections for each response size). We recorded for each experiment how quickly the connection completed (i.e., how long did it take to perform the TCP three-way handshake, the HTTP GET, then download the response and close the connection). As before, we repeated the benchmark three times (without instrument, with Flowcorder, and with a naive version of Flowcorder). The results are visible in Fig. 6.

Fig. 6a shows the median overhead per response size, which is the observed increase in completion time when



(a) Flowcorder has a negligible overhead on the overall completion time when downloading a file. (b) The instrumentation overhead is comparatively higher when downloading one-byte files. (c) Large transfers amortize the overhead as few performance profiles are generated per connection.

Figure 6. Using Flowcorder has almost no application-visible impact on the performance of the Linux TCP stack.

the end-host was being instrumented by Flowcorder. We see that as the size of the HTTP responses increases, the overhead decreases. This result is expected. Indeed, recall that Flowcorder generates at least two performance profiles for each connection, and none in the established state if there are no performance degradations. If the response exceeds a few TCP segments, its completion time is thus dominated by the TCP data transfer, and not by the execution of kernel probes. Fig. 6b thus shows the absolute worst case for these experiments, as the response consists in a single segment. We see that the median increase in the response time in that case is about 0.017%. Fig. 6c shows the overhead with a 1GB response, which exhibits a much lower completion overhead. We also performed experiments over a link with some delay and/or losses, and observed that the overhead in those case was even lower as the response time was completely dominated by the network characteristics.

These benchmarks, show that despite inducing some overhead, Flowcorder has a very low (if not negligible) impact on the performance of connections initiated by applications. This result also holds when instrumenting MPTCP.

5.3 Ensuring accurate measurements

The content of the performance profiles generated by Flowcorder, and thus the accuracy of the measurements, clearly depends on the correctness of our instrumentation of the protocol implementation.

Sources of measurement errors. Flowcorder extracts most of its KPIs by performing raw memory accesses in the kernel’s per-connection states. As the content or layout of these states could vary across kernel versions, this extraction process is thus a first possible source of errors. Values could be read at incorrect offsets, or be decoded incorrectly (e.g., reading only the first 32b of a 64b counter). A second source of possible errors are the assumptions the probes make on the status of the connection. For example, the TCP instrumentation assumes that a connection can be identified by the memory address at which its state resides, which is conveniently passed around as `struct sock *sk` in most functions.

If this assumption is wrong (or no longer holds due to an update), then Flowcorder will produce incorrect measurements, e.g., it might mix up connections, or wrongly assume that a connection received an out-of-order segment.

A third source of errors is the set of probes and their locations. Indeed, as the implementation of the protocol improves over time, the set of functions called for each event (e.g., received segments, timer expiration) and their relative order might change. The most obvious effect of this on Flowcorder would be inconsistent performance profiles (e.g., increasing the number of bytes transferred of a closed connection), or missed events (e.g., missed RTOs).

Preventing measurement errors. To prevent the first source of errors, Flowcorder re-compiles its eBPF code every time probes are inserted. As this compilation process directly happens on the instrumented host, it can use information local to the machine (e.g., headers matching the running kernel, or values in `procfs` to enable or disable the MPTCP instrumentation). This source of measurement errors is thus prevented by design. Incidentally, this re-compilation process also ensures that probes are always inserted at their proper locations, as their offset are also dynamically computed during the eBPF compilation, either by reading the content of `/proc/kallsyms` for kernel symbols, or using the debug symbols of user-space applications.

To prevent the seconds and third types of errors, we built a test suite using Packetdrill [12]. Packetdrill enables us to test protocol implementation using scripts which describe connections. More specifically, those scripts inject crafted packets in a local interface at specific points in time, as well as specify the content of packet(s) that should be sent by an implementation in response to incoming packets or API calls. Packetdrill contains a set of edge test cases for the Linux TCP implementation, and similar test cases for MPTCP are available [68]. As each test case depicts a well-defined connection, we can statically predict the performance profiles that should be produced by Flowcorder when instrumenting that connection. This lets us build integration tests to

validate that Flowcorder accurately instruments protocol implementations as they evolve.

Using this test suite, we were able to ensure that Flowcorder accurately instruments the TCP stack of the Linux kernel from v4.5 to v4.18, and MPTCP v0.93.

6 Flowcorder in a campus network

We now present measurements collected over one month with Flowcorder in a campus network. We deployed Flowcorder in student computer labs, where we run on every host monitoring daemons that instrument the Linux kernel TCP stack, presented in §4, as well as DNS resolutions libraries. Each end-host is dual-stacked and has public addresses.

Viewing the effects of Happy Eyeballs. Fig. 7a shows the repartition of the TCP connections in function of the IP version used. We see that most of the connections are established using IPv6. As major cloud services are very popular amongst students and they all support IPv6, this could be due to Happy Eyeballs [82]. We can confirm that Happy Eyeballs indeed favors connections over IPv6 by looking at Fig. 7b. It compares the median time required to establish new TCP connections depending on the used address family. More specifically, it only contains connections established towards dual-stacked ASes. We see that the time to open a new connection is similar for both address families, despite IPv4 exhibiting many outliers. As Happy Eyeballs gives IPv6 connections a head start of usually 300ms (although some have called to reduce it [7]), this explains why IPv6 is almost always used to reach popular services.

Comparing the performance of different uplinks. Our network is dual-homed. It uses different uplinks for IPv4 and IPv6. We leverage Flowcorder to analyze the difference between the two address families. Fig. 7c shows the median jitter observed for TCP connections. We observe that the jitter experienced by IPv4 connections is higher than for IPv6. This correlates with the trend from Fig. 7b, where IPv4 showed more variations. Finally, to better understand why the IPv4 connection establishment delay had a higher variance, Fig. 7d shows the ratio of connections that were successfully established after losing their initial TCP SYN. We see that this mainly occurs only for IPv4, which might point to an on-site issue with a firewall or congestion of the IPv4 uplink. Overall, these results show that IPv6 connections seem to perform better than IPv4 connections in our campus. This is expected, as only the IPv4 traffic is shaped by our provider.

Comparing the performance of remote cloud services. Another usage for the measurements collected by Flowcorder is to compare the performance when accessing different cloud services. Indeed, as an ISP might have different peering agreements with them, measuring the quality of the connections towards those service can be a factor to decide whether to subscribe to one service or another (or to select a different ISP). For example, Fig. 7e compares the median TCP RTT when accessing two popular cloud services. For these services, a

low RTT is key to ensure a proper level of interactivity. We see that while both services tend to show similar RTT's over IPv4, one of them (P_B) performs much worse when accessed over IPv6². Keep in mind that while Flowcorder uses TCP's estimates to report RTT and jitter, this might not completely reflect the true values to reach the actual server, as there could be middleboxes or TCP proxies present on the path, fiddling with segments.

Detecting a local operational issue. Beside providing external connectivity, our campus network also hosts services such as a DNS resolver or institutional web servers. During our measurement campaign, students were complaining that accessing those web servers was abnormally slow. As these web servers are collocated with the DNS servers, we can thus directly use Flowcorder to compare their performance. Fig. 7f shows the median time to establish a connection to any of these servers. Given that the servers are located a few hundreds of meters away, 30ms to receive a SYN+ACK is a clear performance anomaly, especially compared to the time required to receive a DNS reply. After talking with the network operators, we learned that this problem was due to a faulty load-balancer that was fixed near the end of the observation period.

7 Related work

Monitoring network performance is an age-old topic. Flowcorder draws from three main threads of work.

Collecting transport performance metrics. Passive inference of transport protocol characteristics has been a primary source of measurements for a long time, e.g., inferring per-flow TCP states by analyzing packet headers provided by a network tap (tstat Mellia [54]), or correlating packet traces collected on the end hosts (Deja-vu [1]). More recent approaches tailored to data-centers (e.g., Trumpet [57], Dapper [33]) perform such analyzes in real-time, at the edges of the network (i.e., access switches or virtual machine hypervisors). While these technique provide fine-grained measurements for TCP they will not be applicable to emerging encrypted protocols such as QUIC.

Instrumenting the end-hosts. SNAP [84] or NetPoirot [4] collect an enormous amount of statistics about TCP connections directly from datacenter hosts. By collecting those on a central management system, they can then correlate observations in order to identify the root causes of performance issues (e.g., bottleneck switch or link, or misconfigured of TCP delayed ACK's). Both tools poll event loggers (e.g., Windows EWT, or Linux syslog) every few milliseconds. As such, they are restricted to the measurements provided by those loggers (typically the TCP MIB [64]), with a higher CPU overhead than Flowcorder. Odin [11] is a framework injecting javascript when serving client requests from CDN to perform active

²Further analyzes revealed that the provider's DNS was causing students' requests to use datacenters located on another continent.

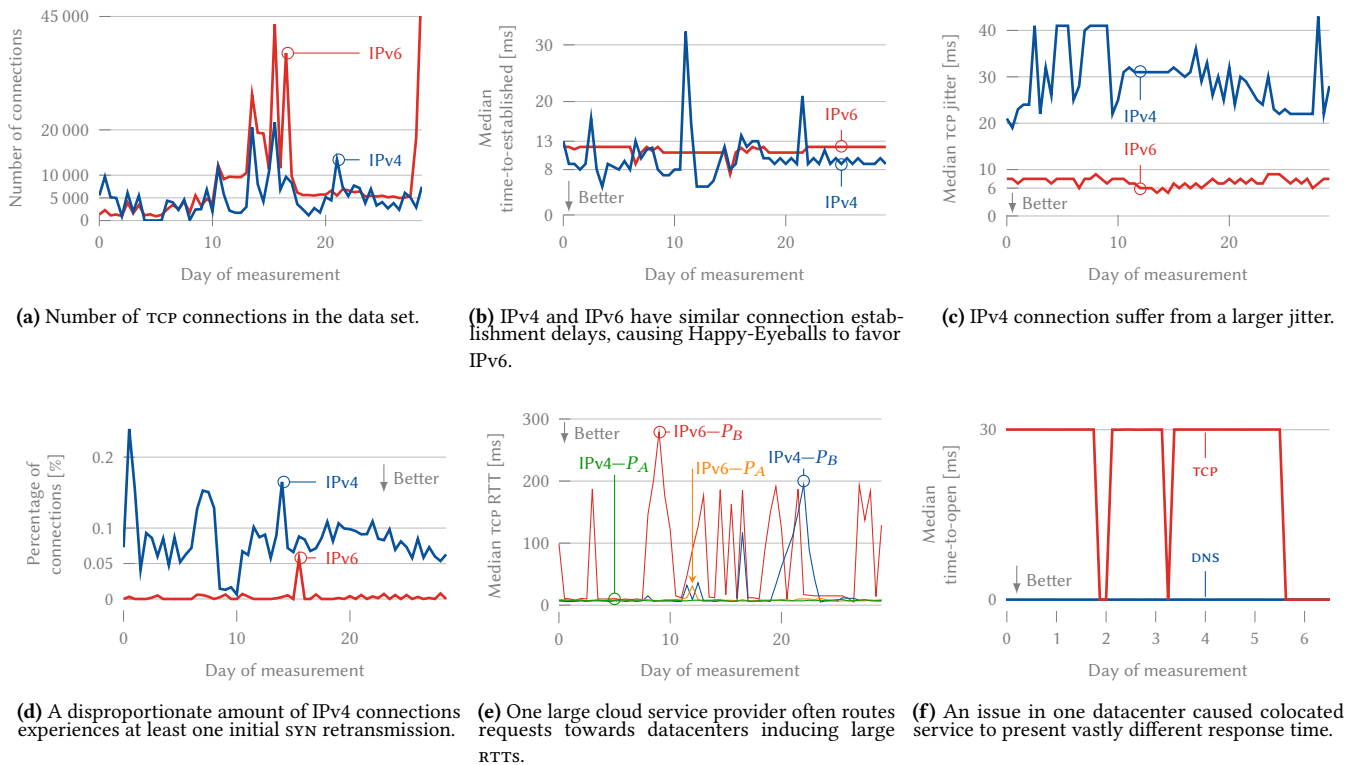


Figure 7. Network performance insights provided by Flowcorder in a dual-stacked, multi-homed, campus network.

measurements. While this approach collects performance metrics as experienced by end-hosts, the measurements that it can records are, by design, much more limited.

Instrumenting protocol implementations. Several tools provide some visibility over the internals of the Linux TCP stack. `tcpprobe` [73] is a kernel module which logs the evolution of the congestion control variable in response to incoming TCP segments. `tcp-tracer` [81] reports the TCP state changes (e.g., `NEW`→`ESTABLISHED`) for all connections. `BCC` [72] provides several small tools, enabling to log some aspects of TCP connections. All of these tools use the same primitives to instrument the TCP stack (i.e., `kprobes`, often combined with `eBPF` handlers), but they are not coupled with enterprise management systems.

8 Conclusion

Flowcorder is a new monitoring framework which directly extracts Key Performance Indicators from the end-hosts, at specific moments in a connection life-cycle. Flowcorder seamlessly integrates with existing Network Management Systems as it generates IPFIX performance profiles. Furthermore, it is future-proof as it readily supports multipath protocols and will also be useable with emerging encrypted protocols. Flowcorder has almost no runtime overhead, and its measurement can easily be analyzed. One future research direction

would be to use the performance profiles generated by Flowcorder to drive tight-control loops on network controllers, to optimize the content of DNS replies (e.g. dynamically preferring the best address family) or to select the best performing provider in multihoming scenarios.

Software artefacts

We release the sources of Flowcorder at <https://github.com/oliviertilmans/flowcorder> under a permissive license. These sources are primarily composed of python (~3300 lines) and restricted c that compiles to `eBPF` (~1900 lines). These include the TCP monitoring daemon which has been tested to work on the Linux kernel from v4.5 to v4.18, its extension to support `MPTCP v0.93`, the DNS monitoring daemon, and scripts to package and deploy them. We also provide a sample IPFIX collector based on an `ELK` stack [24] which comes with preloaded normalization filters. Finally, to ensure the reproducibility of our results, we also provide all scripts used to conduct the benchmarks reported in §5.

References

- [1] Bhavish Aggarwal, Ranjita Bhagwan, Lorenzo De Carli, Venkat Padmanabhan, and Krishna Puttaswamy. 2011. DeJa Vu: Fingerprinting Network Problems. In *CoNEXT '11*. ACM, New York, NY, USA, Article 28, 12 pages. <https://doi.org/10.1145/2079296.2079324>
- [2] Julian Anastasov. [n. d.]. ip-tcp_metrics - management for TCP Metrics. Retrieved 2018 from https://www.linux.org/docs/man8/ip-tcp_metrics.html
- [3] Apple. [n. d.]. Advances in Networking. Retrieved 2018 from <https://developer.apple.com/videos/play/wwdc2017/707/>
- [4] Behnaz Arzani, Selim Ciraci, Boon Thau Loo, Assaf Schuster, and Geoff Outhred. 2016. Taking the Blame Game out of Data Centers Operations with NetPoirot. In *SIGCOMM '16*. ACM, New York, NY, USA, 440–453. <https://doi.org/10.1145/2934872.2934884>
- [5] Brice Augustin, Timur Friedman, and Renata Teixeira. 2011. Measuring multipath routing in the internet. *IEEE/ACM Transactions on Networking (TON)* 19, 3 (2011), 830–840.
- [6] Dtrace authors. [n. d.]. About DTrace. Retrieved 2018 from dtrace.org
- [7] Vaibhav Bajpai and Jürgen Schönwälder. 2016. Measuring the effects of happy eyeballs. In *Proceedings of the 2016 Applied Networking Research Workshop*. ACM, 38–44.
- [8] John Bellardo and Stefan Savage. 2002. Measuring packet reordering. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*. ACM, 97–105.
- [9] O. Bonaventure and S. Seo. 2016. Multipath TCP deployments. *IETF Journal* 12, 2 (2016), 24–27.
- [10] E. Boschi, B. Trammell, L. Mark, and T. Zseby. 2009. *Exporting Type Information for IP Flow Information Export (IPFIX) Information Elements*. RFC 5610. RFC Editor.
- [11] Matt Calder, Manuel Schröder, Ryan Stewart Ryan Gao, Jitendra Padhye, Ratul Mahajan, Ganesh Ananthanarayanan, and Ethan Katz-Bassett. 2018. Odin: Microsoft’s Scalable Fault-Tolerant CDN Measurement System. In *NSDI 18*. USENIX Association.
- [12] Neal Cardwell, Yuchung Cheng, Lawrence Brakmo, Matt Mathis, Barath Raghavan, Nandita Dukkkipati, Hsiao-keng Jerry Chu, Andreas Terzis, and Tom Herbert. 2013. packetdrill: Scriptable Network Stack Testing, from Sockets to Packets.. In *USENIX Annual Technical Conference*. 213–218.
- [13] Neal Cardwell, Yuchung Cheng, C Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. BBR: Congestion-based congestion control. *Queue* 14, 5 (2016), 50.
- [14] Jeffrey D. Case, Mark Fedor, Martin Lee Schoffstall, and James R. Davin. 1990. *Simple Network Management Protocol (SNMP)*. STD 15.
- [15] Taejoong Chung, David Choffnes, and Alan Mislove. 2016. Tunneling for transparency: A large-scale analysis of end-to-end violations in the internet. In *IMC'16*. ACM, 199–213.
- [16] Cisco. [n. d.]. IP SLAs Configuration Guide. Retrieved 2018 from <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/xr-16/sla-xr-16-book.html>
- [17] Cloudflare. [n. d.]. 1.1.1.1 – the Internet’s fastest, privacy-first DNS resolver. <https://1.1.1.1/>
- [18] Jonathan Corbet. 2007. Large receive offload. Retrieved 2018 from <https://lwn.net/Articles/243949/>
- [19] D. Crocker and P. Overell. 2008. *Augmented BNF for Syntax Specifications: ABNF*. STD 68. RFC Editor.
- [20] Quentin De Coninck and Olivier Bonaventure. 2017. Multipath QUIC: Design and Evaluation. In *Conext'17*. ACM, 160–166.
- [21] Piet De Vaere, Tobias Bühler, Mirja Kühlewind, and Brian Trammell. 2018. Three Bits Suffice: Explicit Support for Passive Measurement of Internet Latency in QUIC and TCP. In *IMC'18*.
- [22] Idilio Drago, Rick Hofstede, Ramin Sadre, Anna Sperotto, and Aiko Pras. 2015. Measuring cloud service health using NetFlow/IPFIX: the WikiLeaks case. *Journal of network and systems management* 23, 1 (2015), 58–88.
- [23] Fabien Duchene and Olivier Bonaventure. 2017. Making multipath TCP friendlier to load balancers and anycast. In *2017 IEEE 25th International Conference on Network Protocols (ICNP)*. IEEE, 1–10.
- [24] Elasticsearch. [n. d.]. ELK Stack: Elasticsearch, Logstash, Kibana. Retrieved 2018 from <https://www.elastic.co/elk-stack>
- [25] Andrea Bittau et al. [n. d.]. Tcpcrypt – Encrypting the Internet. Retrieved 2018 from <http://tcpcrypt.org/>
- [26] S. Farrell and H. Tschofenig. 2014. *Pervasive Monitoring Is an Attack*. BCP 188.
- [27] Adrienne Porter Felt, Richard Barnes, April King, Chris Palmer, Chris Bentzel, and Parisa Tabriz. 2017. Measuring HTTPS adoption on the web. In *26th USENIX Security Symposium*. 1323–1338.
- [28] Alessandro Finamore, Marco Mellia, Michela Meo, Maurizio M Munafo, Politecnico Di Torino, and Dario Rossi. 2011. Experiences of internet traffic monitoring with tstat. *IEEE Network* 25, 3 (2011), 8–14.
- [29] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. 2013. *TCP Extensions for Multipath Operation with Multiple Addresses*. RFC 6824.
- [30] The Apache Software Foundation. [n. d.]. ab - Apache HTTP server benchmarking tool. Retrieved 2018 from <https://httpd.apache.org/docs/2.4/programs/ab.html>
- [31] Sriharsha Gangam, Puneet Sharma, and Sonia Fahmy. 2013. Pegasus: Precision hunting for icebergs and anomalies in network flows. In *INFOCOM, 2013 Proceedings IEEE*. IEEE, 1420–1428.
- [32] Jose Garcia-Dorado, Alessandro Finamore, Marco Mellia, Michela Meo, and Maurizio M. Munafo. 2012. Characterization of ISP Traffic: Trends, User Habits, and Access Technology Impact. *IEEE Transactions on Network and Service Management* 9, 2 (Feb 2012).
- [33] Mojgan Ghasemi, Theophilus Benson, and Jennifer Rexford. 2017. Dapper: Data plane performance diagnosis of tcp. In *Proceedings of the Symposium on SDN Research*. ACM, 61–74.
- [34] Google. [n. d.]. QUIC Trace. Retrieved 2018 from <https://github.com/google/quic-trace>
- [35] Sudhanshu Goswami. April 18, 2005. An introduction to KProbes. <https://lwn.net/Articles/132196/>
- [36] Anders Gunnar, Mikael Johansson, and Thomas Telkamp. 2004. Traffic matrix estimation on a large IP backbone: a comparison on real data. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*. ACM, 149–160.
- [37] Chuanxiang Guo, Lihua Yuan, Dong Xiang, Yingnong Dang, Ray Huang, Dave Maltz, Zhaoyi Liu, Vin Wang, Bin Pang, Hua Chen, Zhi-Wei Lin, and Varugis Kurien. 2015. Pingmesh: A Large-Scale System for Data Center Network Latency Measurement and Analysis. In *SIGCOMM '15*. ACM, New York, NY, USA, 139–152. <https://doi.org/10.1145/2785956.2787496>
- [38] Rick Hofstede, Pavel Čeleda, Brian Trammell, Idilio Drago, Ramin Sadre, Anna Sperotto, and Aiko Pras. 2014. Flow monitoring explained: From packet capture to data analysis with netflow and ipfix. *IEEE Communications Surveys & Tutorials* 16, 4 (2014), 2037–2064.
- [39] Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley, and Hideyuki Tokuda. 2011. Is it still possible to extend TCP?. In *IMC'11*. ACM, 181–194.
- [40] R. Housley and R. Droms. 2018. TLS 1.3 Option for Negotiation of Visibility in the Datacenter. (March 2018). Internet draft, draft-rhrtls-tls13-visibility-01.
- [41] IETF doh Working Group. [n. d.]. DNS Over HTTPS. Retrieved 2018 from <https://datatracker.ietf.org/wg/doh/about/>
- [42] IETF quic Working Group. [n. d.]. QUIC. Retrieved 2018 from <https://datatracker.ietf.org/wg/doh/about/>
- [43] IO Visor Project. [n. d.]. eBPF, extended Berkeley Packet Filter. Retrieved 2018 from <https://www.iovisor.org/technology/ebpf>
- [44] ITU. [n. d.]. Introduction to ASN.1. Retrieved 2018 from <https://www.itu.int/en/ITU-T/asn1/Pages/introduction.aspx>

- [45] Sharad Jaiswal, Gianluca Iannaccone, Christophe Diot, Jim Kurose, and Don Towsley. 2007. Measurement and classification of out-of-sequence packets in a tier-1 IP backbone. *IEEE/ACM Transactions on Networking (ToN)* 15, 1 (2007), 54–66.
- [46] Wolfgang John, Sven Tafvelin, and Tomas Olovsson. 2010. Passive internet measurement: Overview and guidelines based on experiences. *Computer Communications* 33, 5 (2010), 533–550.
- [47] Arash Molavi Kakhki, Samuel Jero, David Choffnes, Cristina Nita-Rotaru, and Alan Mislove. 2017. Taking a long look at QUIC: an approach for rigorous evaluation of rapidly evolving transport protocols. In *Proceedings of the 2017 Internet Measurement Conference*. ACM, 290–303.
- [48] The Linux kernel. [n. d.]. perf: Linux profiling with performance counters. Retrieved 2018 from https://perf.wiki.kernel.org/index.php/Main_Page
- [49] The Linux kernel. [n. d.]. struct tcp_info definition. Retrieved 2018 from <https://elixir.bootlin.com/linux/latest/source/include/uapi/linux/tcp.h#L168>
- [50] Adam Langley et al. 2017. The QUIC Transport Protocol: Design and Internet-Scale Deployment. In *SIGCOMM '17*. ACM, New York, NY, USA, 183–196. <http://doi.acm.org/10.1145/3098822.3098842>
- [51] Let's Encrypt Certificate Authority, ISRG. April 17, 2018. Percentage of Web Pages Loaded by Firefox Using HTTPS. Retrieved April 17, 2018 from <https://letsencrypt.org/stats/>
- [52] Bingdong Li, Jeff Springer, George Bebis, and Mehmet Hadi Gunes. 2013. A survey of network flow applications. *Journal of Network and Computer Applications* 36, 2 (2013), 567–581.
- [53] Xiapu Luo, Edmond WW Chan, and Rocky KC Chang. 2009. Design and Implementation of TCP Data Probes for Reliable and Metric-Rich Network Path Monitoring.. In *USENIX Annual Technical Conference*.
- [54] Marco Mellia. 2002. TCP Statistic and Analysis Tool. *IEEE Network* 16, 5 (Sep 2002).
- [55] Microsoft. [n. d.]. Network Tracing in Windows 7: Architecture. Retrieved 2018 from [https://msdn.microsoft.com/en-us/library/windows/desktop/dd569137\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/dd569137(v=vs.85).aspx)
- [56] Microsoft. [n. d.]. NNTTCP-for-Linux. Retrieved 2018 from <https://github.com/Microsoft/nnttcp-for-linux>
- [57] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. 2016. Trumpet: Timely and Precise Triggers in Data Centers. In *SIGCOMM '16*. New York, NY, USA, 129–143.
- [58] Hung X Nguyen and Matthew Roughan. 2013. Rigorous statistical analysis of internet loss measurements. *IEEE/ACM Transactions on Networking (TON)* 21, 3 (2013), 734–745.
- [59] C. Paasch, S. Barr'e, et al. [n. d.]. Multipath TCP in the Linux Kernel. Retrieved 2018 from <https://www.multipath-tcp.org>
- [60] Giorgos Papastergiou, Gorry Fairhurst, David Ros, Anna Brunstrom, Karl-Johan Grinnemo, Per Hurtig, Naeem Khademi, Michael Tüxen, Michael Welzl, Dragana Damjanovic, et al. 2017. De-ossifying the internet transport layer: A survey and future perspectives. *IEEE Communications Surveys & Tutorials* 19, 1 (2017), 619–639.
- [61] K. Pearce. 2014. Multipath TCP Breaking today's networks with tomorrow's protocol. (2014). Presentation at BlackHat 2014.
- [62] Jon Postel. 1981. *Transmission Control Protocol*. STD 7.
- [63] The Chromium project. [n. d.]. How to enable logging in Chromium. Retrieved 2018 from <https://www.chromium.org/for-testers/enable-logging>
- [64] R. Raghunathan. 2005. *Management Information Base for the Transmission Control Protocol (TCP)*. RFC 4022.
- [65] Eric Rescorla. 2018. *The Transport Layer Security (TLS) Protocol Version 1.3*. Internet-Draft draft-ietf-tls-tls13-28. IETF.
- [66] Jan Rühl, Ingmar Poese, Christoph Dietzel, and Oliver Hohlfeld. 2018. A First Look at QUIC in the Wild. In *International Conference on Passive and Active Network Measurement*. Springer, 255–268.
- [67] Omar Santos. 2015. *Network Security with NetFlow and IPFIX: Big Data Analytics for Information Security*. Cisco Press.
- [68] Arnaud Schills. [n. d.]. Packetdrill test suite for MPTCP. Retrieved 2018 from https://github.com/aschills/packetdrill_mptcp
- [69] Vyas Sekar, Nick G Duffield, Oliver Spatscheck, Jacobus E van der Merwe, and Hui Zhang. 2006. LADS: Large-scale Automated DDoS Detection System. In *USENIX Annual Technical Conference, General Track*. 171–184.
- [70] Emile Stephan et al. 2017. QUIC Interdomain Troubleshooting. (July 2017). Internet draft, draft-stephan-quic-interdomain-troubleshooting-00.txt, work in progress.
- [71] The Android project. 2018. Perform network operations using Cronet. Retrieved 2018 from <https://developer.android.com/guide/topics/connectivity/cronet/>
- [72] The IOvisor project. [n. d.]. BPF Compiler Collection (bcc). Retrieved 2018 from <https://www.iovisor.org/technology/bcc>
- [73] The Linux Foundation. [n. d.]. TCP Probe. Retrieved 2018 from <https://wiki.linuxfoundation.org/networking/tcpprobe>
- [74] B. Trammel et al. 2018. Adding Explicit Passive Measurability of Two-Way Latency to the QUIC Transport Protocol. (May 2018). Internet draft, draft-trammell-quic-spin-03, work in progress.
- [75] Brian Trammell and Elisa Boschi. 2011. An introduction to IP flow information export (IPFIX). *IEEE Communications Magazine* 49, 4 (2011).
- [76] Martino Trevisan, Alessandro Finamore, Marco Mellia, Maurizio Munafò, and Dario Rossi. 2017. Traffic Analysis with Off-the-Shelf Hardware: Challenges and Lessons Learne. *IEEE Communications Magazine* (March 2017).
- [77] Piet De Vaere. 2018. *Adding Passive Measurability to QUIC*. Master's thesis. ETH Zurich. MA-2017-16.
- [78] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, and Vern Paxson. 2015. Header enrichment or isp enrichment?: Emerging privacy threats in mobile networks. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization*. ACM, 25–30.
- [79] Daniël van der Steeg, Rick Hofstede, Anna Sperotto, and Aiko Pras. 2015. Real-time DDoS attack detection for Cisco IOS using NetFlow. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*. IEEE, 972–977.
- [80] Tobias Viernickel, Alexander Froemmgen, Amr Rizk, Boris Koldehofe, and Ralf Steinmetz. 2018. Multipath QUIC: A Deployable Multipath Transport Protocol. In *2018 IEEE International Conference on Communications (ICC)*. IEEE, 1–7.
- [81] Weaverworks. [n. d.]. tcptracer-bpf. Retrieved 2018 from <https://github.com/weaverworks/tcptracer-bpf>
- [82] D. Wing and A. Yourtchenko. 2012. *Happy Eyeballs: Success with Dual-Stack Hosts*. RFC 6555. RFC Editor.
- [83] Bahador Yeganeh, Reza Rejaie, and Walter Willinger. 2017. A view from the edge: A stub-AS perspective of traffic localization and its implications. In *Network Traffic Measurement and Analysis Conference (TMA), 2017*. IEEE, 1–9.
- [84] Minlan Yu, Albert G Greenberg, David A Maltz, Jennifer Rexford, Lihua Yuan, Srikanth Kandula, and Changhoon Kim. 2011. Profiling Network Performance for Multi-tier Data Center Applications.. In *NSDI*, Vol. 11. 5–5.