Traffic Trace Artifacts due to Monitoring Via Port Mirroring

Jian Zhang and Andrew Moore
Deptartment of Computer Science, Queen Mary, University of London

***Abstract*—Port-mirroring techniques are supported by many of today's medium and high-end Ethernet switches. The ubiquity and low-cost of port mirroring has made it a popular method for collecting packet traces. Despite its wide-spread use little work has been reported on the impacts of this monitoring method upon the measured network traffic. In particular, we focus upon each of delay and jitter (timing difference), packet-reordering, and packet-loss statistics. We compare the port-mirroring method with inserting a passive TAP (Test Access Point), such as a fibre splitter, into a monitored link. Despite a passive TAP being transparent to monitored traffic, port-mirroring popularity arises from its limited set-up disruption, and (potentially) easier management. This paper documents experimental comparison of traffic using the passive TAP and port-mirroring functionality, and shows that port-mirroring will introduce significant changes to the inter-packet timing, packet-reordering, and packet-loss — even at very low levels of utilisation.**

I. Introduction

Traffic monitoring is crucial to operating all IP networks for many reasons. Whether you want to monitor for security threats, for troubleshooting problems, or for analysis purposes, you need a reliable way to see all of the traffic. Regardless of what analyser or intrusion detection solution you choose, you must decide on a method to give your monitoring equipment physical access to the network traffic.

There are three common ways to fulfill this:

- Attach an analysis or monitoring device to a passive TAP (Test Access Port) inserted into the target link(s).
- Redirect traffic from the target link to analysis or monitoring equipment attached via a network switch. The network switch duplicates the required traffic from the monitored data-stream, sending the replicas out a dedicated interface/port (in Cisco terminology, a Switch Port ANalyser, or SPAN, and described as port mirroring hereafter in this paper).
- A hybrid technique approach using a dedicated monitor inserted into target link(s). The port aggregator replicates traffic from one or more of these links and transmits the replicated traffic through a separate interface to the monitoring and analysis equipment. We do not analyse the port-aggregator approach specifically in this paper but we consider it would share some of the disadvantages of the port-mirroring approach with some limited protection against packet loss provided by the port aggregator's built-in memory buffer.

This work has been supported by the EPSRC projects GR/S93714/02 and GR/T10510/02. Corresponding author: andrew.moore@dcs.qmul.ac.uk

The advantage of the port-mirroring solution is its cost as this feature is included for free with most managed switches available on the market, it is relatively easy-to-use and may be configured remotely. These advantages indicate why the port-mirroring technique has been a popular method for collecting packet traces for various purposes ([1], [2], [3]). However, the traffic trace artifacts incurred due to monitoring via port mirroring has not been well studied, which renders the risk that analysis results based on the traffic trace collected via port mirroring method could be biased or simply incorrect. In this paper, our focus is to answer three fundamental questions about the impact of port-mirroring in terms of timing, reordering and loss:

1) The port-mirroring method needs to buffer packets from the monitored link until they can be sent over the mirroring link. What, then, are the exact differences between the timing of the original packet streams on the monitored link and the timing of their counterparts on the mirroring link? This is particularly important for any analysis which requires the packet inter-arrival times(IATs).
2) Are the original packet sequences on the monitored link maintained when the streams are put onto the mirroring link by port mirroring method? In other words, does the reordering occur on the mirroring link, and, if it does, what are the reordering statistics? This question is important to any analysis of packet orders and the the reordering issue also affects the above timingalso impacts IATs.
3) Since the buffer at the mirrored port can be overflowed, the port-mirroring method might lose packets. In this paper, we quantify the properties of packet loss arising due to monitoring via port mirroring.

On the other hand, the passive TAP technique has been regarded as most effective and accurate measurement method due to its transparency and robustness (TAPs never drop packets regardless of network conditions). We carried out a number of experiments to try to answer the above questions by comparing the original packet traces collected via the passive TAP and the mirrored traces collected via the port mirroring. Firstly, we constructed the experiment network by using two Cisco Catalyst switches (2950) and a number of workstations.

Three types of packet traces were generated and fed into the network to investigate the timing, reordering and loss statistics under different traffic patterns and traffic loads. Then we repeated our experiments in a second test network constructed using one Cisco Catalyst switch 2950 and one HP switch 2824 to investigate whether the observed statistics are also valid for other switches. Our main findings about the traffic trace artifacts in terms of timing difference, packet-reordering and packet-loss[1] due to monitoring via port mirroring are as follows:

- The traffic trace artifact in terms of timing difference exists under different traffic load level using different traffic patterns on different switches when the traces are collected via port-mirroring method.
- Under a certain traffic load level, the statistics of timing difference exhibit quite-similar characteristics for the aggregated traffics in either direction of a monitored link.
- The mean value of timing-difference statistics increases with an increase in the traffic load on the monitored link and, although a certain timing difference value could be dominant to some extent under a traffic load level, the values of timing difference will span a range of microseconds even under lower traffic loads;
- When the traces are collected via port mirroring method, a significant percentage of packets get reordered for the aggregated traffic in both directions of a monitored link even under low levels of utilisation.
- Packet reordering is observed for the aggregated traffic in a single direction of a monitored link especially under higher traffic loads on the link.
- The reordering time and reordering number statistics exhibit different characteristics for the traffic streams with different packet sizes. And the characteristics of reordering-time and reordering-number for the aggregated traffic in both directions of a monitored link are consistent under different traffic loads.
- Both the statistics of reordering time and reordering number span a range of values although a certain value can be dominant to some extent.

Using the port-mirroring method, M. Arlitt *et al.*[1] reported their estimated number of lost packets by watching for the gaps in the TCP sequence numbers where they collected the packet traces via the port-mirroring method. As the paper's main purpose was the analysis of TCP reset behavior in the Internet, it did not provide more insights of the losses incurred using port-mirroring technique. Although some white papers [4], [5] predict the possible impacts on traces collected via port mirroring, none of them has given the experiment data to support their predictions. To the best of our knowledge, our paper is the first effort to make a comprehensive study on traffic-trace artifacts incurred due to monitoring via port mirroring.

Paper orgainization follows this, in Section 2, the methods

[1]Section 2 provides the definitions of "timing difference", "packet reordering" and "packet loss" used in this paper.

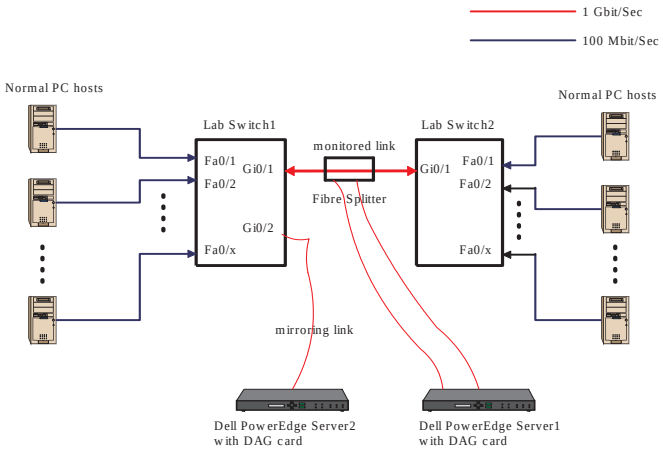

Fig. 1. The network and measurement topology used in our experiments.

used in the network experiments to investigate the impacts of port-mirroring on collected traces are discussed. The experiment details and the results are presented in Section 3. Section 4 discusses some implications of our findings on the usage of the port mirroring technique for the future network measurement and monitoring. Section 5 provides concluding remarks and directions for future work.

II. Methods to Investigate the Traffic Trace Artifacts Incurred Via Port-Mirroring

This section discusses the measurement methods we used in our experiments to investigate the traffic-trace artifacts incurred due to monitoring via port mirroring. We focus on examining whether any traffic-trace artifact: timing difference, packet-reordering and packet-loss, will occur when traffic traces are collected using port-mirroring method and then documentg the artifact characteristic. The passive TAP technique has been regarded as the most effective and accurate method to collect traffic traces due to its robustness and transparency to the measurement data. Therefore, we study the traffic trace artifacts by comparing the port-mirroring method with inserting a passive TAP (Test Access Point), such as a fiber splitter, into a monitored link.

The network and measurement topology used in our experiments is shown in Fig 1. Two ethernet switches (lab-switch1 and lab-switch2) are linked via a 1Gbps port (Gi0/1, Fig 1) and each of them connects a number of PC hosts, which act as the traffic generators, its 10/100Mbit ports. Moreover, lab-switch1 is configured to mirror the traffics traversing its port Gi0/1 in both directions to another 1Gbps port (Gi0/2, Fig 1), i.e., the port mirroring feature is enabled. Thus, the monitored link is the 1 Gbps point-to-point link between lab-switch1 and lab-switch2. Two Dell PowerEdge 2850 servers are deployed as the capture engine and each of them is equipped with a DAG 4.5G2 card. The clocks of the two DAG cards are locked together and synchronised to the clock of one PowerEdge server. A fiber splitter is inserted into the monitored link to enable a DAG card to capture the original traffic streams on the link and another DAG card to capture the traffic streams on the mirroring link.

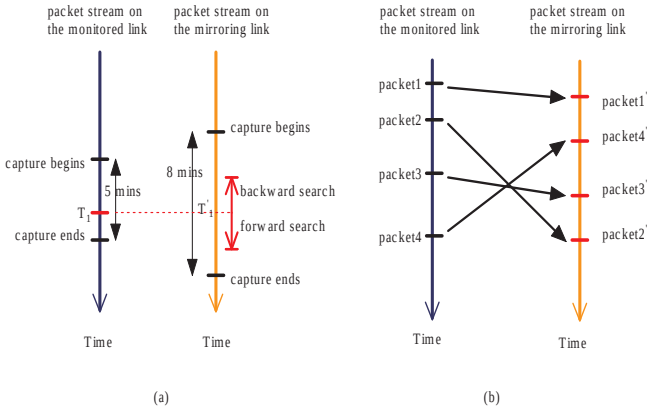

Fig. 2. The measurement methods in our experiments.

To further investigate whether the statistics of timing difference, packet-reordering and packet-loss will be affected by packet sizes, we used three different traffic patterns which are used to generate traffics by those PC hosts. The first traffic pattern is an 8-minute synthetic packet trace consisting of packets which are distinguishable from each other and have the same packet size of 1500 octets (maximum ethernet packet size). The second traffic pattern is an 8-minute synthetic packet trace consisting of packets which are distinguishable from each other and have the same packet size of 46 octets (minimum ethernet packet size). The third traffic pattern is a 15-minute real ethernet trace [2] with variable packet sizes (called the *real* ethernet trace in this paper) and we replaced only IP IDs of the packets in the trace to make distinguishable from each other. Then, during every experiment, a number of PC hosts connected to the switches will replay one of the three packet patterns using the tool *tcpfire* [6] to generate the network traffics. Note that, when a traffic pattern is copied to a PC host for replaying, the source MAC addresses of the packets in the traffic pattern are changed to the PC NIC's MAC address so that the traffic stream generated by different hosts is distinguishable. Moreover, we (hard-)configure one of the spare ports on lab-switch1 to be the sink of all the packets received from lab-switch2 and one of lab-switch2's spare ports to be the sink of all the packets received from lab-switch1. Thus, the PC hosts connected to the switches do not need to take careful of receiving packets.

The way we measure the terms "timing difference", "packet-reordering" and "packet-loss" in our experiments is explained below. The exact meanings of these terms in this paper are given by examples. As illustrated in Fig. 2(a), we begin capturing the traffic stream on the mirroring link first and finish it later to guarantee that the traffic stream captured on the monitored link will be found on the mirroring link. After the above capture processes, two traffic trace files will be obtained: one trace file contains the traffic data on the monitored link (called the original traffic trace file [3] hereafter) and another contains the traffic data on the mirroring link (called the mirroring traffic trace file hereafter). Then, we will try to find every single packet of the original traffic trace file in the mirroring trace file by comparing the packets of the two trace files. For a packet with a timestamp T_1 of the original trace file, we try to find it by searching around T_1 for a certain period [4] backwards and forwards in the mirroring trace file (see Fig. 2(a)) and for the case that the packet is not found after the above searching, it will be counted as the packet-loss incurred due to monitoring via port-mirroring, otherwise, its counterpart with a timestamp like $T_1^{'}$ will be found in the mirroring trace file. The timing difference of a packet and its counterpart incurred due to monitoring via port mirroring is defined as follows: Time Difference = $T_1^{'} - T_1$.

In addition, we define a packet (e.g., *packet A*) to get reordered if the packet, which appears earlier than another packet (e.g., *packet B*) on the monitored link, appears later than that packet (*packet B*) on the mirroring link. Obviously, a number of packets which appear later than *packet A* on the monitored link can cause *packet A* to get reordered. Below we call the packet, which not only causes *packet A* to get reordered but also the timestamp of which counterpart found on the mirroring link is the smallest among the counterparts of those packets that cause *packet A* to be reordered, as the causing-reordering-packet of *packet A*. Then, the reordering time of *packet A* is defined as the subtraction of the timestamp of the counterpart of *packet A* and the timestamp of the counterpart of the causing-reordering-packet of packet A, and the reordering number of *packet A* is defined as the number of packets which cause *packet A* to get reordered.

For instance, we assume that packet1, packet2, packet3 and packet4 of the original traffic trace file are found to match the packets packet1$^{'}$, packet2$^{'}$, packet3$^{'}$ and packet4$^{'}$ of the mirroring traffic trace file respectively (see Fig. 2(b)). The timestamps of packet1, packet2, packet3 and packet4 are denoted as $T_{packet1}$, $T_{packet2}$, $T_{packet3}$ and $T_{packet4}$ and it is assumed that $T_{packet1} \leq T_{packet2} \leq T_{packet3} \leq T_{packet4}$. Moreover, the timestamps of packet1$^{'}$, packet2$^{'}$, packet3$^{'}$ and packet4$^{'}$ are denoted as $T^{'}_{packet1}$, $T^{'}_{packet2}$, $T^{'}_{packet3}$ and $T^{'}_{packet4}$ and it is assumed that $T^{'}_{packet1} \leq T^{'}_{packet4} \leq T_{packet3}{}^{'} \leq T^{'}_{packet2}$). Based on the above definitions, packet4 and packet3 cause packet2 get reordered and packet4 also causes packet3 gets reordered. Furthermore, if we assume that packet4 is the causing-reordering-packet of both packet2 and packet3, then the reordering time of packet2 is $T^{'}_{packet2} - T^{'}_{packet4}$ and the reorder number of packet2 is 2, and the reordering time of packet3 is $T^{'}_{packet3} - T^{'}_{packet4}$ and the reordering number of packet3 is 1.

III. Experiments and Results

This section presents the results of network experiments we conducted to investigate the traffic-trace artifacts incurred due to monitoring via port-mirroring method. Due to the

[2]This real ethernet trace was taken from a day-time trace of a 100-user University CS-department.

[3]the packets which are not generated by PC hosts using one traffic pattern (e.g., packets originating from the switches) will first be filtered out by comparing the source and destination MAC addresses.

[4]This search time period has been tested and validated in our experiments and it is different for different traffic patterns.

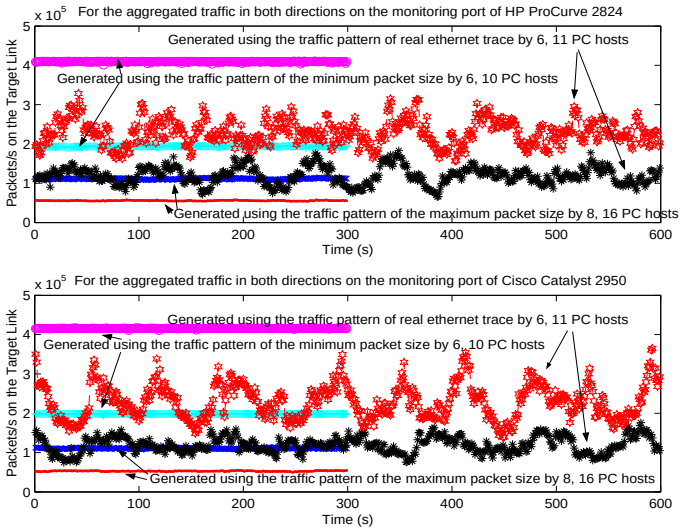

Fig. 3. The packet distributions of the traffic captured on the monitored link.

page limitation, we report only a subset of our experimental results here, where two typical traffic load levels (moderately loaded: around 60%-70% utilisation of the monitored link and moderately overloaded: around 130%-140%) utilisation of the monitored link) were generated using each of the three traffic patterns [5] respectively. In addition, as mentioned in Section 2, the PC hosts connected to the lab-switch1 and lab-switch2 (see Fig. 1) are responsible for generating traffics into the network. In our experiments, the traffic loads generated on the monitored link were altered by changing the number of PC hosts connected to the switches. Furthermore, the traffic-trace artifacts arising due to port mirroring were studied on two ethernet switches: Cisco Catalyst Switch 2950 and HP ProCurve 2824, to illustrate our findings are valid across different switches.

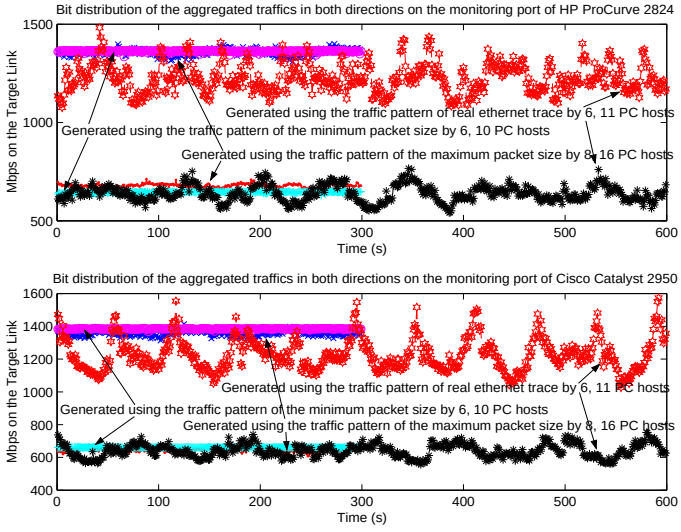

Fig. 4. The bit distributions of the traffics captured on the monitored link.

First of all, two Cisco Catalyst switches (2950G-48-EI), each of which has 48 10/100 Ethernet ports and 2 GBIC module slots, were equipped as the lab-switch1 and lab-switch2 shown in the network and measurement topology (see Fig. 1), respectively. Secondly, a HP ProCurve switch 2824, which has 24 10/100/1000Base-T RJ-45 ports and 4 mini-GBIC Gigabit ports, was deployed as lab-switch1 and a Cisco Catalyst switch 2950G-48-EI as lab-switch2 and the experiments were repeated on them. Moreover, we will call the direction from lab-switch1 to lab-switch2 on the monitored link the TX direction and the opposite one from lab-switch2 to lab-switch1 the RV direction hereafter throughout this paper. The packet and bit distributions of the traffics captured on the monitored link in our experiments are shown in Figs. 3-4. It can be seen from the figures that the traffics are almost static when generated using the traffic patterns with the maximum and minimum ethernet packet sizes, whereas the traffics are much more bursty when generated using the traffic pattern of real ethernet trace. This is what we expected. Note that, for the traffic patterns with the maximum and minimum ethernet packet sizes, a 5-mins traffic trace was captured on the monitored link and a 8-mins traffic trace was captured on the mirroring link, respectively; for the traffic pattern using the real ethernet trace, a 10-mins traffic trace was captured on the monitored link and a 12-mins traffic trace was captured on the mirroring link.

After capturing the traffic traces on the monitored and mirroring links, we compare them using the methods introduced in Section 2. Then, the set of packets which appear on the monitored link and are also found in the mirroring link can be obtained, so is the set of packets which appear on the monitored link but are not found in the mirroring link. For each found packet, its timing difference, reordering time and reordering number will be calculated based on their definitions in Section 2. Next, the statistics of timing difference will be calculated as follows: the time interval between the minimal and maximal values of timing difference of all found packets is divided into a number of 1-microsecond bins and then the number of packets of each timing value is counted. Importantly, the normalised count is used in this paper: the normalised count equals the count of packets for which the timing-difference values fall into a specific 1-microsecond bin divided by the total number of packets observed on the mirroring link. The statistics for reordering time and reordering number are calculated using this method.

A. Maximum-size ethernet packet size

This subsection reports the experiments where the traffic pattern with the maximum ethernet packet size was replayed by PC hosts to generate traffics into the switches.

Figs. 5-6 show the statistics of timing difference of the aggregated traffics in either TX or RV direction [6] incurred due to monitoring via port mirroring for this traffic pattern. It can be seen clearly from Figs. 5-6 that the traffic trace artifact in terms of timing difference exists under different

[5] See Section 2 for the descriptions of the traffic patterns used in our experiments

[6] This means that the traffics traversing the monitored link are filtered and only the traffics in either TX or RV direction are considered

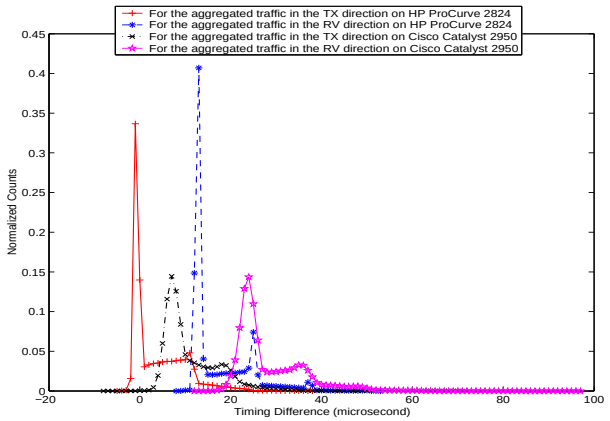

Fig. 5. Packet reordering time — maximum-size ethernet trace (8 hosts.)

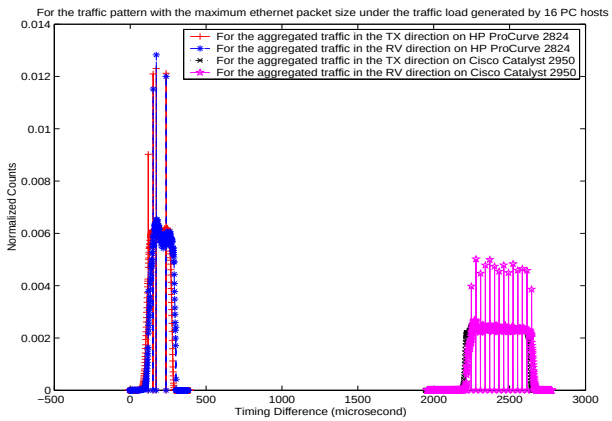

Fig. 6. Packet reordering time — maximum-size ethernet trace (16 hosts.)

TABLE I
THE PERCENTAGE OF REORDERED PACKETS FOR THE AGGREGATED TRAFFICS IN EITHER TX OR RV OR BOTH DIRECTIONS

	Cisco Catalyst 2950		HP ProCurve 2824	
The percentage of packets which get reordered	Under traffic load of 8 PC hosts	Under traffic load of 16 PC hosts	Under traffic load of 8 PC hosts	Under traffic load of 16 PC hosts
Both directions	21.080%	39.071%	17.099%	35.732%
For TX only	0.00055%	0.00103%	0%	0.00039%
For RV only	0%	0.00123%	0%	0.00096%

traffic load level on different switches for this traffic pattern. Moreover, under a certain traffic load on a switch, the statistics of timing difference exhibit very similar characteristics for the aggregated traffics in either TX or RV direction. Furthermore, the mean value of timing difference increases with the increase of the traffic load on the monitored link and the values of timing difference will span a range of microseconds even under lower traffic loads. For the Catalyst Switch 2950, the mean value of timing difference can reach around 2400 microseconds when the traffic load is around 1360Mbps on the target link.

In addition, Figs. 5-6 show that the mean value of timing difference for the TX direction is a little smaller (around 20 microseconds) than the one for RV direction. We believe it is due to the fact that the traffic stream in the TX direction can be put onto the mirroring link by lab-switch1 immediately while the traffic stream in the RV direction needs to travel to lab-switch1 on the monitored link first. We conjecture that the delay is the cumulation of both marshalling delays and the propogation delays as this grade of switch would implement a store-

Figs. 7-8 and Table 1 show the statistics of reordering time and reordering number incurred due to monitoring via port mirroring for this traffic pattern. For the aggregated traffic in both directions [7], it can be noticed that a significant percentage of packets get reordered. Moreover, Figs. 7-8 indicate that there is a step of around 12 microseconds between consecutive reordering time values. We believe it might be due to the fact that this traffic pattern with the maximum ethernet packet size (1500 octets) needs around 12 microseconds to get transmitted

[7]This means that the traffics traversing the monitored link in both directions are considered

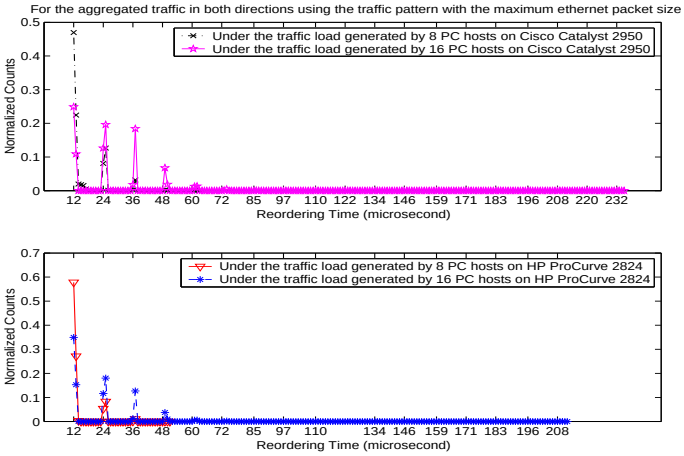

Fig. 7. Packet reordering time — maximum-size ethernet trace.

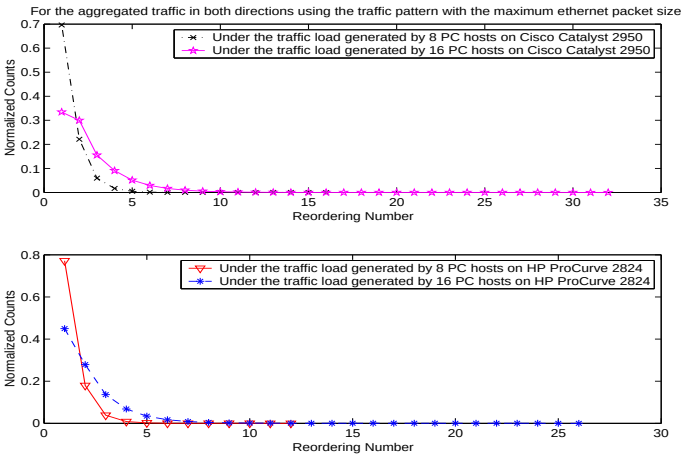

Fig. 8. Packet reordering number — maximum-size ethernet trace.

through the 1 Gigabit link. Furthermore, the characteristics of reordering time and reordering number of the aggregated traffics in both directions are consistent under different traffic loads on different switches for this traffic pattern. Fig. 7-8 also show that the reordering time and reordering number span a range of values although a certain value might be dominant to some extent. Surprisingly, we also observe that packet-reordering exists for the aggregated traffic in either TX or RV direction especially under higher traffic loads (see Table 1). Due to the limited page space, the statistics of the reordering

time and reordering number of the aggregated traffic in a single direction (either TX or RV direction) are not included in this paper. However, the above observations for reordering time and reordering number are also valid for them.

The statistics of packet loss incurred due to monitoring via port mirroring is as follows: there is no packet loss under the traffic load generated by 8 PC hosts; there are 26.738% and 26.965% packets on the monitored link lost on the mirroring link under the traffic load generated by 16 PC hosts when port mirroring is made on Cisco Catalyst Switch 2950 and HP ProCurve 2824, respectively.

B. Minimum-size ethernet packet size

This subsection reports the experiments where the traffic pattern with the minimum ethernet packet size was replayed by PC hosts to generate traffics into the switches.

The statistics of timing difference of the aggregated traffics in either TX or RV direction are presented in Fig. 9-10. It is also clearly indicated in the two figures that the traffic-trace artifact in term of timing difference exists under different traffic load level on different switches for this traffic pattern. Under a certain traffic load on a switch, the statistics of timing difference also exhibit very similar characteristics for the aggregated traffics in either TX or RV direction. Furthermore, the mean value of timing difference also increases with the increase of the traffic load on the monitored link and the values of timing difference will span a range of microseconds even under lower traffic loads although the mean values of timing difference are much smaller for this traffic pattern than the ones for the traffic pattern with the maximum ethernet packet size. The phenomenon that the mean value of timing difference for the TX direction is a little smaller than the one for RV direction is also maintained for this traffic pattern.

Figs. 11-12 and Table 2 show the statistics of reordering time and reordering number incurred due to monitoring via port mirroring for this traffic pattern. It is also noticed that a significant percentage of packets get reordered for the aggregated traffics in both directions. We also observe that packet-reordering exists for the aggregated traffic in either TX or RV direction under higher traffic loads. Moreover, the characteristics of reordering time and reordering number are consistent under different traffic loads on the same switch for this traffic pattern and the reordering time and reordering number also span a range of values although a certain value might be dominant to some extent. We observed no packet-reordering existing for the aggregated traffic in a single direction on HP ProCurve 2824. However, packet-reordering exists for the aggregated traffic in the TX direction on Cisco Catalyst Switch 2950 for this traffic pattern (see Table 2).

When comparing the characteristics of the timing difference and reordering for this traffic pattern with the one for the traffic pattern with the maximum ethernet packet size, we can conclude that the traffic trace artifacts in terms of timing difference and packet reordering exhibit different characteristics for the traffic patterns with different packet sizes and the mean values of their statistics will normally increase with the increase of

TABLE II
THE PERCENTAGE OF REORDERED PACKETS FOR THE AGGREGATED TRAFFICS IN EITHER TX OR RV OR BOTH DIRECTIONS

	Cisco Catalyst 2950		HP ProCurve 2824	
Percentage of packets which get reordered	Under traffic load of 6 PC hosts	Under traffic load of 10 PC hosts	Under traffic load of 6 PC hosts	Under traffic load of 10 PC hosts
Both directions	8.842%	45.702%	11.624%	24.782%
For TX only	0%	29.249%	0%	0%
For RV only	0%	0%	0%	0%

mean packet size in the observed traffics. Surprisingly, we observed that there is no loss for this traffic pattern under the two traffic load levels when port mirroring is made either on Cisco Catalyst Switch 2950 or HP ProCurve 2824. We believe it might be due to the advantages of short packet size: the internal buffer of the port-mirroring port (e.g., Gi0/2, see Fig. 1) might not easily get overflowed for this minimum ethernet packet size even if the average traffic load exceeded the bandwidth of the monitored link.

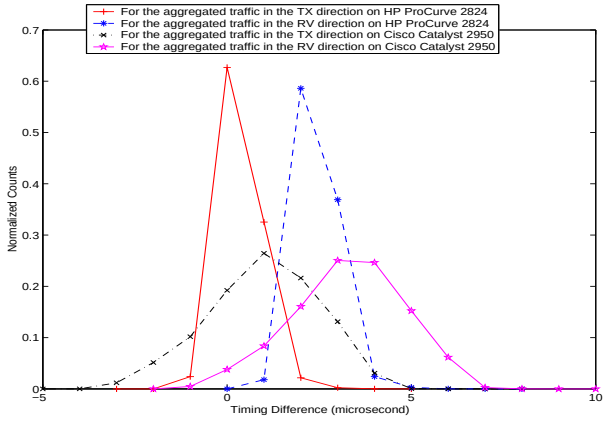

Fig. 9. Packet reordering time — minimum-size ethernet trace (6 hosts.)

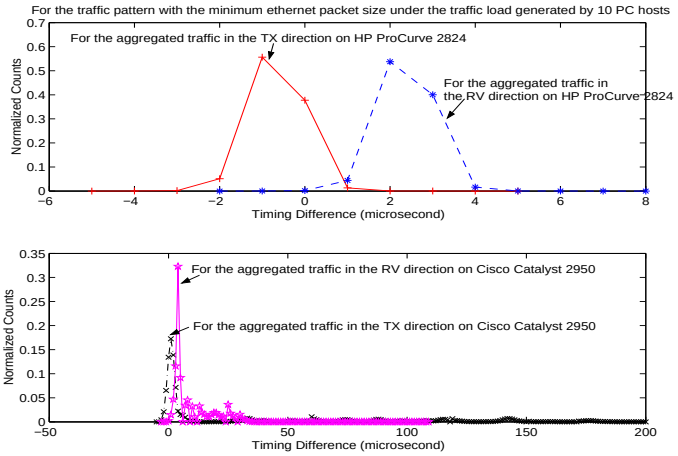

Fig. 10. Packet reordering time — minimum-size ethernet trace (10 hosts.)

C. Real ethenet packet trace

This subsection reports the experiments where the traffic pattern of real ethernet trace was replayed by PC hosts to generate traffics into the switches.

The statistics of timing difference of the aggregated traffics in either TX or RV direction are presented in Fig. 13-14. When combining the results in the above two subsections, we can concluded the following findings about the traffic trace artifact

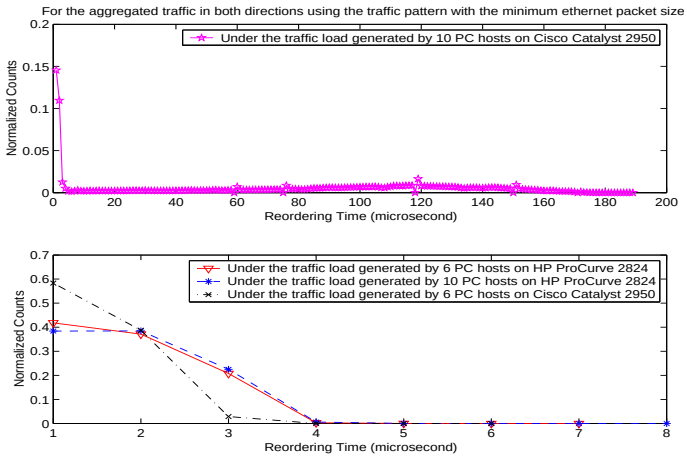

Fig. 11. Packet reordering time — minimum-size ethernet trace.

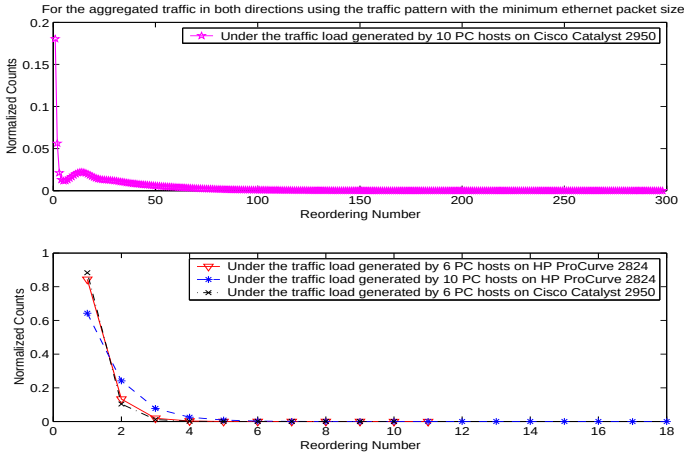

Fig. 12. Packet reordering number — minimum-size ethernet trace.

in terms of timing difference incurred due to monitoring via port mirroring:

- The traffic trace artifact in terms of timing difference exists under different traffic load level using different traffic patterns on different switches when the traces are collected via port mirroring method.
- Under a certain traffic load generated using a traffic pattern, the statistics of timing difference exhibit very similar characteristics for the aggregated traffics in either TX or RV direction.
- The mean value of timing difference statistics increases with the increase of the traffic load on the monitored link and it also increases with the increase of the mean packet size in the observed traffic trace. Moreover, the mean value of timing difference for the TX direction is a little smaller than the one for RV direction.
- The values of timing difference will span a range of microseconds even under lower traffic loads.

Figs. 15-16 and Table 3 show the statistics of reordering time and reordering number incurred due to monitoring via port mirroring for this traffic pattern. When combining the results in the above two subsections, the findings about the traffic trace artifact in terms of packet reordering incurred due to monitoring via port mirroring can be concluded as follows:

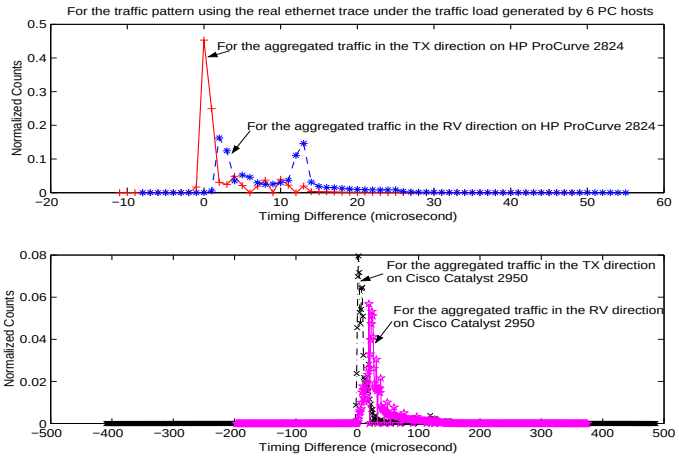

Fig. 13. Packet reordering time — real ethernet trace (6 hosts.)

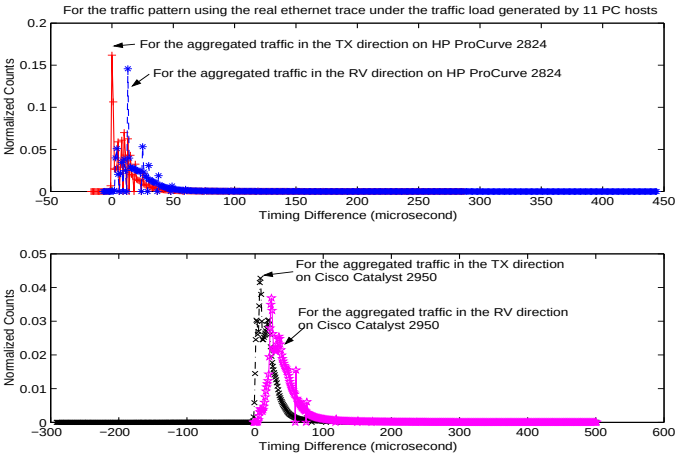

Fig. 14. Packet reordering time — real ethernet trace (11 hosts.)

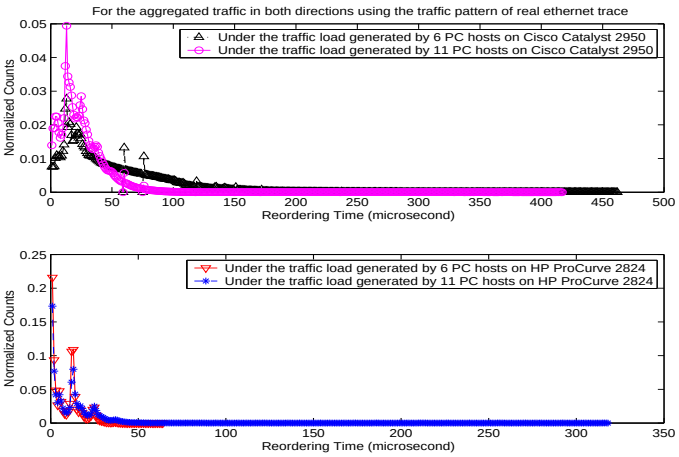

Fig. 15. Packet reordering time — real ethernet trace.

TABLE III
THE PERCENTAGE OF REORDERED PACKETS FOR THE AGGREGATED TRAFFICS IN EITHER TX OR RV OR BOTH DIRECTIONS

	Cisco Catalyst 2950		HP ProCurve 2824	
Percentage of packets which get reordered	Under traffic load of 6 PC hosts	Under traffic load of 11 PC hosts	Under traffic load of 6 PC hosts	Under traffic load of 11 PC hosts
Both directions	30.739%	34.457%	12.078%	21.339%
For TX only	7.777%	0.0573%	0.00012%	0.00102%
For RV only	0.0036%	0%	0%	0.00008%

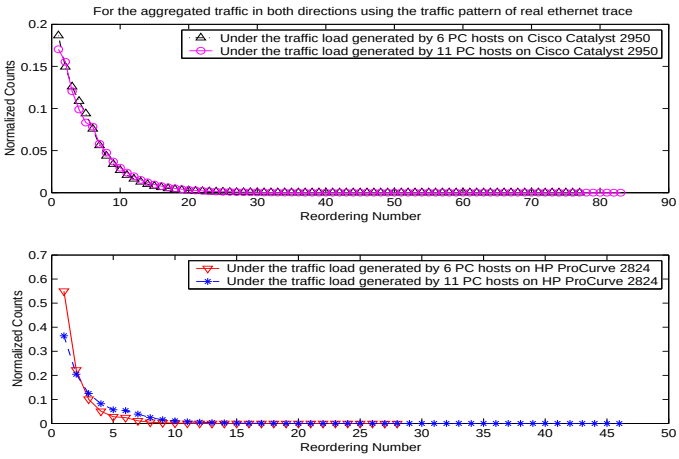

Fig. 16. Packet reordering number — real ethernet trace.

- When the traces are collected via port mirroring method, a significant percentage of packets gets reordered for the aggregated traffics in both directions even under low levels of utilisation.
- Packet-reordering might exist for the aggregated traffic in a single direction of the monitored link especially under higher traffic loads.
- The characteristics of reordering time and reordering number for the aggregated traffics in both directions of the monitored link are consistent under different traffic loads.
- Both the reordering time and reordering number span a range of values although one of them is dominant.

For this traffic pattern, there is no packet loss under the traffic load generated by 6 PC hosts; there are 0.0031% and 0.0077% packets on the monitored link get lost on the mirroring link under the traffic load generated by 11 PC hosts when port mirroring is made on Cisco Catalyst Switch 2950 and HP ProCurve 2824, respectively. When combining the packet-loss results in the above two subsections, we believe that the statistics of packet loss incurred due to monitoring via port mirroring is consistent under a certain traffic load no matter on which switch the port mirroring is made.

IV. Discussions

Based on the findings on the traffic trace artifacts in terms of timing difference, packet reordering and packet loss incurred by port mirroring, we discuss some potential implications of using the port-mirroring method for the future network measurement and monitoring. Firstly, as the timing difference between the original packet streams and the mirrored packet streams does exist and it could span a range of values under all traffic load levels, the inter-arrival time statistics of a packet trace would get biased if the trace is collected via the port mirroring method. Moreover, further analysis results could get biased if they depend on the accuracy of the inter-arrival time statistics. Thus, for network monitoring or analysis which needs highly accurate inter-arrival time statistics inferred from the collected trace data, more accurate methods, for example, the passive TAP technique, should be used.

Secondly, since quite a number of packet reorderings do happen in the packet traces collected via the port mirroring method, it will not only affect the inter-arrival time statistics, but also bias all the analysis results which are based on the packet arrival sequences, e.g., the analysis of TCP reset behavior based on TCP sequence number. Therefore, for this kind of analysis, we also recommend using the more accurate method to collect their data traces.

V. Conclusion

The port-mirroring technique is supported by most of today's switches and is a popular method for collecting packet traces for various purposes due to its wide availability and low cost. However, its side effects or impact on the collected traffic traces in terms of timing difference, packet reordering and packet losses, have been little known by the networking community. In this paper, we carried out the well-designed experiments to try to make a comprehensive study on the traffic trace artifacts incurred due to monitoring via port mirroring. The experiment results show that, when the traces are collected via port mirroring method, the traffic trace artifact in term of timing difference does exist and under different traffic load level using different traffic patterns on different switches, Furthermore, a significant percentage of packets get reordered for the aggregated traffic in both directions of a monitored link even under low levels of utilisation. We also document the statistics of the timing difference, packet reordering and packet losses observed in our experiments, which suggests that more-accurate methods should be used to collect the packet traces if the network monitoring and/or analysis needs to infer highly accurate inter-arrival time statistics or to rely on accurate packet arrival sequences.

Future work

We see this investigation as only the beginning of such an enquiry into the impact of port-mirroring. While not detailed here, we found inconsistent results for LACP link channels and consider this very imporant for a future investigation.

Future presentations of this topic would include a complete presentation on port-mirroring when the overall mechanism is applied to only part of the data stream: selections of host, protocol, and port, or single directions of flow.

Thanks: We thank the anonymous reviewers for their feedback. We also thank Matt Burnstein and Tavinda Jandu for their technical assistance and Ralphe Neill, Awais Awan and Wei Li for their assistance in reading early versions of this paper.

References

[1] M. Arlitt and C. Williamson, "An Analysis of TCP Reset Behaviour on the Internet," ACM Computer Communication Review, Vol.35, No.1 pp.37-44, Jan. 2005.
[2] M. Arlitt, B. Krishnamurthy and J. C. Mogul, "Predicting short-transfer latency from TCP arcana: Atrace-based validation," ACM/USENIX IMC'05, pp.213-226, Oct. 19-21, 2005, Berkeley, CA, USA.
[3] K.T. Chen, C.Y. Huang, P. Huang and C.L. Lei, "Quantifying Skype User Satisfaction," ACM Sigcomm 2006, pp.399-410, Sept. 12-14, Pisa, Italy.
[4] "Analyzing Full-Duplex Networks Through SPANs, Port Aggregators, and TAPs," White paper, Network Instruments (www.networkinstruments.com).
[5] J. Spooner and S. Donnelly, "Monitoring and Transport without compromise," White paper, Endace Technology Ltd.
[6] "Tcpfire," http://www.cl.cam.ac.uk/Research/SRG/