

Caliper: Precise and Responsive Traffic Generator

Monia Ghobadi*, Geoffrey Salmon*, Yashar Ganjali*, Martin Labrecque†, J. Gregory Steffan†

*Department of Computer Science, †Department of Electrical and Computer Engineering

University of Toronto, Toronto, ON, Canada

{monia, geoff, yganjali}@cs.toronto.edu, {martinl, steffan}@eecg.toronto.edu

Abstract—This paper presents Caliper, a highly-accurate packet injection tool that generates precise and responsive traffic. Caliper takes live packets generated on a host computer and transmits them onto a gigabit Ethernet network with precise inter-transmission times. Existing software traffic generators rely on generic Network Interface Cards which, as we demonstrate, do not provide high-precision timing guarantees. Hence, performing valid and convincing experiments becomes difficult or impossible in the context of time-sensitive network experiments. Our evaluations show that Caliper is able to reproduce packet inter-transmission times from a given arbitrary distribution while capturing the closed-loop feedback of TCP sources. Specifically, we demonstrate that Caliper provides three orders of magnitude better precision compared to commodity NIC: with requested traffic rates up to the line rate, Caliper incurs an error of 8 ns or less in packet transmission times. Furthermore, we explore Caliper’s ability to integrate with existing network simulators to project simulated traffic characteristics into a real network environment. Caliper is freely available online.

I. INTRODUCTION AND MOTIVATION

Making any changes to the Internet infrastructure is extremely difficult, if possible at all. Any new network component, protocol, or design implemented on a global scale requires extensive and accurate testing in sufficiently realistic settings. Real network experiments are extremely difficult: network operators usually do not like any modifications to their network, unless the proposed changes have been tested exhaustively in a large scale network. The only remaining option for testing the impact of a given change is using testbeds for network experiments. In this paper, we present the design, implementation, and evaluation of Caliper, a precise and responsive traffic generator based on the NetFPGA platform with highly-accurate packet injection times. Caliper can be easily integrated with various software-based traffic generation tools. Caliper injects dynamically created packets, and thus, it can react to feedback and model the closed-loop behavior of TCP and other protocols. The ability to produce live traffic makes Caliper useful to explore a variety of what-if scenarios by tuning user, application, and network parameters.

There are many challenges associated with performing valid experiments in network testbeds. Generating realistic and responsive traffic that reflects different network conditions and topologies is one of such key challenges. To perform network experiments, researchers often use a collection of commodity Linux machines as traffic generators. However, creating a large number of connections in order to accurately model the

traffic shape in networks with thousands of flows is difficult for several reasons. Below, we summarize some of these key challenges.

Using pre-recorded network traces: It is not always possible to use packet level network traces. The reason is that pre-recorded packet traces do not maintain the feedback loop behavior between the network and traffic sources. For example TCP operates in a closed-loop fashion and it reacts to the congestion feedback it receives from the network. A pre-recorded TCP trace does not react to new network conditions such as packet drops and to perform closed loop experiments, one needs to collect traces at higher levels.

Using network simulation tools: The complexity of traffic generation increases when trying to capture the heterogeneity of link capacities using only a limited number of physical machines. While network simulation tools can be very helpful in understanding the impact of a given change to a network, their predictions might not be accurate due to their simplified and restricted models and settings. For example, to perform router buffer sizing experiments, it is not obvious how to precisely simulate a real production router with many stages of buffering and a unique architecture [1].

Using commodity hardware: Commodity hardware does not guarantee the precision of generated traffic, which is bounded by the system timer resolution.¹ In addition, there are differences in implementation and default settings (and sometimes lack of a way to change those settings) in commodity hardware. Our experiments with Linux using an Intel NIC with three different drivers leads us to believe that not only does these device drivers require different parameters for adjusting their operations, but also the outcome is different. Seemingly similar experiments might result in different outcomes due to differences in devices. For example, Figure 1 shows the differences between the CDFs of packet inter-arrival times when the Interrupt Coalescence (IC) option is turned off for an Intel NIC card with three different drivers.

Hence, it is intrinsically difficult to perform *time-sensitive* network experiments with confidence on the accuracy of packet injection times. Time-sensitive experiments are those that need high-precision timings for packet injections into the network. Experimenting with congestion control algorithms [2], buffer sizing in Internet routers [1], and denial of service attacks which use low-rate packet injections [3] are all examples of time-sensitive experiments, where a subtle

¹A Linux kernel is typically capable of providing resolutions of 1 ms. In comparison, a packet of 1500 bytes on a 1 Gbps link has a transmission time of less than 12 μ s.

A preliminary version of this work appeared in the NetFPGA Developers Workshop 2009.




Fig. 1. CDF of packet inter-arrival times for different NIC card drivers.

variation in packet injection times can change the results significantly [4].²

As an alternative, commercial traffic generators such as Spirent Communications’ Avalanche box [5], Ixia [6], and Endace DAG cards [7] are useful for some experiments, but they have their own drawbacks. They are usually very expensive and their proprietary nature makes them inflexible for research purposes. Prasad *et al.* [8] describe differences observed between a TCP Reno packet sequence generated by Avalanche traffic generator and the expected behavior of the standard TCP Reno protocol.

A more precise and desirable solution for research community is using hardware based packet generators such as the traffic generator by Covington *et al.* [9] (hereafter referred to as the Stanford Packet Generator or SPG). SPG is based on *NetFPGA* [10], a PCI-based programmable board containing an FPGA, four gigabit Ethernet ports, and memory. SPG generates more precise traffic by accurately replicating the transmission times recorded in a *pcap* trace file, similar to the operation of the *tcpreplay* software program; this method eliminates the dependence between the generated traffic and the NIC model. While the traffic that SPG generates is more precise than many prior approaches, it has several limitations. The closed-loop feedback for TCP sources (and any other protocol that depends on the feedback from the system) is not kept because the trace files are based on past measurements. Furthermore, replaying a prerecorded trace on a link with different properties (such as capacity and buffer size) does not necessarily result in realistic traffic without performing non-trivial adjustments. Finally, SPG can only (i) replay the exact packet inter-arrival times provided by the trace file, or (ii) produce fixed inter-arrival times between packets; i.e., ignoring the variation of packet timings from the original trace.

In this paper, we present the design, implementation, and evaluation of Caliper, a precise and responsive traffic generator based on the *NetFPGA* platform with highly-accurate packet injection times that can be easily integrated with various

²For instance, consider a simple router buffer sizing experiment with buffer size of 90 packets and a 10 *Gbps* link transmitting 1500 bytes packets to the buffer. In this case, having 120 μ s packet injection inaccuracy can lead to transmission of 100 back-to-back packets affecting the experiment by filling up the buffer and causing unnecessary packet drops.

software-based traffic generation tools. Caliper provides a key feature that makes it useful in a large range of network experiments: Caliper injects dynamically created packets, and thus, it can react to feedback and model the closed-loop behavior of TCP and other protocols. Note that characterizing real-traffic is not the goal of this work, instead, our objective is packet injection accuracy that complements existing work on traffic characterization and realistic traffic generation.

Our evaluations demonstrate that Caliper is able to reproduce packet inter-arrival times from a given arbitrary distribution with high accuracy while capturing the closed-loop feedback of TCP sources. We present the accuracy of Caliper with various packet arrival rates and demonstrate that with requested traffic rates up to the line rate, the maximum error that Caliper incurs is 8 *ns* which is the resolution of the measuring system’s clock. We further demonstrate integration of Caliper with existing software-based traffic generators as well as the *ns-2* network simulator.

II. PRECISE TRAFFIC GENERATION

Caliper’s main objective is to precisely control the transmission times of live packets which are created in the host computer, continually streamed to the *NetFPGA*, and transmitted on the wire. The generated packets are sent out of a single Ethernet port of the *NetFPGA*, according to any given sequence of requested inter-transmission times. It is important to note that because packets are streamed, the generator can immediately change the traffic in response to feedback. Unlike previous works that replay packets with prerecorded transmission times from a trace file, Caliper generates live packets and supports closed-loop traffic. Therefore, Caliper can easily be coupled with existing traffic generators (such as *Iperf* [11]) to improve their accuracy at small time scales.

A. Caliper’s Components

Caliper is built on *NetThreads* [12], a platform we have created for developing packet processing applications on FPGA-based devices and the *NetFPGA* in particular. *NetThreads* is primarily composed of FPGA-based multithreaded processors, providing a familiar yet flexible environment for software developers: programs are written in C, and existing applications can be ported to the platform. In contrast with a PC or NIC-based solution, *NetThreads* is similar to a custom hardware solution since it offers direct network I/O and allows the programmer to specify accurate timing requirements.

Caliper’s components are illustrated in Figure 2 showing the life cycle of a packet through the system, from creation to transmission. First, a user space process or a kernel module on the host computer determines when a packet should be sent. A description of the packet, containing the transmission time and all the information necessary to assemble the packet is sent to the *NetFPGA* driver. In the driver, multiple packet descriptions are combined and copied to the *NetFPGA* card. Combining descriptions reduces the number of separate transfers required and is necessary for sending packets at the line rate of 1 *Gbps*. From there, packet descriptions are processed in software on the *NetThreads* soft multithreaded processors. Each software




Fig. 2. Components of Caliper packet generator.

thread assembles packets in the output memory. Next, a selected thread sends all of the prepared packets in the correct order at the requested transmission times. Finally, the hardware pipeline of the NetFPGA transmits the packets onto the wire. Note that users can edit all parts of Caliper to modify and extend its functionality [13].

In the rest of this section we briefly explain each stage of a packet’s journey through Caliper. We also describe the underlying limitations and challenges that influenced our design. Due to lack of space, we omit hardware implementation details and refer the interested reader to our technical report [14].

Packet Creation to Driver: The reasons and context of packet creation are application-specific. To produce realistic traffic, we envision that a network simulator, such as Swing [15], will decide when to send each packet. This simulation may be running in either a user space process, like `ns-2` [16], or a Linux kernel module, as in ModelNet [17]. To easily allow either approach, we send packets to the NetFPGA driver using Linux NetLink sockets, which allow arbitrary messages to be sent and received from either user space or the kernel. In Section III we describe examples of using Caliper with Iperf traffic generator and our own user space program. In [14], we describe and evaluate a prototype that we develop to allow packets from the `ns-2` simulator to be sent on a real network using Caliper. At this stage, the messages sent to the NetFPGA driver do not contain the entire packet as it will appear on the wire. Instead, packets are represented by minimal descriptions which contain the size of the packet and enough information to build the packet headers. The parts of the payload that are not set will be zeroed when the packet is eventually transmitted.

Driver to NetThreads: We modified the driver provided with NetFPGA to support Caliper. Its main task is to copy the packet descriptions to the NetFPGA card using DMA over the PCI bus. It also assembles the packet headers and computes checksums. To obtain the line rate throughput, the driver combines the headers of multiple packets and copies them to the NetFPGA in a single DMA transfer. Next, the NetFPGA hardware pipeline stores them into the input memory of the NetThreads system.

NetThreads to Wire: This last part of Caliper runs as software on the NetThreads platform inside the NetFPGA. The driver sends its messages containing the headers of multiple packets and their corresponding transmission times. Then, Caliper prepares these packets for transmission and sends them at the appropriate times.



Fig. 3. Topology of experiments.

B. Integration with Existing Tools

Caliper is intended to be integrated with software traffic generators. Since Caliper is acting as a network card device in the kernel, it can transmit packets generated within the Linux kernel with any software packet generator (i.e., ping, Iperf, etc.) according to user-specified inter-arrival times. When using TCP sources, Caliper can transmit live TCP connections and closed-loop sessions. As a result, the generated traffic becomes “responsive” to changing network conditions or competing application traffic by capturing the congestion feedback of TCP sources and any other Linux implemented protocols. For example, one desirable advantage of using Caliper for TCP sources is the ability to define an inter-packet gap to provide TCP pacing [18] in hosts. TCP pacing addresses the burstiness of TCP traffic by minimizing the possibility of overflow in router buffers [1]. Towards this goal, Caliper is able to adjust precisely the interval between outgoing packets and produce smoothed and stable traffic. In Sections III-A and III-C, we demonstrate that Caliper provides three orders of magnitude better precision compared to commodity NIC when using Iperf to generate TCP and UDP traffic.

III. EVALUATION

In this section we evaluate the performance of Caliper by focusing on its accuracy and flexibility features. We set-up our experiments to reflect Caliper’s intended use: to complement existing traffic generators by allowing them to precisely control when packets are transmitted. Hence, we present our experiments where the most important metric is the accuracy of packet transmission times. We also present measurements of an existing network emulator which clearly demonstrate the need that Caliper fulfills. We further present a prototype that integrates Caliper with the `ns-2` simulator.

We perform our evaluations using Dell Power Edge 2950 servers running Debian GNU/Linux 5.0.1 (codename Lenny) each with an Intel Pro/1000 Dual-port gigabit network card and a NetFPGA card. The topology of our experiments is illustrated in Figure 3. In each test, there is a single server sending packets and a single server receiving packets via a NetFPGA-based router in the middle that measures the packets inter-arrival times. In the experiment described in Section III-D, the router is replaced with a server running a software network emulator which routes packets between the sender and receiver.

Since Caliper’s main goal is to transmit packets exactly when requested, the measurement accuracy is vital to the evaluation. Measuring arrival times in software using `tcpdump` or similar applications is imprecise: generic NICs combined with

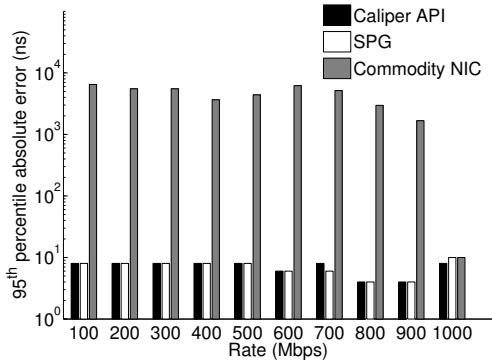


Fig. 4. Comparing the 95th percentile of absolute error ($|D_R - D_M|$) between Caliper, SPG, and commodity NIC when injecting UDP packets.

OS overheads are intrinsically inaccurate at the level we operate [4]. Therefore, we use a NetFPGA router to measure packet inter-arrival times at the middle node (router). The NetFPGA router is configured with the “event capturing module” of the NetFPGA router design[1], [19] which supports instrumenting the router’s output queues.³ when a packet arrives, departs or is dropped, the system clock time of the NetFPGA, which has an 8 ns granularity, is recorded. To reduce overhead, the NetFPGA router batches multiple events in packets that are periodically sent out and we obtain the packet inter-arrival times at the router by subtracting successive timestamps in the payload of those event packets.

A. Sending UDP Packets at Fixed Intervals

The simplest test case for Caliper is to generate UDP packets with a fixed inter-transmission time. Comparing the requested inter-transmission time with the observed inter-arrival times demonstrates Caliper’s degree of precision. As explained in Section II, Caliper leverages software running on what has previously been a hardware-only network device, the NetFPGA. Even executing software, NetThreads should provide sufficient performance and control for precise packet generation.

To evaluate the above criteria we compare Caliper’s transmission times against those of Stanford’s Packet Generator (SPG), which is implemented on the NetFPGA solely in hardware. Moreover, we demonstrate the lack of precision when using a commodity NIC transmitting Iperf traffic. Figure 4 shows the 95th percentile of absolute error between the measured inter-arrival times (D_M) and the requested inter-transmission times (D_R) corresponding to various packet transmission rates (T_R). It is important to note that the 95th percentile error is a more conservative metric than the average error as it captures the 5% largest errors. For each transmission rate, we send 1,500,000 UDP packets of size 1518 bytes

³To increase the accuracy of the timestamps even more, we removed two parts of the router pipeline that could add a variable delay to packets before they reach the output queues. This is possible because we are only interested in measuring packets that arrive at a particular port and the routing logic is unnecessary.




Fig. 5. CDF of measured inter-arrival times compared with an input Pareto distribution. Only a single curve is visible since the two plots match entirely.

(including Ethernet headers) using Caliper, SPG, or an Intel commodity Ethernet NIC. To generate constant bit rate traffic over the commodity NIC we use the Iperf traffic generator with rate T_R . We then capture a portion of traffic in a trace file and replay it with SPG while configuring SPG with the exact same packet inter-arrival time that we used with Caliper, D_R .

As Figure 4 illustrates, for all range of transmission times up to 1 Gbps, the 95th percentile absolute error is around 8 ns for both Caliper and STG. The clock period of the sending and measuring NetFPGA systems is 8 ns, and hence an error of 8 ns implies that most of the inter-transmission times are within one clock cycle (the measurement resolution). This shows that even though NetThreads is executing software, it still allows precise control on when packets are transmitted. On the other hand, note that the commodity NIC’s error is almost three orders of magnitude higher than both Caliper and STG. At 1 Gbps rate, we notice that the error is minimum for Caliper, STG as well as the commodity NIC case because the network is operating at its maximum utilization and packets are sent and received back-to-back.

Although both Caliper and STG packet generators are of similar accuracy, SPG has a limitation that makes it unsuitable for the role we intend for Caliper. The packets sent by SPG must first be loaded onto the NetFPGA as a pcap file before they can be transmitted. This two-stage process means that SPG can only replay relatively short traces that have been previously captured.⁴ Although SPG can optionally replay the same short trace multiple times, it can not dynamically be instructed to send packets by a software packet generator or network emulator using a series of departure times that are not known a priori. Caliper, on the other hand, can be used to improve the precision of packet transmissions streamed by any existing packet generation software.

B. Variable Inter-arrival Times and Packet Sizes

Another advantage of Caliper is its ability to generate packets with an arbitrary sequence of inter-arrival times and

⁴The largest memory on the board is 64 MB which is only about 0.5 seconds of traffic at the 1 Gbps line rate.




Fig. 6. CDF of measured inter-arrival times when both packets sizes and requested inter-arrival times vary. Only a single curve is visible since the two plots match entirely.

sizes that are both essential parts of performing realistic large scale experiments. Figure 5 shows the CDFs of both the requested and the measured transmission times for an experiment with 4000 packets with inter-arrival times following a Pareto distribution. Interestingly, only a single curve is visible in the figure since the two curves match entirely (for clarity we add crosses to the figure at intervals along the input distribution’s curve). As we will demonstrate in Section III-D, this property of Caliper is exactly the component that the network emulators need. Caliper can take a list of packets and transmission times and send the packets when requested. The crucial difference between Caliper and SPG is that SPG has a separate load phase that prevents it from being used by network emulators.

As another example, Figure 6 shows the CDFs of the requested and the measured transmission times when the requested inter-arrival of packets follows the spike bump pattern probability density function observed in the study on packet inter-arrival times in the Internet by Katabi *et al.* [20]. In this distribution, a flow traverses a low bandwidth bottleneck with an inter-arrival of 8 *ms* followed by a high bandwidth bottleneck. Moreover, to demonstrate Caliper’s ability in generating packets with variable sizes, we choose the packet sizes according to another realistic distribution from the same study: 50% are 1518 bytes, 10% are 612 bytes, and 40% are 64 bytes. Note that, again, Caliper generates the traffic exactly as expected and hence only one curve is visible.

C. Generating Responsive Traffic

As explained in Section III-E, Caliper has the ability to receive packets from the Linux network stack and hence it can be used to produce live TCP connections and closed-loop sessions. In this section, we evaluate the performance of Caliper to inject smoothed TCP packets (paced TCP) at precise time intervals and compare the precision of using TCP Iperf traffic with Caliper and a commodity NIC.⁵ As in [1], we use the Precise Software Pacer (PSPacer) [21] package

⁵Note that in this set of experiments we are unable to include SPG due to its open-loop limitation.

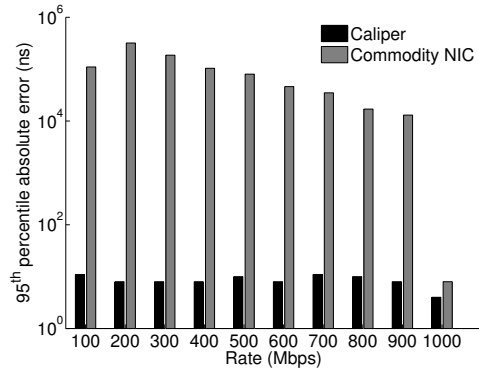


Fig. 7. Comparing the 95th percentile of absolute error ($|D_R - D_M|$) between Caliper and a commodity NIC when injecting TCP packets.

as a loadable kernel module to enforce pacing while using the commodity NIC. The challenges to accomplish precise packet pacing are discussed in [19]. PSPacer paces packets by injecting gap packets between the real packets. By knowing the speed of the link and controlling the number and size of the gap packets, PSPacer controls the timing of packets in software without using timers. The trade-off is that packets are being sent at the line rate through a regular NIC even when the data rate has been limited by pacing. In both experiments, we use an unmodified version of TCP, as implemented by the Linux network stack.

Similar to experiments in Section III-A, we calculate the absolute error ($|D_R - D_M|$) between the measured inter-arrival times (D_M) and the requested inter-transmission times (D_R) corresponding to the requested packet transmission rate (T_R) of Iperf. As illustrated in Figure 7, Caliper improves the 95th percentile of absolute error by almost three orders of magnitude compared to the commodity NIC. Hence, Caliper’s accuracy enables researchers to perform live and time-sensitive network experiments with confidence on the accuracy of packet injection times. As in Figure 4, at 1 *Gbps* rate, the error of both Caliper and the commodity NIC is minimal because packets are sent and received back-to-back.

D. Accuracy of Existing Software Network Emulators

The goal of network emulators is to allow arbitrary networks to be emulated inside a single machine or using a small number of machines. Each packet departure time is calculated based on the packet’s path through the emulated network topology and on interactions with other packets. The result of this process is an ordered list of packets and corresponding departure times. How close the actual transmission times are to these ideal departure times is critical for the precision of the network emulator.

Existing software network emulators have been built on Linux and FreeBSD [17], [22]. To minimize overhead, they process packets in the kernel and use a timer or interrupt firing at a fixed interval to schedule packet transmissions. They effectively divide time into fixed-size buckets, and all packets scheduled to depart in a particular bucket are collected and

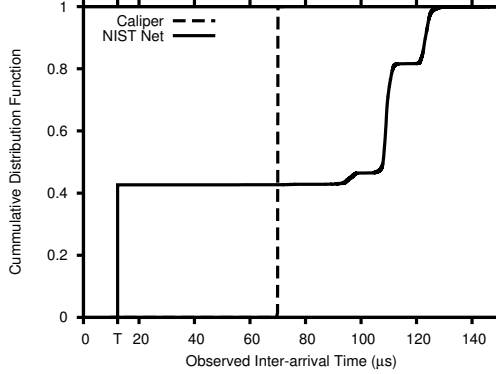



Fig. 8. Effect of NIST Net adding delay to packets sent $70 \mu s$ apart. $T = 12.304 \mu s$ is the time it takes to transmit a single 1472-byte packet at $1 Gbps$.

sent at the same time. Clearly, the bucket size controls the scheduling granularity; i.e., packets in the same bucket will essentially be sent back-to-back.

To quantify the scheduling granularity problem, we focus on the transmission times generated by NIST Net[22], a representative network emulator. Here, we generate MTU-sized UDP packets at a fixed arrival rate using Caliper. The packets are received by a server running NIST Net, pass through the emulated network, and are routed to a third server which measures the resulting packet inter-arrival times. NIST Net is configured to add $100 ms$ of delay to each packet. Although adding a delay to every packet is a simple application of a network emulator, by varying the input packet inter-arrival times, NIST Net’s scheduler inaccuracy is clearly visible.

Figure 8 is a CDF of the measured intervals between packet arrivals in NIST Net’s input and output traffic where the requested packet inter-transmission time is $70 \mu s$. As shown, Caliper accurately delivers packets to NIST Net in the intermediate sever, but NIST Net is unable to preserve the precision. Here a packet is sent by Caliper to NIST Net, and thus should depart from NIST Net, every $70 \mu s$. This interval is smaller than the fixed timer interval used by NIST Net, which has a period of $122 \mu s$ [22], thus NIST Net will either send the packet immediately or in the next timer interval. Consequently, in Figure 8, 40% of the packets are received back-to-back if we consider that it takes just over $12 \mu s$ to transmit a packet of the given size on the wire (the transmission time of a single packet is marked with a “T” on the X-axis). Very few packets actually depart close to the correct $70 \mu s$ interval between them. Most of the remaining intervals are between $100 \mu s$ and $140 \mu s$. Note that the server running NIST Net is using a commodity NIC which also plays a negative role in preserving the packet inter-transmission times.

Even when the interval between arriving packets is larger than NIST Net’s bucket size, the actual packet transmission times are still incorrect. We repeated the same experiment with inter-transmission times of $640 \mu s$ and $700 \mu s$ arrivals and observed that in both cases, 70% of the intervals are actually either $610 \mu s$ or $732 \mu s$, which are multiples of NIST Net’s $122 \mu s$ bucket size. It is only possible for NIST Net to



```

1.#Caliper’s interval in seconds:
2.set caliper_interval 0.001
3.#define the nodes n0, n1, n2, and n3
4.#define the links (n0, n2), (n2, n3), and (n3, n1)
5.#obtain the queue of the specific caliper queue:
6.set caliper_queue_ [$ns simplex-link-op $n2 $n3 queue]
7.#call use-caliper function:
8.$caliper_queue_ use-caliper
9.#set the physical IP/MAC addresses mapping table:
10.$ns insert_nat IP_N2 IP_N3 PORT_N2 PORT_N3 MAC_N2 MAC_N3
11.#Create a UDP agent and attach it to node n0
12.#Create a CBR traffic source and attach it to udp0
13.#set the rate of the CBR source:
14.$cbr0 set interval_ $caliper_interval

```

Fig. 9. Illustration of a simple topology expressed in our integration of `ns-2` with Caliper.

send packets either back-to-back or with intervals that are multiples of $122 \mu s$. When we vary the inter-arrival time of the input traffic between $610 \mu s$ and $732 \mu s$, it only varies the proportion of the output intervals that are either $610 \mu s$ or $732 \mu s$.

The cause of the observed inaccuracies is not specific to NIST Net’s implementation of a network emulator. Any software that uses a fixed-size time interval to schedule packet transmissions will suffer similar failures at small time scales, and the generated traffic will not be suitable for experiments that are sensitive to the exact inter-arrival times of packets. The exact numbers will differ, depending on the length of the fixed interval. To our knowledge, Modelnet[17] is the software network emulator providing the finest scheduling granularity of $100 \mu s$ with a 10KHz timer. Although higher resolution timers exist in Linux that can schedule a single packet transmission relatively accurately, the combined interrupt and CPU load of setting timers for every packet transmission would overload the system. Therefore, our conclusion is that an all-software network emulator executing on a general-purpose operating system requires additional hardware support (such as the one we propose) to produce realistic traffic at very small time scales.

E. Network Emulation by Integrating Caliper with `ns-2`

Simulation and testbed construction represent the two most important methodologies available to network researchers for the design and evaluation of both novel and existing networking elements. Employing an emulation capability in network simulation provides the ability for real-world traffic to interact with a simulation. Since many researchers are already familiar with the Network Simulator `ns-2`, this is a useful tool to test real network devices together with simulated networks. Such integrations will enable researchers to repeat simulation experiments under different link and environment conditions.

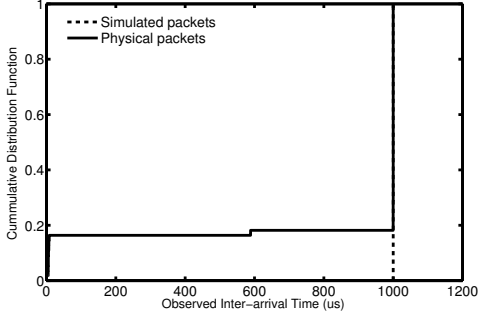


Fig. 10. CDF of simulated packets’ and physical packets’ inter-transmission times when integrating Caliper with `ns-2`.

Compared to previous attempts to connect `ns-2` to a real network [23], our integration of Caliper with `ns-2` enables generating real packets with transmission times that match the `ns-2` simulated times even on very small time scales. In our integration, we mark a particular link in `ns-2`’s simulated network to be mapped to physical link(s). When the simulation starts, the simulated packets are built and traverse the simulated links and nodes until they reach the specific marked queue (`caliper_queue`) connecting the simulation and physical worlds. At this point, Caliper builds real packets and transmits them according to their simulated inter-arrival times. The mapping between simulated node IDs and physical IP/MAC addresses and port numbers is also specified in the simulation configuration file. On the other end of the `caliper_queue`, there is another simulation running, which receives the physical packet, and fires the appropriate simulation event indicating that the corresponding simulated packet has been received.

Our implementation only requires a few additional commands to a simulation program written in the Tcl language which makes it extremely convenient to work with. As an example, Figure 9 illustrates a simple topology expressed in our extension. Lines 1 to 6 define the topology and the nodes conventionally in `ns-2` language. In line 8, we define the link between node 2 (n_2) and node 3 (n_3) as `caliper_queue` so that when packets depart the link’s queue the simulator will send the packet to Caliper’s driver. In line 10, we create a mapping table for the physical IP and MAC addresses corresponding to the physical ends of wire between n_2 and n_3 . Finally in line 14 we assign a traffic source to n_0 with the packet inter-arrival time defined in line 2 (1 *ms*).

Figure 10 compares the CDF of measured physical packets’ inter-transmission times with the simulated packet logs from the `ns-2` trace file. The `ns-2` sources are set to send UDP packets with 1 *ms* inter-transmission times and our integration of Caliper with `ns-2` is able to preserve the inter-arrival times of 80% of the packets. One of the challenges to integrate Caliper with simulation software is synchronizing the simulation time and real-time. In order to keep the difference between real-time and simulation time minimum, we pause the simulation scheduler and as a result, we maintain the same simulation clocks between the simulated world and the real world. We envision that the main cause for the inaccuracy in

20% of the packets is due to our inefficiency in synchronizing real and simulated times.

IV. DISCUSSION AND FUTURE WORK

The limitations of Caliper stem from copying packets between the host computer and the NetFPGA over the 32-bit, 33-MHz PCI bus, which has a bandwidth of approximately 1 *Gbps*. As explained in Section II, the payloads of packets sent by Caliper are usually all zeros, which requires sending only the packet headers over the PCI bus. This is sufficient for network experiments that do not involve packet payloads. A larger body of experiments ignore most of the packet payloads except for a minimal amount of application-layer signaling between sender and receiver. To support this, arbitrary custom data can be added to the start of any packet payload. This additional data is copied to the NetFPGA card and is included in the packet. In the future, we plan to allow a number of predefined packet payloads to be copied to the NetFPGA in a preprocessing phase to later be attached to outgoing packets without the need to repeatedly copy them over the PCI bus. We envision this feature would support many experiments where multiple flows send packets with the same or similar payloads.

We are working on extending the Caliper software on NetThreads to enable more features while preserving the precision of the traffic. For NetThreads applications in general, the maximum achievable packet rate depends on the amount of computations done per packet and hence it also is a function of the packet size (the shortest packets are the worst case leaving less cycle budget per packet: the 125MHz clock allows for 1 cycle per packet byte per processor). Finally, we made both NetThreads and Caliper available as free software to download [13].

V. RELATED WORK

There have been many software- and hardware-based packet generators presented in the literature. Some of the software workload generators try to characterize network traffic by empirically deriving models for web traffic [24], [25], or other network applications, such as TELNET, SMTP, NNTP, and FTP [26]. Cao *et al.* [27] model the HTTP traffic and parameterize the network characteristics such as the round-trip times at the clients rather than capturing it empirically. Netspec [28] builds source models to generate traffic for TELNET, FTP, HTTP, voice and video.

One popular way to generate traffic for testbeds is through packet traces from existing networks. RAMP [29] generates high bandwidth traces using a simulation environment involving source-level models for HTTP and FTP. Rupp *et al.* [30] introduce a packet trace manipulation framework for testbeds. They present a set of rules to manipulate a given network trace, for instance, stretch the duration of existing flows, add new flows, change packet size distributions, etc. Hernandez *et al.* [31] generate realistic TCP workloads using a one-to-one mapping of connections from the original trace to the test environment. Swing [15] can create responsive, closed-loop traffic with similar burstiness characteristics on multiple time scales to existing traces by estimating wide-area

network characteristics. Another popular way to create realistic traffic is to use network emulators. For example, Swing uses ModelNet [17] to emulate a network of links each with its own bandwidth, delay and drop probability. Unfortunately, relying on network emulators has its own limitations. The network is emulated in software, and the position of packets within the network is only updated e.g. 10000 times per second or once every 100 μ s for ModelNet. Thus, the packets sent do not have high precision timings as roughly 8 MTU sized packets can be transmitted at 1 Gbps in 100 μ s. Using the discrete event simulator ns-2 as the network emulator also suffers from similar timing issues. Mahrenholz *et al.* [32] recommend several modifications to ns-2 to improve the accuracy of its network emulation feature. The NetFPGA-based packet generator introduced by Covington *et al.* [9] can reliably replay a trace file and capture packets at Gbps line rate. However, as mentioned in Section I the traces are based on prior recording and it would be difficult to extrapolate them to closed-loop traffic and other workload/topology scenarios.

There are a number of other available software tools for traffic generation such as Harpoon [33] and tcplib [34]. However, they are all designed to match the property distributions of a trace at a coarse granularity and none attempt to guarantee the behavior of traffic at short time scales. They ignore the unavoidable timing issues introduced by the users' hardware and OS choices. Our efforts are complementary to the above mentioned works as we focus only on constructing real packets and providing exact transmission times. To the best of our knowledge, we present the first framework for generating precise closed-loop traffic providing guarantees for inter-transmission times at very short time scales.

VI. CONCLUSIONS

Generating realistic traffic in network testbeds is challenging yet crucial for performing valid measurement experiments. Software network emulators schedule packet transmission times in software, incurring unavoidable inaccuracy for inter-transmission intervals in the sub-millisecond range – hence they are insufficient for experiments sensitive to the inter-arrival times of packets. In this paper we present Caliper, a precise and responsive traffic generator built on NetFPGA board. Caliper allows packets generated on the host computer to be sent with extremely accurate inter-transmission times and is designed to be integrated with existing software traffic generators and network emulators. We demonstrate Caliper's precision and integration with existing software to generate traffic that is realistic and accurate at almost all time scales. In our experiments, the maximum error that Caliper incurs is around 8 ns which is the NetFPGA's clock cycle time and also our measurement resolution. Overall, Caliper allows researchers to perform experiments that were previously infeasible.

REFERENCES

[1] Beheshti, N., Ganjali, Y., Ghobadi, M., McKeown, N., Salmon, G.: Experimental study of router buffer sizing. In: IMC. (2008) 197–210
 [2] Chen, Y., Griffith, R., Liu, J., Katz, R.H., Joseph, A.D.: Understanding TCP incast throughput collapse in datacenter networks. In: WREN. (2009) 73–82

[3] Kuzmanovic, A., Knightly, E.W.: Low-rate TCP-targeted denial of service attacks. In: Proceedings of ACM SIGCOMM. (2003) 75–86
 [4] Beheshti, N., Ganjali, Y., Ghobadi, M., McKeown, N., Naous, J., Salmon, G.: Performing time-sensitive network experiments. In: ANCS. (2008) 127–128
 [5] Website: Avalanche Traffic Generator <http://www.spirentcom.com/>.
 [6] Website: Ixia: Leader in IP performance testing <http://www.ixiacom.com/>.
 [7] Website: Endace <http://www.endace.com/>.
 [8] Prasad, R., Dovrolis, C., Thottan, M.: Evaluation of Avalanche traffic generator. Technical report, Georgia Institute of Technology (2007)
 [9] Covington, G.A., Gibb, G., Lockwood, J.W., McKeown, N.: A packet generator on the NetFPGA platform. In: FCCM. (2009) 235–238
 [10] Naous, J., Gibb, G., Bolouki, S., McKeown, N.: NetFPGA: reusable router architecture for experimental research. In: PRESTO. (2008) 1–7
 [11] Website: Iperf <http://sourceforge.net/projects/iperf/>.
 [12] Labrecque, M., Steffan, J.G., Salmon, G., Ghobadi, M., Ganjali, Y.: NetThreads: Programming NetFPGA with threaded software. In: NetFPGA Developers Workshop. (2009)
 [13] Website: Caliper's wiki page <http://netfpga.org/foswiki/bin/view/NetFPGA/OneGig/PreciseTrafGen>.
 [14] Ghobadi, M., Salmon, G., Ganjali, Y., Labrecque, M., Steffan, J.G.: Caliper: Precise and responsive traffic generator. Technical Report TR10-UT-SNL-10-10-00, University of Toronto (2010)
 [15] Vishwanath, K.V., Vahdat, A.: Realistic and responsive network traffic generation. In: SIGCOMM. (2006) 111–122
 [16] Website: The Network Simulator - ns-2 <http://www.isi.edu/nsnam/ns/>.
 [17] Vahdat, A., Yocum, K., Walsh, K., Mahadevan, P., Kostić, D., Chase, J., Becker, D.: Scalability and accuracy in a large-scale network emulator. SIGOPS (2002) 271–284
 [18] Zhang, L., Shenker, S., Clark, D.D.: Observations on the dynamics of a congestion control algorithm: The effects of two-way traffic. In: CCR. (1991) 133–147
 [19] Beheshti, N., Ganjali, Y., Ghobadi, M., McKeown, N., Naous, J., Salmon, G.: Performing time-sensitive network experiments. Technical Report TR08-UT-SNL-09-10-00, University of Toronto (2008)
 [20] Katabi, D., Blake, C.: Inferring congestion sharing and path characteristics from packet interarrival times. Technical report, MIT (2001)
 [21] Takano, R., Kudoh, T., Kodama, Y., Matsuda, M., Tezuka, H., Ishikawa, Y.: Design and evaluation of precise software pacing mechanisms for fast long-distance networks. In: PFLDNet. (2005)
 [22] Carson, M., Santay, D.: NIST Net: a Linux-based network emulation tool. SIGCOMM Computer Communication Review 33(3) (2003) 111–126
 [23] Fall, K.: Network emulation in the Vint/NS simulator. In: ISCC'99. (1999) 244
 [24] Barford, P., Crovella, M.: Generating representative web workloads for network and server performance evaluation. SIGMETRICS Performance Evaluation Review 26(1) (1998) 151–160
 [25] Mah, B.A.: An empirical model of HTTP network traffic. In: INFOCOM '97. (1997) 592
 [26] Paxson, V.: Empirically derived analytic models of wide-area TCP connections. IEEE/ACM Transactions on Networking 2(4) (1994) 316–336
 [27] Cao, J., Cleaveland, W.S., Gao, Y., Jeffay, K., Smith, F.D., Weigle, M.: Stochastic models for generating synthetic HTTP source traffic. In: INFOCOMM. (2004)
 [28] Lee, B.O., Frost, V.S., Jonkman, R.: NetSpec 3.0 source models for Telnet, FTP, voice, video and WWW traffic. Technical Report ITTC-TR-10980-19, University of Kansas (1997)
 [29] Sen, S., Wang, J.: Analyzing peer-to-peer traffic across large networks. In: IEEE/ACM Transactions on Networking. (2002) 219–232
 [30] Rupp, A., Dreger, H., Feldmann, A., Sommer, R.: Packet trace manipulation framework for test labs. In: IMC'04. (2004)
 [31] Hernandez-campos, F., Donelson, F., Jeffay, S.K.: Generating realistic TCP workloads. In: Computer Measurement Group International Conference. (2004) 273–284
 [32] Mahrenholz, D., Ivanov, S.: Real-time network emulation with ns-2. In: DS-RT'04. (2004)
 [33] Sommers, J., Barford, P.: Self-configuring network traffic generation. In: IMC '04. (2004) 68–81
 [34] Danzig, P.B., Jamin, S.: tcplib: A library of TCP internetwork traffic characteristics. Technical Report USC-CS-91-495, University of Southern California (1991)