CHAPTER 6

Reproducing Traffic

Sometimes the appropriate response to reality is to go insane.

— Philip K. Dick (1928–1982), Valis

Dissertations are not finished; they are abandoned.

— Frederick P. Brooks, Jr. (1931–)

This chapter examines the statistical characteristics of source-level trace replay experiments, comparing them to those of their corresponding original traces. As discussed in Chapter 5, and illustrated in Figure 5.1, a packet header trace $\mathcal{T}_h$ and its source-level trace replay can be compared at two levels. The first level is how well the set of connection vectors $\mathcal{T}_c$ extracted from $\mathcal{T}_h$ are preserved by the trace replay experiments. This means to collect a packet header trace $\mathcal{T}_h'$ from the replay and extract a new set of connection vectors $\mathcal{T}_c'$. Section 5.2 presented a comparison of $\mathcal{T}_c$ and $\mathcal{T}_c'$ for three traces. It demonstrated that the characteristics of $\mathcal{T}_h$ captured by $\mathcal{T}_c$ are accurately reproduced by the traffic generation method and its implementation. The second level at which traces and their replays can be compared is to directly extract statistics from $\mathcal{T}_h$ and $\mathcal{T}_h'$. If these statistics are *reasonably* close, we can say that the traffic generation method *reproduces* the original traffic using closed-loop traffic generation. This is the type of comparison discussed in this chapter. As we will show, source-level trace replay generally results in a good approximation of the statistical characteristics of the original traffic, which supports the use of the a-b-t model as a foundation for realistic traffic generation.

6.1 Beyond Comparing Connection Vectors

The main goal of this dissertation is to improve the state-of-the-art in closed-loop traffic generation by developing a better approach to source-level modeling. In particular, we presented in Chapter 3 the

sequential and concurrent versions of the a-b-t model, which provide a first method for describing source-level behavior in an application-independent manner. We also discussed an efficient data acquisition algorithm for extracting a-b-t connection vectors from the packet headers of TCP connections. The first way in which we justified our source-level model was by examining connections from different applications, and demonstrating that their source-level descriptions in terms of a-b-t connection vectors properly captured their source-level behavior. The second way in which we can justify the model is to study the traffic generated using this model. If generated traffic is shown to closely approximate original traffic, this would strongly support the claim that the a-b-t model is a good description of source behavior. In other words, given that the statistical characteristics of $\mathcal{T}_h$ are obviously a function of source behavior, being able to generate a $\mathcal{T}_h'$ statistically similar to $\mathcal{T}_h$ would confirm the quality of $\mathcal{T}_c$ as a description of the original source behavior.

Comparing $\mathcal{T}_h$ and $\mathcal{T}_h'$ is however a subtle exercise. The actual replay of $\mathcal{T}_c$ necessarily requires choosing a set of network-level parameters, such as round-trip times and TCP receiver window sizes, for each TCP connection in the source-level trace replay experiment. The exact set of packets and their arrival times is a direct function of these parameters, as explained in Chapter 4. As a consequence, if we were to conduct a source-level trace replay using arbitrary network-level parameters, we would obtain a $\mathcal{T}_h'$ with little resemblance to the original $\mathcal{T}_h$. The replayed a-b-t connection vectors may be a perfect description of the source behavior driving the original connections, but the generated $\mathcal{T}_h'$ would still be very different from the original $\mathcal{T}_h$. To address this difficulty, the replay should incorporate network-level parameters individually derived from $\mathcal{T}_h$ for each connection. In Chapter 4, we described and evaluated methods for measuring three important network-level parameters: round-trip time, TCP receiver window size and loss rate. While this set of parameters is by no means complete, it does include the main parameters that affect the average throughput of a TCP connection, [PFTK98]. In this chapter, we examine the results of several source-level trace replay experiments, showing that the generated traffic is remarkably close to the original traffic. This is a strong justification of our source-level modeling approach, since it demonstrates that the closed-loop replay of a-b-t connection vectors provides a good approximation of the original traffic.

Incorporating network-level properties is important, but it is critical to understand the main shortcoming of this approach. The goal of our work is not to make the generated traffic $\mathcal{T}_h'$ identical to the original traffic $\mathcal{T}_h$, which could be accomplished with a simple packet-level replay. The goal is to develop a closed-loop traffic generation method based on a rich characterization of source behavior. Comparing $\mathcal{T}_h$ and $\mathcal{T}_h'$ is a means to understand the quality of traffic generation method, where quality is considered

to be higher as the original trace is more closely approximated. By construction, traffic generated using source-level trace replay can *never* be identical to the original traffic. The statistical properties of original packet header traces are the result of multiplexing a large number of connections into a single link, and these connections traverse a large number of different paths with a variety of network conditions. It is simply not possible to fully characterize this environment and reproduce it in a laboratory testbed or in a simulation. This is both because of the limitations of passive inference from packet headers, and because of the stochastic nature of network traffic. Source-level trace replay can never incorporate every factor that shaped $\mathcal{T}_h$, and therefore differences between $\mathcal{T}_h$ and $\mathcal{T}_h'$ are unavoidable. Still, finding a *close* match between an original trace and its replay, even if they are not identical, constitutes strong evidence in favor of our a-b-t model and our data acquisition and generation methods. It also demonstrates the feasibility of generating realistic network traffic in a closed-loop manner that resembles a rich traffic mix.

Besides evaluating source-level trace replay by comparing original traces and their replays, this chapter also considers whether *detailed* source-level modeling is necessary to achieve high-quality traffic generation. This is accomplished by comparing traffic generated using $\mathcal{T}_c$ (*i.e.*, replaying connection vectors and network-level parameters) and traffic generated using a simplified version of $\mathcal{T}_c$ with *collapsed epochs*, which we will name $\mathcal{T}_c^{coll}$. Formally, given a sequential connection vector $C_i = (e_1, e_2, \ldots, e_n)$, $n \geq 1$, with epoch tuples of the form $e_j = (a_j, ta_j, b_j, tb_j)$, we define the version of C_i with collapsed epochs as

$$C_i^{coll} = ((\sum_{i=1}^{n} a_i, 0, \sum_{i=1}^{n} b_i, 0)).$$

The only a-type ADU size in the resulting connection vector is the total amount of data sent from the connection initiator to the connection acceptor, and the only b-type ADU size is the total amount of data sent from the connection acceptor to the connection initiator. No quiet time is part of a connection vector after collapsing its epochs. Similarly, given a concurrent connection vector $C_k = (\alpha, \beta)$, where

$$\alpha = ((a_1, ta_1), (a_2, ta_2), \ldots, (a_{n_a}, ta_{n_a}))$$

and

$$\beta = ((b_1, tb_1), (b_2, tb_2), \ldots, (b_{n_b}, tb_{n_b})),$$

we define the version of C_k with collapsed epochs as

$$C_k^{coll} = ((\sum_{i=1}^{n_a} a_i, 0), (\sum_{i=1}^{n_b} b_i, 0)).$$

Traffic generated according to $\mathcal{T}_c^{coll}$ does not incorporate any internal source-level structure of connections, *i.e.*, epochs and inter-ADU quiet times are ignored. For this reason, we say that the collapsing of epochs "removes" detailed source-level modeling. Note however that even if epochs are collapsed, the total amount of data transferred in each direction does not change. The results in this chapter demonstrate that traffic generated using $\mathcal{T}_c$ is substantially closer to the original traffic than traffic generated using $\mathcal{T}_c^{coll}$.

The evaluation of source-level trace replay presented in this chapter examines the results of replaying five traces. These traces were first considered in Section 3.5: Leipzig-II, UNC 1 PM, UNC 1 AM, UNC 7:30 PM and Abilene-I. Our analysis compares the statistical characteristics of each of these traces and their replays using the following metrics:

- time series of byte throughput,
- time series of packet throughput,
- Body and tail of the marginal distribution of byte throughput,
- Body and tail of the marginal distribution of packet throughput,
- Wavelet spectrum (logscale diagram),
- Estimated Hurst parameter and its confidence interval, and
- time series of the number of active connections.

These metrics were introduced in Section 4.2. For each original trace, we compare four different replays, conducted using *tmix* and *usernet* in the testbed shown in Figure 5.2. The first replay is the *lossless replay*, which replayed the a-b-t connection vectors in $\mathcal{T}_c$, giving each TCP connection its measured round-trip time and TCP receiver window sizes. The second replay is the *lossy replay*, which was identical to the first one, but it also applied random packet dropping to each TCP connection according to its measured loss rate. The third replay, is the *lossless replay with collapsed epochs*, which replayed the a-b-t connection vectors after they had their epochs collapsed, and it also gave each connection its measured round-trip time and TCP received window sizes. The fourth replay is the *lossy replay with collapsed epochs*, which was identical to the third one but incorporated loss rates. We will often refer to the first two replays as *full replays* and to the second two replays as *collapsed-epochs replays*.

It is important to note that our method for incorporating losses into the experiments, random dropping according to the measured probability of loss per connection, is not consistent with closed-loop

traffic generation. We are by no means suggesting that loss rates should be incorporated in this manner into regular networking experiments that require closed-loop traffic generation. In such experiments, losses should only be the result of congestion on network links and buffering limitations. If this is the case, the endpoints generating synthetic traffic can not only react to the network conditions (*e.g.*, reducing sending rates when congestion is detected), but also modify them (*e.g.*, reducing overall congestion thanks to the lower sending rates). This is the right approach to reproduce the essential feedback loop in TCP which should be used in empirical studies of TCP performance.

However, loss is an important factor in TCP behavior (see Section 4.1.3), so our lossy experiments should result in a $\mathcal{T}_h'$ that is closer to the original $\mathcal{T}_h$. By incorporating losses, we eliminate one possible cause of divergence between original and replayed traces which could confuse our assessment of our source-level modeling approach. Comparing lossless and lossy replays enables a more systematic evaluation of our traffic modeling and generation methods, and it also helps to understand the impact of loss rates on the generated traffic. Losses are shown to have only a minimal effect on some traces and for some metrics, but a much more substantial effect on others.

The analysis in this section confirms the high-quality of the synthetic traffic generated using source-level trace replay. Our analysis reveals some (mostly minor) differences between original traffic and replay traffic. While we put forward some hypotheses about the cause of these differences, their confirmation requires further analysis. This additional work, which would involve both analysis and experimentation, would certainly be enlightening. It would tell us more about the limitations of our approach, and even about the inherent limitations of testbed experimentation. However, we have chosen not to pursue this avenue here. As discussed above, our goal is not to generate a $\mathcal{T}_h'$ equal to $\mathcal{T}_h$, but to convincingly demonstrate the benefits of our closed-loop traffic generation method. We believe this chapter achieves this goal, so we do not present any further analysis beyond the comparison of five traces and their four types of source-level replays using a rich set of metrics.

6.2 Source-level Replay of Leipzig-II

6.2.1 Time Series of Byte Throughput

The first trace we consider in this chapter is Leipzig-II. It has a duration of 2 hours and 45 minutes, and its average throughput is relatively low. We will first consider the traffic received by Leipzig's hosts

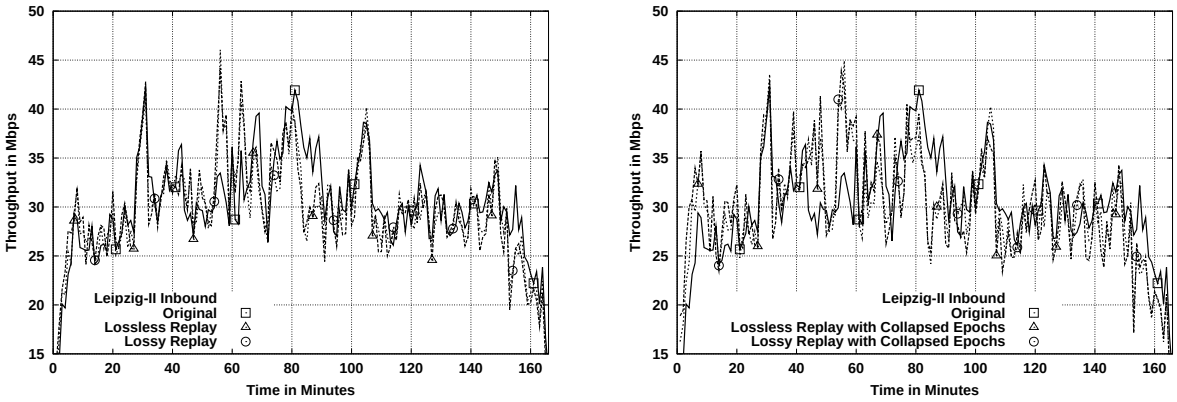

Figure 6.1: Byte throughput time series for Leipzig-II inbound and its four types of source-level trace replay.

from Internet hosts, *i.e.*, in the inbound direction with respect to the University of Leipzig. Figure 6.1 compares the original time series of byte throughput (solid line) and four different source-level trace replays (dashed lines). The plot on the left shows the full replays of $\mathcal{T}_c$ with and without imposing loss rates using *usernet* . The plot shows that the original time series is highly bursty[1], even when 1-minute bins are considered. Both replays closely approximate the original traffic, showing a strikingly good match in most regions. It also shows very little difference between lossless and lossy replays. This suggests that losses had a very moderate impact in the original trace, at least regarding the time series of byte throughput.

We also observe in the left plot of Figure 6.1 several major throughput spikes, *e.g.*, in minutes 25 and 105, that are very closely approximated by both replays. It is clear that the source-level nature of these spikes was accurately captured by our modeling approach. In a few other regions, the original and the replayed traces do not match so well. We have for example a spike in the throughput of the replays in minute 55 that was not present in the original traffic. This suggests that, for some number of connections active in that region of the trace, our model did not capture a significant limitation of throughput that was present in the original trace. This limitation could be at the source level or at the network level, but there is no way to know without further analysis. Given our traffic generation methods, we can however say that loss is very unlikely to be behind this difference, since both lossless and lossy replays show the same spike. We can also observe the opposite phenomenon in several locations, such as minutes 90 and 152, were we find ditches in the throughput of the replays. Here our measurement and modeling approach seems to be imposing an artificial limitation to the throughput of one or more connections.

[1]The term *bursty* does not have a unique meaning. In this paragraph, it simply refers to high variability. Some authors consider traffic more bursty as its long-range dependence becomes stronger [WP98], while others as its marginal distribution becomes less Gaussian [SRB01]. We make use of these more formal definitions, discussed in Chapter 4, in later sections.

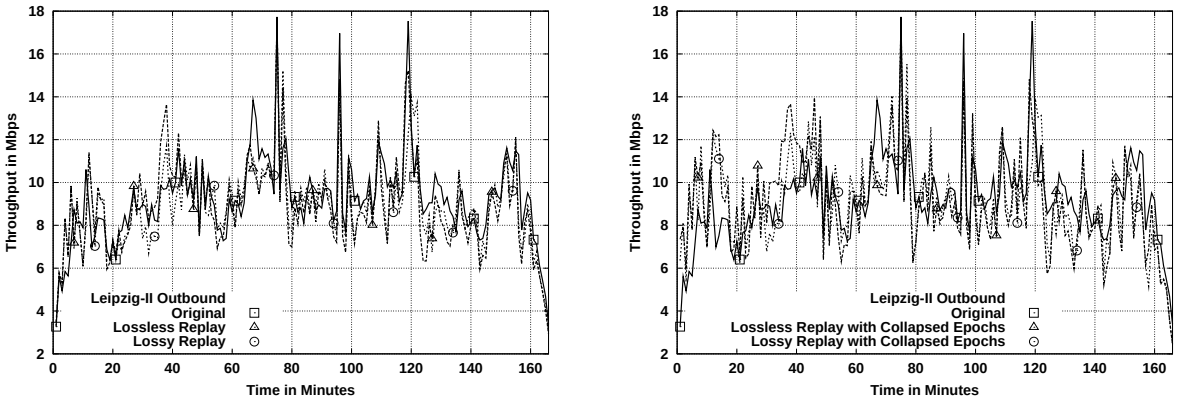

Figure 6.2: Byte throughput time series for Leipzig-II outbound and its four types of source-level trace replay.

While this suggests that further refinement is possible, the plot clearly shows that our approach result in an excellent approximation of the original byte arrival process and its overall burstiness.

The right plot of Figure 6.1 compares the original time series of byte throughput and the ones from the lossless and lossy replays with collapsed epochs. The approximation is also generally good, but the replays appear more bursty, which seems rather significant given the high level of aggregation (1-minute bins). The replays with collapsed epochs results in several new spikes in which the replay is well above the original throughput. This means that removing the source-level structure enabled artificially higher throughputs for some number of replayed connections. Despite these difficulties, it is important to note that the collapsed-epochs replay achieves a reasonably good approximation of original throughput with a much simpler source-level model. The collapsed-epochs replays could then be sufficient for some kinds of experimental studies in which only a good reproduction of the time series of byte throughput is required.

The time series of byte throughput in the outbound direction is studied in Figure 6.2. The comparison of the original and the full replays is found in the left plot. As we observed for the opposite direction, the time series from the replays closely track the original one, and losses do not have a significant impact. We find a number of sharp spikes and ditches from the original traffic that are well reproduced by the replays, *e.g.*, minutes 88, 97 and 143. We also find some artificial ones not present in the original, notably the spike in the replay on minute 38 and the ditch around minute 70. The right plot compares the original and the collapsed-epochs replays, which are again shown to be somewhat more bursty that the full replays throughout the trace.

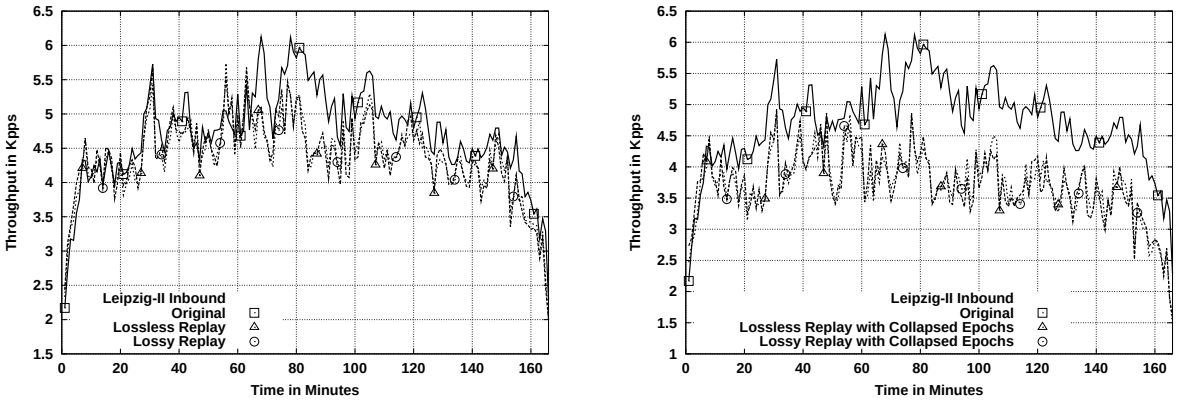

Figure 6.3: Packet throughput time series for Leipzig-II inbound and its four types of source-level trace replay.

6.2.2 Time Series of Packet Throughput

The analysis of the time series of packet throughput reveals larger differences between original and replayed traffic. Figure 6.3 shows the time series for the inbound direction. The comparison of the time series from the original trace and those from the full replays reveals a close approximation for the first 60 minutes, and a consistently lower packet throughput for the rest of the trace. The replays generally have between 2% and 5% less packets per 1-minute bin that the original trace, although they mostly track the original shape. The right plot shows that the collapsed-epochs replays result in far lower packet throughput for the entire trace, between 20% and 40% below the original. This clearly shows that the detailed modeling of source-level structure accomplishes a more realistic traffic generation in terms of the number of generated packets. The main reason is the modeling of epochs, which often increases the number of segments per connection. Replaying an epoch with non-zero ADU sizes necessarily involves sending two packets, even if the sizes of the ADUs are very small. An epoch involves a necessary exchange of data, so at least one packet is used to carry the ADU a_i from the initiator to the acceptor, and another one to carry the ADU b_i from the acceptor the initiator. This means for example that a connection with 10 epochs, and ADUs with a size of 100 bytes in both directions requires 20 packets to be fully replayed. On the contrary, the collapsed-epochs version of this connection can be replayed with a single pair of packets, since the 10 ADUs in each direction can fit into a single TCP segment (it is only 1,000 bytes). Another reason for the more realistic time series of packet throughput when the full replay is used is the modeling of quiet times. Quiet times between two ADUs sent in the same direction (see Section 3.1.2) can also result in a larger number of packets per connection, since they often prevent consecutive small ADUs from sharing packets.

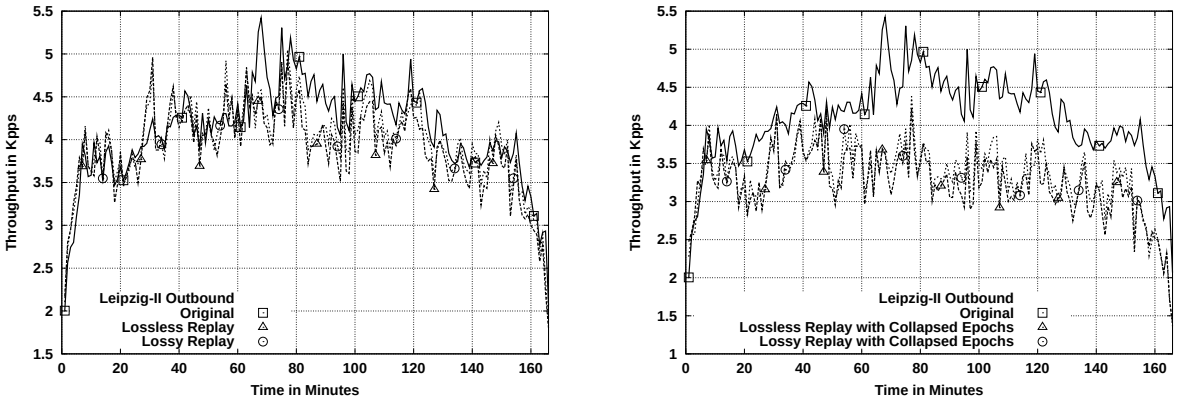

Figure 6.4: Packet throughput time series for Leipzig-II outbound and its four types of source-level trace replay.

While the results in Figure 6.3 convincingly demonstrate a substantially more realistic traffic generation with the full model, there is still some room for improvement. We can think of several possible refinements, which should improve the approximation. First, we made no attempt to model the Maximum Segment Size (MSS) supported by the path of each TCP connection. Instead of relying on the default size derived from Ethernet's MTU (1,500 bytes), as we do in our experiments, it seems possible to collect MSS information for each connection and extend *tmix* to make use of these measurements[2]. Connections replayed using smaller MSS values would frequently require more packets to be replayed. Second, the measurement techniques we used to determine ADU boundaries for data sent in the same direction rely on a constant inter-ADU quiet time threshold equal to 500 milliseconds. Some applications may be using smaller quiet times between their writes, which could result in a larger number of packets per connection. Simply reducing the threshold is problematic, since this would increase the number of spurious splits of ADUs due to network delays (rather than application behavior). To avoid this, we could make the inter-ADU quiet time threshold a multiple of the measured round-trip time. Given the typical distributions of round-trip times (see Section 4.1.1), this method would reduce the threshold for most connections and increase the sensitivity of the measurements. Another approach is to study segment sizes, using non-full segments to mark ADU boundaries. This would require some further refinement, since non-full segments can easily come from application writes which are not a multiple of the MSS. Two consecutive non-full segments are for example far more likely to mark a true ADU boundary.

The lesson is similar for the outbound direction results, which are shown in Figure 6.4. The left plot

[2]MSS is a system-wide constant in FreeBSD, so generating traffic that preserves per-connection MSS is not directly possible with our current implementation. However, there is a relatively simple way to extend our method to support per-connection MSS values. We could use a first step to group connections with the same MSS and then assign each group to a host configured with that MSS. Fortunately, only a few MSS values are common on the Internet, so it seems feasible to implement this extension without increasing the number of hosts.

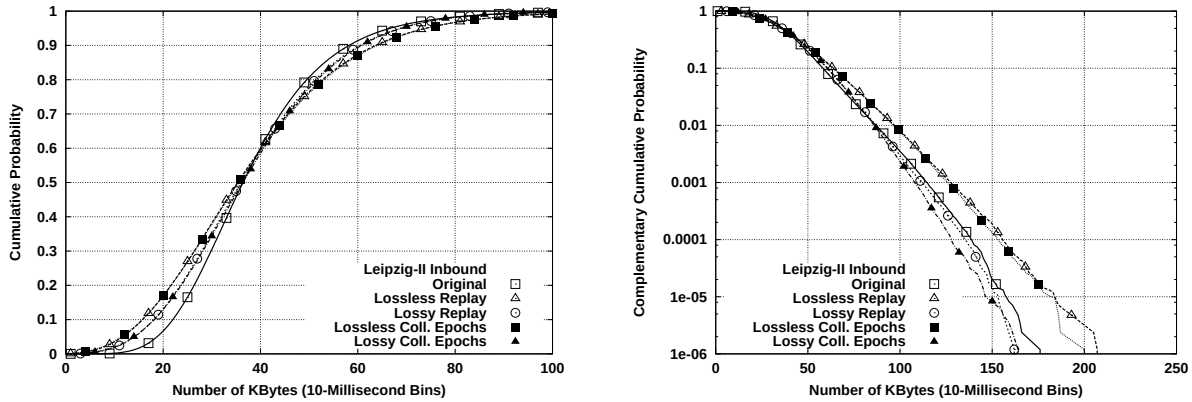

Figure 6.5: Byte throughput marginals for Leipzig-II inbound and its four types of source-level trace replay.

shows that the full replays are generally a good approximation to the original, but they exhibit a somewhat lower number of packets in some regions. On the contrary, collapsed-epochs replays consistently show a far lower number of packets.

The reader may be puzzled by the finding of very similar shapes for the inbound and outbound time series of packet throughput, which show spikes and ditches located at the same minutes. This is due TCP's acknowledgment mechanism, which forces TCP endpoints to at least send one acknowledgment for each pair of data segments received. As consequence, a connection that sends a large number of data segments in one direction, creating a spike in the time series, must necessarily receive a large number of acknowledgments in the opposite direction, creating a similar spike.

6.2.3 Marginal Distributions

One important limitation of the type of analysis in the previous section is the use of a relatively coarse level of aggregation (1-minute bins). The obvious question is whether the close match between original traffic and its source-level replays is also found at finer scales, which are arguably more important for some kinds of studies, such as router queuing evaluation. Given the highly bursty nature of the throughput time series, simply plotting the time series at finer levels of aggregation just makes the plots completely unreadable. In this section, we rely on a different kind of analysis to examine the difference between original and replayed traffic at a finer level of aggregation. Instead of the 1-minute bins used in the previous section, this section examines throughput using CDFs of the marginal distributions extracted from time series of 10-milliseconds bins. Section 4.2.2 further discusses the reasoning behind this type of analysis.

Figure 6.5 plots the marginal distributions of the byte throughput in the inbound direction, showing the data for the original time series and the four types of replay. The left plot shows the body of the marginal distributions using CDFs in linear axes. The right plot shows the tail of the marginal distributions using CCDFs in a logarithmic y-axis. The plot of the tail provides information about the 10-millisecond bins with the highest throughput, giving us a better sense of how well the most "aggressive" regions (*i.e.*, with the highest throughput) of the time series are reproduced by the replays. The vast majority of the plot comes from throughputs that are relatively uncommon, *e.g.*, half of the plot shows data from only 0.1% of the distribution. On the contrary, the plot of the body provides information about the most common bins, showing the entire distribution without focusing on any particular region. These two visualizations are complementary. The body plot shows the overall match, which is relevant for experiments in which producing a realistic range of fine-scale throughputs is important. The tail plots shows the extremal match, which is relevant for experiments in which reproducing the magnitude and frequency of peak throughputs is important. None of these plots says anything about the dependency structure of the time series, which is important and that we study in a later section using wavelets. While wavelets are a powerful analysis tool, marginals are far easier to interpret in networking terms.

The left plot shows the original data using a solid curve marked with white squares, and the replay data using dashed curves. The full replay experiments are marked with white symbols, and the collapsed-epochs replay experiments with black symbols. We can make several observations about this plot. The position of the original curve with respect to the replay curves defines two different regions in the plots. Below 40 KB, the distribution from the original data is slightly heavier than those from the replays. Above 40 KB, the distribution is slightly lighter. This means that the replays tended to be less concentrated around the central value than the original data, For example, the number of bins with 10 KB is negligible in the original data, but corresponds to between 2% and 5% of the bins in the replays. We could therefore say that the replays are somewhat more bursty, in the sense that we find more bins with small values and more bins with large values in the CDFs from the replays than in the CDFs from the original data. The exact reason is unclear, but we can make a hypothesis. We know from the previous section that the total number of bytes is similar in original and replay time series. This means that the presence of a larger number of bins with more bytes in the replay must necessarily be accompanied by a larger number of bins with fewer bytes to compensate. Connections in the replay are exposed to more homogeneous delays (primarily because round-trip times are fixed), which gives replayed connections a chance to achieve higher throughput. In the aggregate, and when considering such fine scales, the presence of one or a few replay connections with higher throughput than originally observed creates bins

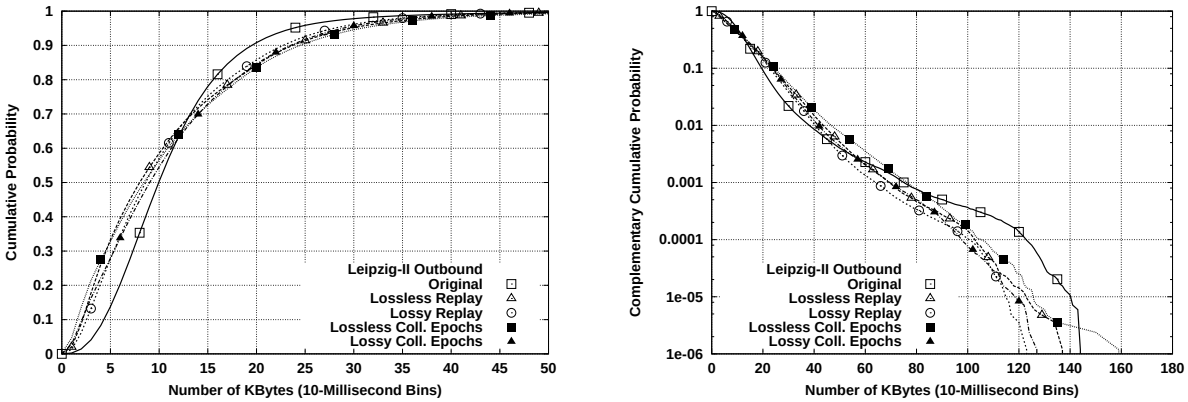

Figure 6.6: Byte throughput marginals for Leipzig-II outbound and its four types of source-level trace replay.

with more bytes, which are part of the upper portion of the body of the marginal distribution. Faster connections run out of data sooner, in turn creating bins with fewer bytes than originally observed, which show up in the lower portion of the body of the marginal distribution. Therefore, the somewhat milder conditions in the replay can explain the wider spread of marginal distributions from the source-level trace replay experiments.

Another observation from the plot of the bodies is that the collapsing of the epochs of the replayed connection vectors has no effect on the marginal distribution of byte throughput. This is an interesting finding, given that we did find a difference for the plots in Figure 6.1. It means that the slightly more bursty replays with collapsed epochs come from a less realistic correlation structure rather than from a fine-grain difference in the values of the bins. The plot also shows that the distributions from the lossy replays are slightly closer to the original than those from the lossless ones. This is evidence in support of the statement in the previous paragraph regarding the impact of more complex network dynamics, which make the highest throughput of many connections lower in the original trace. Adding losses has precisely this effect, making the marginal distributions from the replays closer to the marginal distribution from the original.

The analysis of tails in the right plot confirms the last observation. The plot of the body shows a lighter second half of the distribution. The plot of the tails shows heavier tails from the lossless experiments, and slightly lighter tails from the lossy experiments. The tail from the lossy full replay is actually an excellent fit of the original data. Lossless replays gave some connections the opportunity to reach higher throughputs, which in turn created bins with a larger number of bytes than in the original. Adding losses avoided this problem. In general, the results in Figure 6.5 are very reassuring.

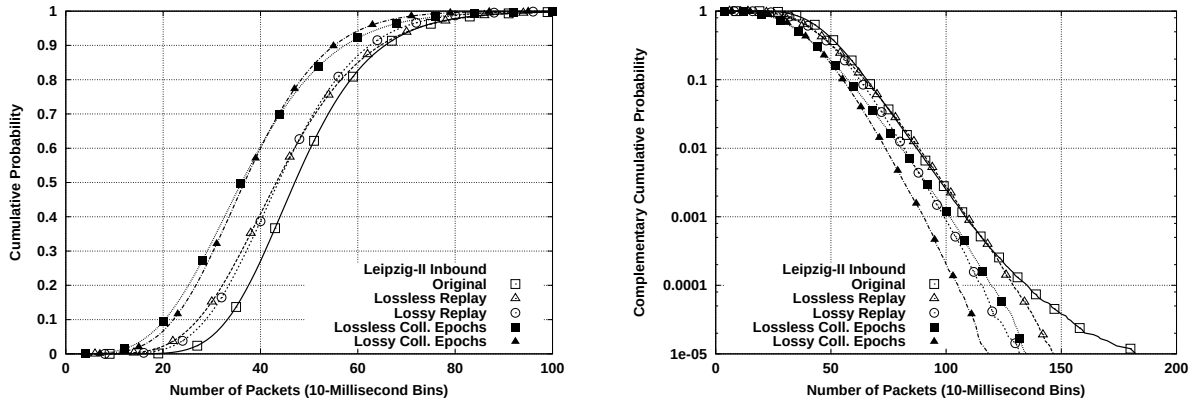

Figure 6.7: Packet throughput marginals for Leipzig-II inbound and its four types of source-level trace replay.

The marginal distributions for the time series of byte throughput in the outbound direction are shown in Figure 6.6. The bodies of distributions (left plot) exhibit a substantial tail, which makes them less Gaussian than distributions from the inbound data. As in the previous case, the range of bin sizes with a significant number of samples is wider for the replays than for the original. The relative difference seems slightly larger in this case, although the absolute difference is of the same magnitude. Lossy replays are again slightly closer to the original.

The tails of the marginal distributions shown in the right plot are not as close to a straight line as those found for the inbound direction. The shape of the tail is most complex for the original data, especially in the region above 90 KB. All of the replays achieve a good match below 90 KB, but are substantially lighter than the original above that value. The reason is unclear. The different shape can easily be due to the characteristics of a few connections (given the very small probabilities considered). The four replays result in similar tails.

As in the previous section, we follow our analysis of byte throughput with an analysis of packet throughput. The marginal distributions for the inbound direction are shown in Figure 6.7. The comparison of the bodies reveals a quite different result for packet throughput. In general, the distributions from the replays are significantly lighter than the distribution from the original. The difference is far larger for the collapsed-epochs replays. The reason was already discussed in the previous section. Collapsing epochs can often reduce the number of segments in a connection, since it enables connections to combine small ADUs from different epochs into a single ADU, increasing packet utilization. Our full replay, while much closer than the collapsed-epochs replay, is still lighter than the original. The possible extensions described in the previous section could improve the match further. Note also that the improvement when losses are used is quite minor, so retransmissions are not likely to explain the difference between

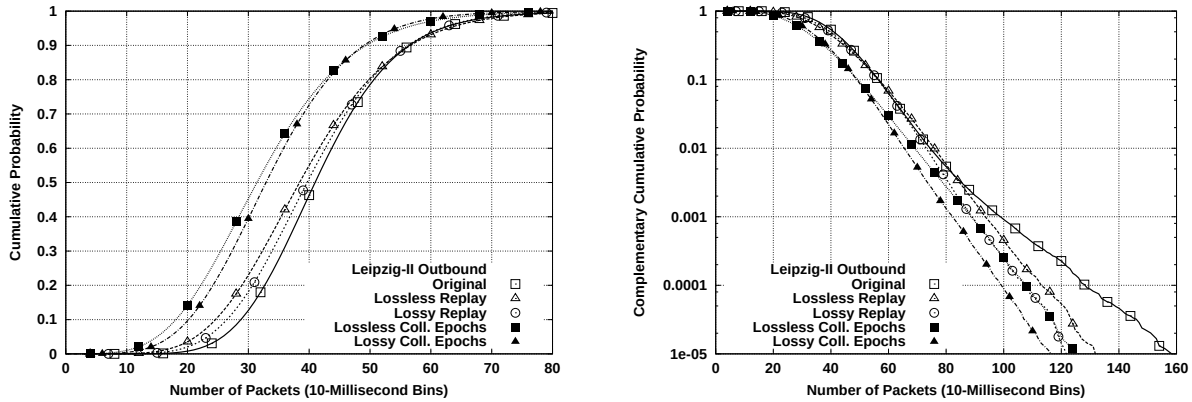

Figure 6.8: Packet throughput marginals for Leipzig-II outbound and its four types of source-level trace replay.

original and replay distributions.

The tails of the marginal distributions from the replays are lighter than those from the original data. Interestingly, the best match is achieved by the lossless replay with fully characterized epochs rather than by the lossy replay. The match is excellent below 10^{-4}. Above this value, the shape of the tail from the original data is less linear, which could be caused by a small number of connections with characteristics that we do not model well. Lossy replays result in significantly lighter tails, as expected given the loss-induced reduction in connection throughput.

Figure 6.8 shows the same analysis for the outbound direction. Collapsed-epochs replays again resulted in bodies that are substantially lighter than the body of the original distribution. In contrast, the full replay achieved a much closer approximation, even overlapping the original distribution for the largest values. Adding losses to the experiments made the replays only a bit closer to the original. This is a strong indication that source-level structure, and not loss/retransmission, is behind the differences between original and replay trace. We can distinguish two regions in the plot of the tails. Below 80 Kpps, the replays with fully characterized epochs provide an excellent match, while those with collapsed epochs result in significantly lighter tails. Above 80 Kpps, the slope of the tail from the original trace is far higher than the slopes of the tails from the replays.

6.2.4 Long-Range Dependence

Another way of looking at the time series of byte and packet arrivals is to study the characteristics of the time series for a wide range of time scales. This can be accomplished using scaling analysis tools,

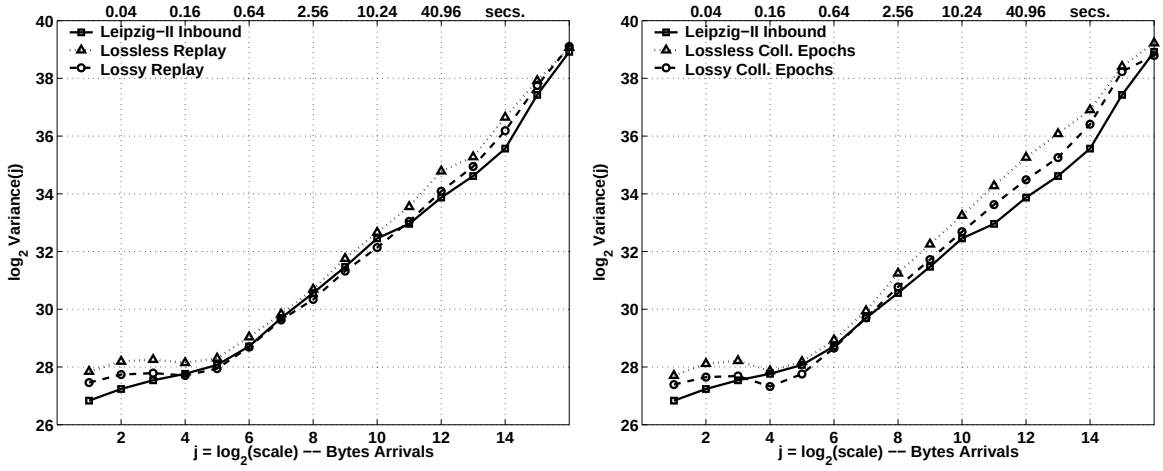

Figure 6.9: Wavelet spectra of the byte throughput time series for Leipzig-II inbound and its four types of source-level trace replay.

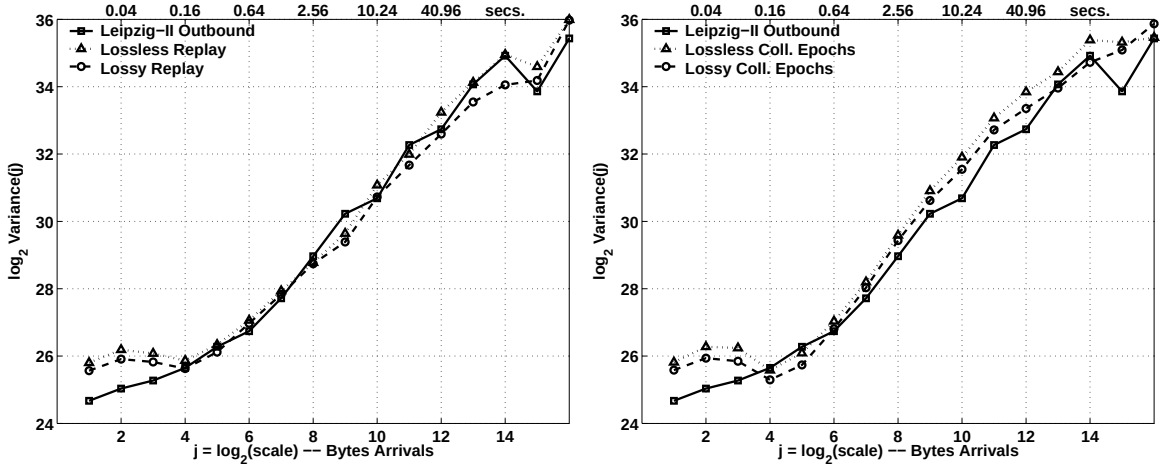

Figure 6.10: Wavelet spectra of the byte throughput time series for Leipzig-II outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
Leipzig-II	H	C. I.	H	C. I.
Original	0.9201	[0.8990, 0.9412]	0.9973	[0.9762, 1.0184]
Lossless Replay	0.9863	[0.9652, 1.0074]	1.0475	[1.0264, 1.0686]
Lossy Replay	0.9583	[0.9372, 0.9794]	0.9832	[0.9621, 1.0043]
Lossless Coll. Epochs	0.9986	[0.9775, 1.0197]	1.0473	[1.0262, 1.0684]
Lossy Coll. Epochs	0.9668	[0.9457, 0.9879]	1.0083	[0.9872, 1.0294]

Table 6.1: Estimated Hurst parameters and their confidence intervals for the byte throughput time series of Leipzig-II and its four types of source-level trace replay.

such as the wavelet transform, which was introduced in Section 4.2.3. In this section, we use wavelet spectrum plots and Hurst parameters estimates to compare the scaling of the arrival processes found in original and replay traces. Figure 6.9 shows the wavelet spectra of the time series of byte arrivals in the inbound direction. The left plot reveals an excellent match between the original and the full replays. The linear region between octaves 6 and 14 is very similar in the three spectra. This tells us that the kind of long-range dependence found in the original and in the replay traces is very similar. If we equate burstiness to long-range dependence, we can say that the generated traffic faithfully reproduced the burstiness of the original traffic. The finest time scales show a somewhat larger difference between octaves 1 and 5. The spectrum of the original data starts at a lower energy level than the spectra of the replay data. It also shows a linear trend with an upward slope, which is far less clear in the replay data.

The exact cause of the small difference is not completely clear. Our additional experiments strongly suggest that it is due to more complex network-level characteristics in the Internet than in the network testbed. We conducted a large set of experiments (not reported here) which betrayed that the energy levels at the finest time scales are dominated by round-trip times and other network-level parameters[3]. The slightly better match achieved with the lossy replay is consistent with this claim. Further work on network-level modeling may help improve the match, but it is beyond the scope of this dissertation. The approximation seems acceptable for most experimental studies.

The wavelet spectra of the collapsed-epochs replays is similar to the wavelet spectrum of the original trace, as shown in the right plot of Figure 6.9. The spectra from the replays exhibits a slightly higher slope in the linear region, and a slightly worse approximation of the fine-scale region. The benefit of modeling source-level behavior is relatively small, in terms of scaling behavior, for this trace, but present nonetheless.

Figure 6.10 shows the analysis of the wavelet spectra of the time series of byte throughput in the outbound direction. One interesting observation is that the wavelet spectrum of the original is far from the expected straight line. This is due to the low mean throughput on this direction. A handful of connections can have a large impact in the aggregate throughput, which makes the aggregate less stable, showing a less clear scaling. The full replays are very close to the original in the scaling region, but show a larger gap at fine scales. The collapsed-epochs replays result in a slightly worse approximation.

Estimated Hurst parameters for the byte throughput time series are shown in Table 6.1. The original

[3]More specifically, we learned that the range of the distribution of round-trip times determines the knee of the spectrum, while the distribution of window size determines the level of energy at the finest scales. Related results from web traffic simulations can be found in [FGHW99].

trace exhibits a smaller estimated Hurst parameter than the replays. The estimate for the lossy replay is however within the confidence interval of the original for the outbound and very close to the upper bound for the inbound. In general, lossless replays have higher Hurst parameters than lossy replays, and the replays with collapsed epochs have somewhat higher Hurst parameters than the full replays. Note also that several estimated Hurst parameters for the outbound direction are above 1, with the lossless replay even having the lower bound of the confidence interval above 1. Non-stationarities, properly captured by the source-level trace replay, may be behind this extreme burstiness. It is important to note that non-stationarity, even if present, does not change the fact that our computation of wavelet energy and Hurst estimates is identical in all cases. This makes the comparative results meaningful, at least in relative terms.

Figure 6.11 shows the wavelet spectra for the time series of packet throughput in the inbound direction. As in the case of byte throughput, the spectra of the replays are quite similar to the spectrum of the original, especially in the linear region. The spectra of the collapsed-epochs replays are somewhat farther from the original spectrum than the ones from the full replays. The slope of the linear region is again higher for the collapsed-epochs replays, and the difference is also larger at the finest scales.

The analysis of the packet throughput in the output direction shown in Figure 6.12 reveals a close approximation of the original spectrum by the full replays. Collapsed-epochs replays are slightly worse. Note also that the spectrum of the original trace is smoother here than in Figure 6.10. The phenomenon that distorted the linear scaling in the original time series of byte throughput seems far less significant for the time series of packet throughput.

Table 6.2 presents the estimates of Hurst parameters and confidence intervals for the original and replay time series of packet throughput. The original and the lossy full replays have almost identical estimated Hurst parameters for the inbound direction, while the other replays show higher Hurst parameters. The estimated Hurst parameter of the lossy full replay is again the closest one to the original estimate for the outbound direction. It is somewhat lower than the original, but within the confidence interval. The other replays show significantly higher estimated Hurst parameters. Note also that the estimated Hurst parameters for the outbound direction do not go above 1 in this case.

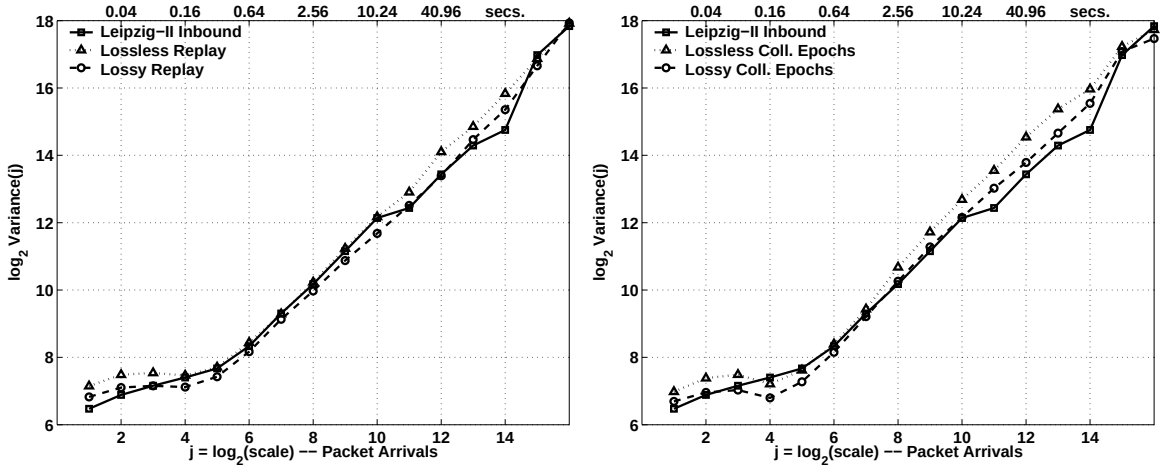

Figure 6.11: Wavelet spectra of the packet throughput time series for Leipzig-II inbound and its four types of source-level trace replay.

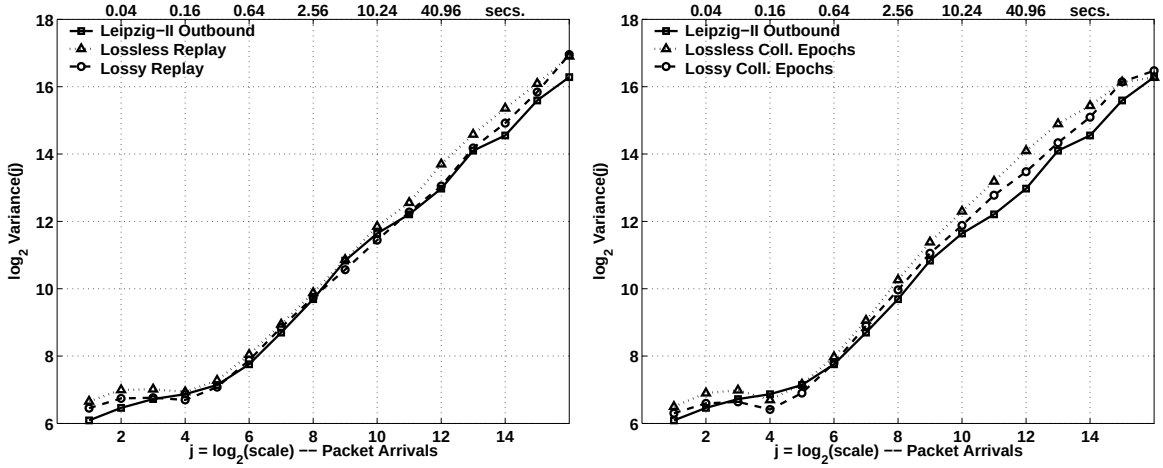

Figure 6.12: Wavelet spectra of the packet throughput time series for Leipzig-II outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
Leipzig-II	H	C. I.	H	C. I.
Original	0.9208	[0.8975, 0.9442]	0.9399	[0.9165, 0.9633]
Lossless Replay	0.9716	[0.9482, 0.9950]	0.9701	[0.9468, 0.9935]
Lossy Replay	0.9271	[0.9038, 0.9505]	0.9194	[0.8961, 0.9428]
Lossless Coll. Epochs	0.9883	[0.9649, 1.0116]	0.9925	[0.9692, 1.0159]
Lossy Coll. Epochs	0.9587	[0.9353, 0.9820]	0.9635	[0.9402, 0.9869]

Table 6.2: Estimated Hurst parameters and their confidence intervals for the packet throughput time series of Leipzig-II and its four types of source-level trace replay.

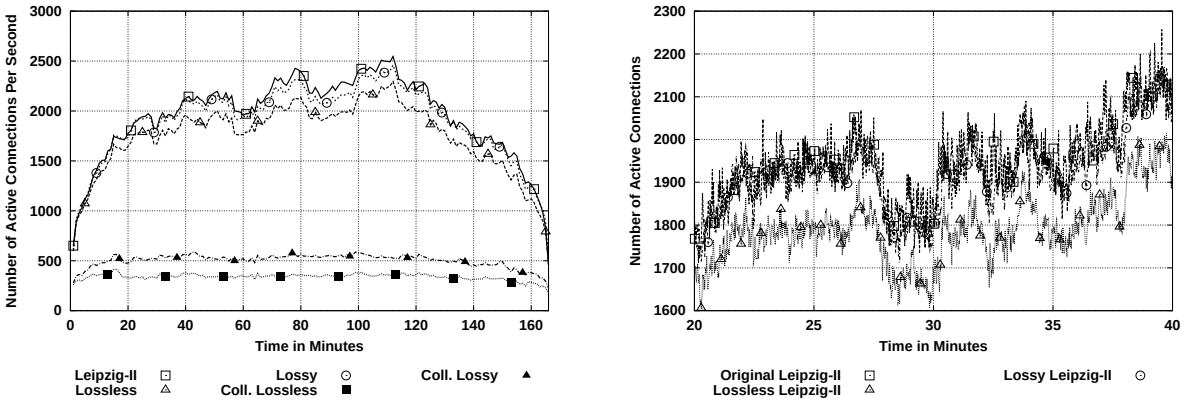

Figure 6.13: Active connection time series for Leipzig-II and its four types of source-level trace replay.

6.2.5 Time Series of Active Connections

The final metric we examine in this chapter to evaluate how closely original and generated traffic match is the time series of active connections. The left plot in Figure 6.13 shows the time series from the original trace using a solid line, and the time series from the four replays using dashed lines. The first observation from this plot is that the collapsed-epochs replays resulted in a strikingly lower number of active connections that the full replays. Since the number of connections replayed in both types of the replay is the same, this difference is due to the substantially shorter durations of the connections replayed with their epochs collapsed. The collapsing of epochs increases connection durations, because quiet times and epoch structure disappear. Epochs require at least one round-trip time to be replayed (see Section 3.1.1). As a result, the number of active connections is several times smaller in the collapsed epochs replays than in the original trace. On the contrary, the number of active connections observed in the full replays is far closer to the original.

The left plot of Figure 6.13 also provides a good illustration of the impact of replaying losses on the quality of the approximation. The number of active connections increases substantially when loss rates are used in the generation, both in the case of collapsed-epochs replays and full replays. However, it is clear from this plot that collapsing epochs has a far more substantial impact on the number of active connections than incorporating losses, at least for the Leipzig-II trace. Given how carefully our replay reproduced the main network-level parameters that affect TCP throughput (round-trip time, window size and loss rates), this result strongly suggest that traffic generated without any modeling of epoch structure and quiet time has an unrealistically low number of active connections.

While the lossless full replay achieves a reasonable approximation of the original time series, the

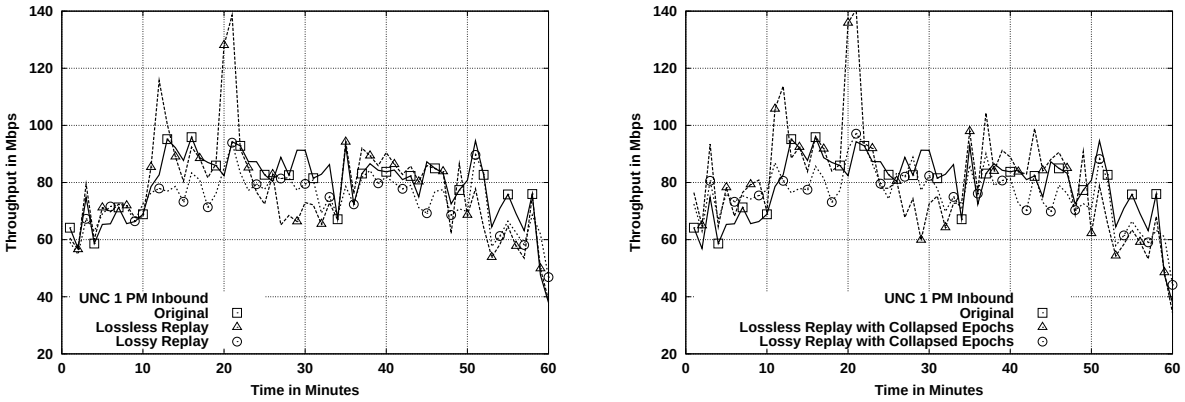

Figure 6.14: Byte throughput time series for UNC 1 PM inbound and its four types of source-level trace replay.

lossy full replay is almost a perfect match. The difference is always below 100 connections, which can be considered an outstanding result. It is clear that generating traffic using a combination of detailed source-level models and primary network-level parameters makes the number of active connections very realistic. Note also that this is not only true for the coarse scale (1 minute) at which the left plot of Figure 6.65 is displayed, but also at the finer scale (5 seconds) in the right plot. Notice for example how closely the replay tracks the significant variability in the original time series.

6.3 Source-level Replay of UNC 1 PM

6.3.1 Time Series of Byte Throughput

Figure 6.14 shows the time series of byte throughput for UNC 1 PM in the inbound direction, revealing a good match between original and replayed traces. Lossless replays with and without collapsed epochs are generally closer than lossy replays, which are often 10 to 20 Mbps below the original. However, lossless replays show large spikes (minutes 14 and 21) that are not found neither in the original trace nor in the lossy replays. The lossy replays are actually very close to the original in the neighborhood of these spikes (*e.g.*, between minutes 20 and 28). Interestingly, the time series for Leipzig-II shown in Figure 6.14 did not reveal a significant difference between lossless and lossy replays. Finding an explanation for this phenomenon requires further analysis, but this plot certainly justifies our choice of comparing the original trace to lossless and lossy versions of its source-level trace replay. Without a lossy replay, we would be tempted to conclude from the artificial throughput spikes in lossless replay that our source-level model is not properly reproducing an end-point limitation that was present in the original environment.

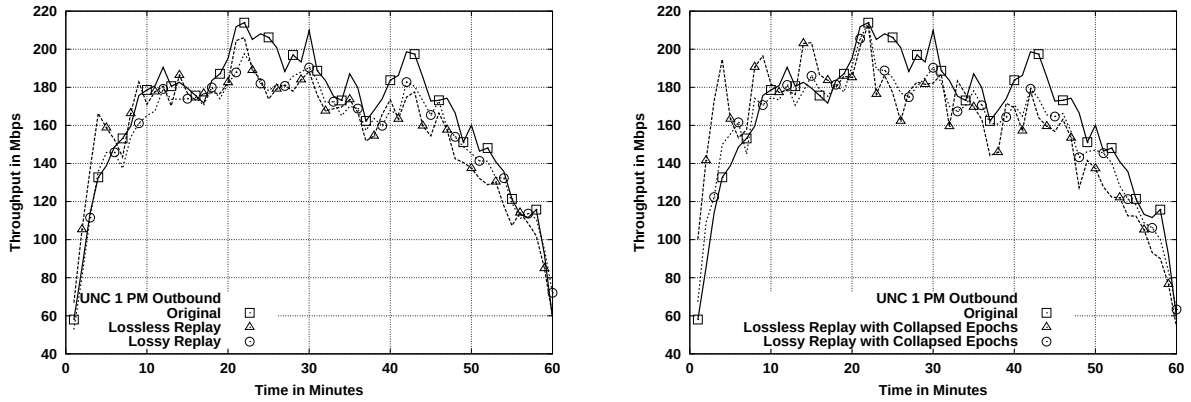

Figure 6.15: Byte throughput time series for UNC 1 PM outbound and its four types of source-level trace replay.

However, the lossy replay, by showing that adding losses eliminates this spikes, demonstrates that they are purely due a network-level parameter and not to a limitation of the a-b-t model. Once again, we are not naively advocating for incorporating open-loop losses into traffic generation experiments, but addressing a difficulty that significant loss can create when trying to understand how realistic our modeling of the traffic source is. Simply relying on a lossless replay can be misleading, as this example demonstrates.

As in the full replay case, the lossless collapsed-epochs replay shows two large spikes that are not present in the lossy collapsed-epochs replays. The general impression from the plot is that collapsing epochs moderately increases the burstiness of the replay. Note for example the larger spike in the minute 5, the spikes in minutes 36 and 44, and the large ditch in minute 29. The collapsed-epochs lossy replay is quite similar to the full lossy replays, but we find a few periods where the approximation of the original throughput is slightly worse. For example, the collapsed-epochs replay shows a drop of byte throughput in minute 40 that is not present in the full lossy replay.

Figure 6.15 reveals somewhat different lessons from the time series of byte throughput in the outbound direction of UNC 1 PM. Regarding the full replays shown in the left plot, we see that the lossless replay has only one significant spike above the original traffic. One reason behind this finding is that the much higher average byte throughput makes spikes due to a few connections far less significant in relative terms.

Both full replays are generally slightly below the byte throughput of the original trace. The reason is not completely clear, but it suggests that the replay has a somewhat lighter distribution of connection throughputs, which makes the aggregate throughput slightly lower. If the replay is continued beyond

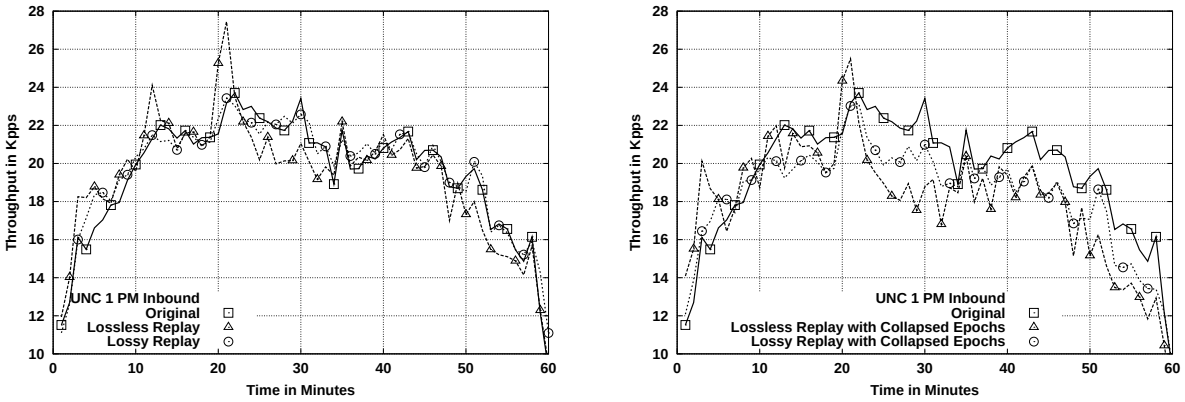

Figure 6.16: Packet throughput time series for UNC 1 PM inbound and its four types of source-level trace replay.

minute 60, we do observe connections that remain active for a few more minutes and transfer enough data to account for the difference between the time series. We examined the logs from the generator hosts and confirmed that no overload occurred during the experiments, so the cause seems to be some artificial limit on the throughputs of the connections in our replay. One cause could be the overestimation of quiet times discussed in Section 5.2.1. Another possible cause is that the replays did not take into account the specific MSS of each connection. Every connection was given the FreeBSD default value (1,460 bytes), which is the most common one on the Internet. However, it could be the case that a significant fraction of the segments were carried in TCP connections with a smaller MSS. These connections would then have higher control overhead, making their transferring of the same payload result in more bytes and therefore higher aggregate throughput. Given the small size of TCP headers, it is unlikely that the extra overhead would result in more than a few additional Mbps.

The results from the replays with collapsed epochs are similar, although we observe several additional spikes in the case of the lossless replay. The lossy replay does not show these spikes, but it is still below the original for most of the time series. Interestingly, it provides a closer approximation in some regions, such as between minutes 10 to 22. We can argue that this is an accidental improvement due to the artificially larger throughputs that a fraction of the connections achieves after their epochs are collapsed.

6.3.2 Time Series of Packet Throughput

The analysis of the packet throughput in the inbound direction shown in Figure 6.16 reveals a number of interesting characteristics. Both lossless replays show substantial spikes above the original packet throughput. This is consistent with the similar finding for byte throughput. We also observe that

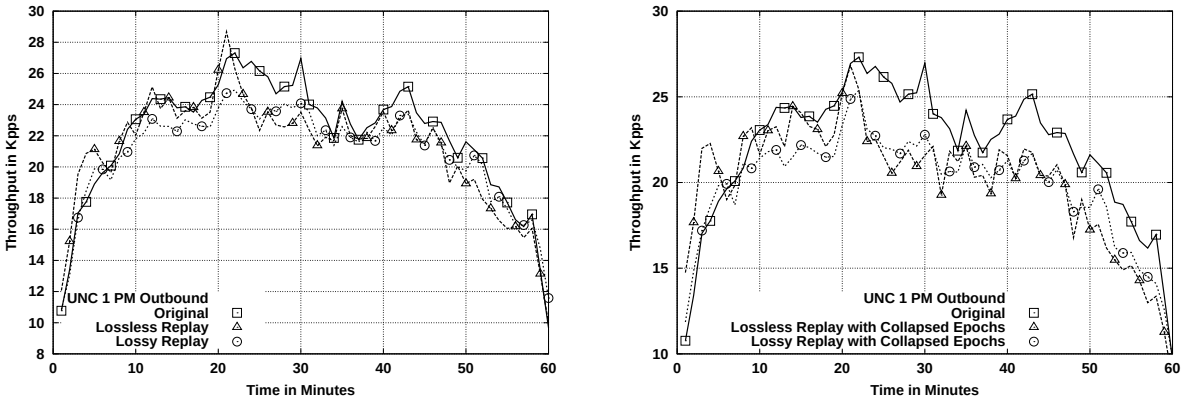

Figure 6.17: Packet throughput time series for UNC 1 PM outbound and its four types of source-level trace replay.

collapsed-epochs replays generated a substantially smaller number of segments than full replays. As in the case of the analysis of the Leipzig-II replay shown in Figure 6.3, the lack of detailed source-level modeling in the collapsed-epochs replays makes traffic less realistic in terms of the aggregate packet throughput. In contrast, the lossy full replay shows an excellent match for most of the time series. This result is different from the Leipzig-II one, where the full replays achieved a good approximation, but were still below the original packet throughput. Adding per-connection losses had a very minor impact on the Leipzig-II packet throughput, but the effect is substantial in the UNC 1 PM replay, where we observe increments of up to 2,000 packets per second. This result demonstrates the effectiveness of our source-level modeling method, and also justifies our effort to incorporate losses in the replay in order to study the realism of our modeling approach.

Figure 6.17 examines packet throughput in the outbound direction. Unlike the inbound direction, adding losses does not have a substantial impact here, and the aggregate packet throughput remains below the original trace even for the lossy full replay. As discussed in Section 6.2.2, this could be due to some limitations of our data acquisition algorithm in terms of how well it infers source-level characteristics, or to the use of the default MSS for all connections. As in previous cases, collapsed-epochs replays generate a substantially lower number of packets than full replays, which are far closer to the original packet throughput.

6.3.3 Marginal Distributions

The marginal distribution of byte throughput for the inbound direction of UNC 1 PM and its replays are shown in the Figure 6.18. The bodies of the distributions show that lossy replays provide a better

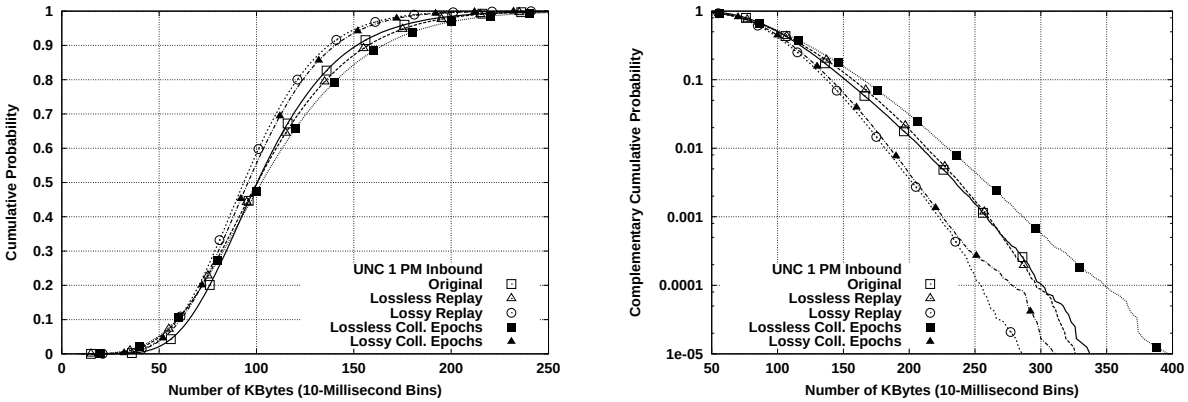

Figure 6.18: Byte throughput marginals for UNC 1 PM inbound and its four types of source-level trace replay.

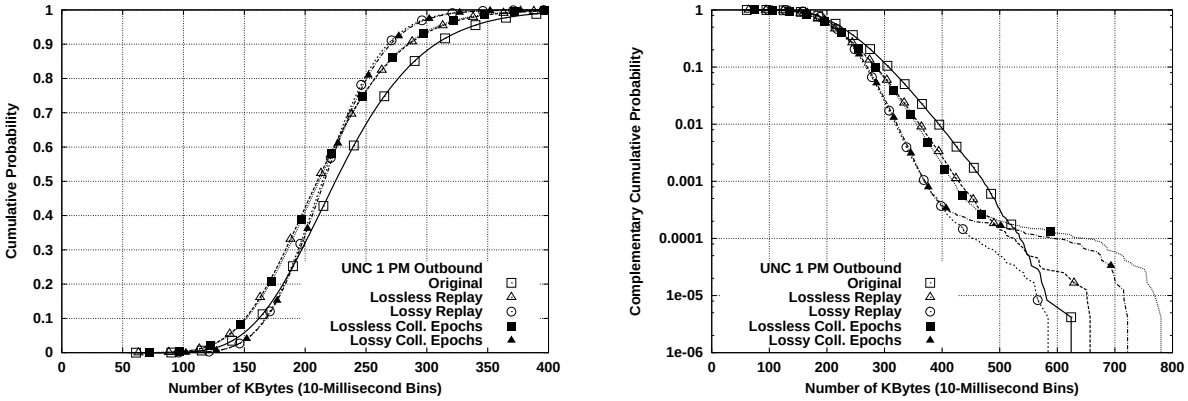

Figure 6.19: Byte throughput marginals for UNC 1 PM outbound and its four types of source-level trace replay.

approximation, although they are slightly heavier than the original. Interestingly, the analysis of the time series in Section 6.3.1 showed lower aggregate throughput from lossy replays, which seems inconsistent with the heavier bodies in the marginal distribution. The explanation is given by the plot of the tails of the marginals, which shows far lighter tails from the lossy replays. The way in which losses were incorporated in the experiments limited peak throughput substantially at the fine scales considered in the marginal plots. This is because the probability of artificial losses increases linearly with throughput, which is not generally true for real conditions. On the contrary, the lossless full replay reproduced the tail very accurately, demonstrating that the experimental environment and generation method are perfectly capable of reproducing the observed peak throughputs. It seems likely that further refinements in the implementation of per-connection losses, making them less open-loop, could make the tails closer to the original.

The marginal distributions in the outbound direction, which are shown in Figure 6.19, reveal a

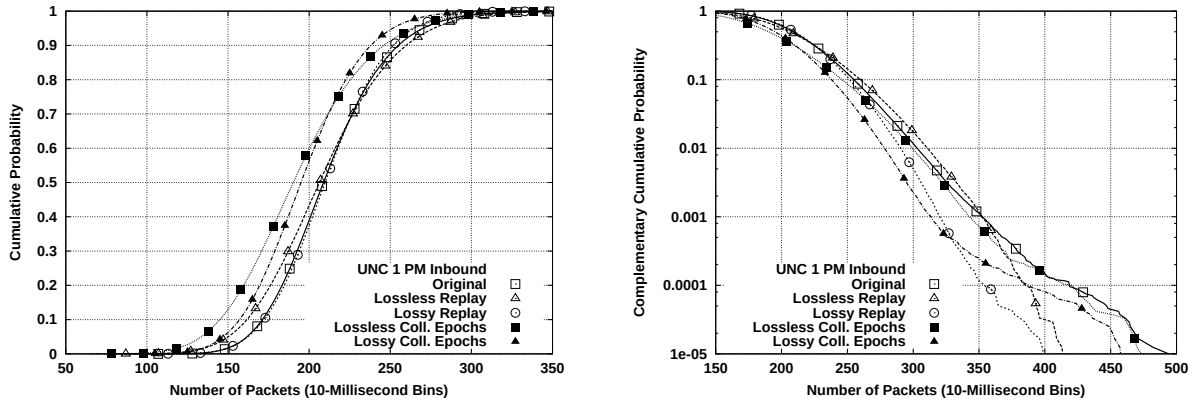

Figure 6.20: Packet throughput marginals for UNC 1 PM inbound and its four types of source-level trace replay.

somewhat worse approximation. We can distinguish three regions in the plot of the bodies. For values below 175 KB, lossless replays are lighter than the original, while lossy ones are heavier. Above 175 KB, all replays are lighter, which shows that the finding of lower aggregate byte throughput in Section 6.3.1 is due to overall lower throughputs at fine scales (rather than only to lighter tails). In the region after 175 KB, we can also observe that lossy replays are heavier below 275 KB and lighter above that. The marginal distributions from the lossy replays are less concentrated around the mean value, and are therefore somewhat more bursty, which is consistent with the similar finding for Leipzig-II (see Section 6.2.3).

Regarding the tails, we observe that for probabilities below 0.00075, the tail of original marginal is substantially heavier than the tails of the replay marginals. For probabilities above that, the collapsed-epochs replays show a major change in the shape of the distributions, being far heavier than the original for the largest values. We did not encounter a similar phenomenon in the Leipzig-II replays, where lossy collapsed-epochs replays always had a lighter tail than the lossless full replay. The number of 10-millisecond bins with very high throughput is larger for collapsed-epochs replays than for the full replays. Note that this artifact is only visible by looking at the tails of the marginals, and not at their bodies or at the time series of byte throughput.

The marginal distributions of packet throughput for UNC 1 PM inbound are shown in Figure 6.20. As observed for Leipzig-II, and as we may expect from 6.16, collapsed-epochs replays result in bodies that are significantly lighter than the body of the original marginal. Full replays are far closer, being the lossy full replay an excellent approximation of the original distribution. Interestingly, the tails reveal a rather different picture. Below 350 Kpps, the lossy replays have lighter tails than the original, especially in the case of the lossy full replay. Lossless replays closely approximate the original tail. Above 350

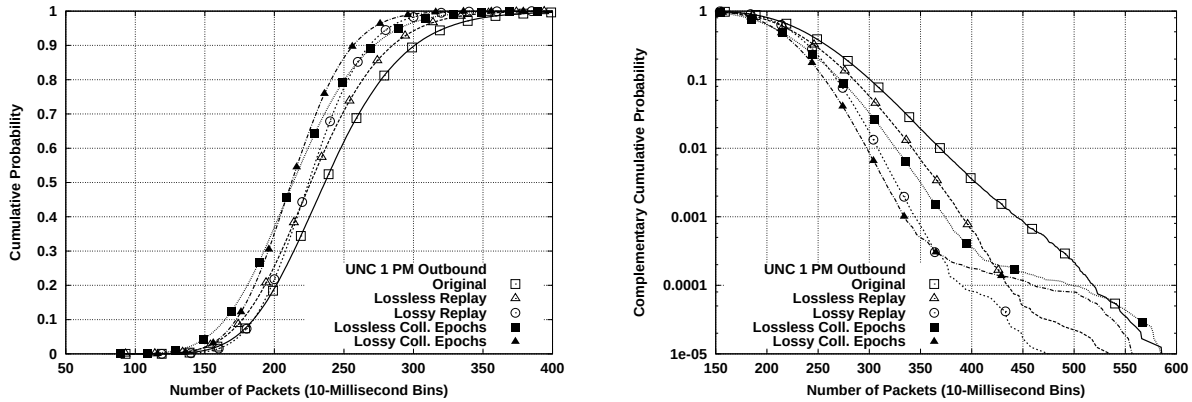

Figure 6.21: Packet throughput marginals for UNC 1 PM outbound and its four types of source-level trace replay.

Kpps, both full replays are lighter than the original, while the collapsed-epochs replays reproduce the probability of very high throughput bins accurately.

Figure 6.21 shows the marginal distributions of packet throughput in the opposite direction. All replays are lighter than the original, being the lossy full replay the closest one. The tails from the replays are also significantly lighter than the original tail. We also observe a similar change in the tail of the collapsed-epochs replays, which are very close to the original for the largest values.

6.3.4 Long-Range Dependence

The left plot of Figure 6.22 shows that the wavelet spectrum of the original byte throughput in the inbound direction is well approximated by both full replays for lower and medium octaves. The finding of this good match at the lower octaves differs from the result for the replay of the Leipzig-II trace, where this part of the wavelet spectrum was not so well approximated. The lossy replay shows less energy for octaves 8 and above, while there is a significant jump in the energy of the lossless replay for octaves 12 and above. In the right plot, the lossless collapsed-epochs replays shows substantially more energy for octaves above 4, while the lossy replay provides a better approximation.

For the outbound direction, the left plot of Figure 6.23 reveals a better approximation of the finest scales by the lossless full replay, while both full replays closely match the original spectrum at coarser scales. The right plot shows that both collapsed-epochs replays have less energy at the finest scales, with a rather sharp ditch for octaves 5 and 6 that was not present in the original. This ditch was far less pronounced in the full replays. Beyond the finest scales, the lossless collapsed-epochs replay

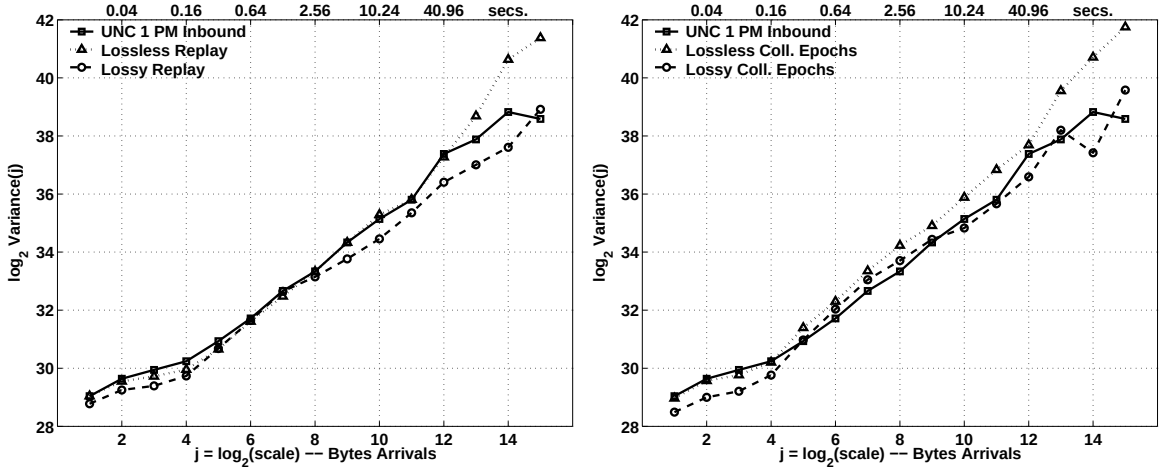

Figure 6.22: Wavelet spectra of the byte throughput time series for UNC 1 PM inbound and its four types of source-level trace replay.

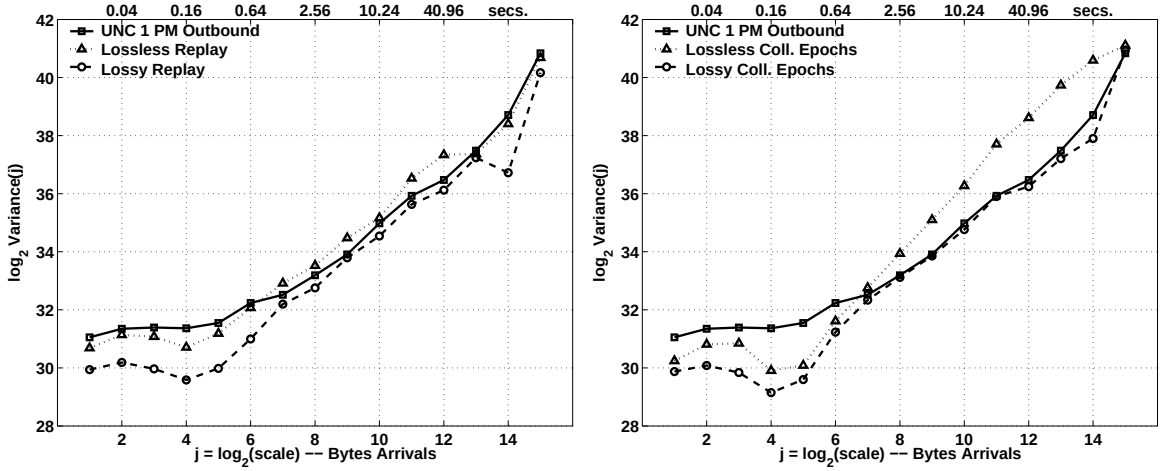

Figure 6.23: Wavelet spectra of the byte throughput time series for UNC 1 PM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 1 PM	H	C. I.	H	C. I.
Original	0.9557	[0.9113, 1.0002]	0.9717	[0.9272, 1.0161]
Lossless Replay	0.9632	[0.9188, 1.0077]	0.9585	[0.9141, 1.0030]
Lossy Replay	0.9118	[0.8674, 0.9563]	0.9306	[0.8861, 0.9750]
Lossless Coll. Epochs	0.9521	[0.9077, 0.9966]	1.0170	[0.9726, 1.0615]
Lossy Coll. Epochs	0.8441	[0.7996, 0.8885]	0.8657	[0.8212, 0.9101]

Table 6.3: Estimated Hurst parameters and their confidence intervals for the byte throughput time series of UNC 1 PM and its four types of source-level trace replay.

is a poor match of the original, while the lossy one provides a close approximation. This high impact of losses in the collapsed-epochs replay, far larger than in the full replay case, suggests a significant interaction between loss and long-range dependence when traffic is not generated according to a detailed source-level model. In other words, endpoints that generate traffic according to less realistic models (without epochs) are artificially more aggressive than Internet sources. This makes them more sensitive to lossy environments, since losses can more sharply decrease their higher throughput. This can result in experiments that overestimate the impact of losses on performance.

The estimated Hurst parameters and their confidence intervals shown in Table 6.3 are somewhat surprising. In the inbound direction, the estimated Hurst parameter of the original trace is most closely approximated by the lossless replays. The lossy full replay is slightly lower, and the lossy collapsed-epochs replay is far lower. The same is true in the opposite direction, at least for the lossless replays. It is difficult to interpret the meaning of these estimates in the context of the previous results. On the one hand, we found large spikes in the time series of byte throughput that suggest substantially higher burstiness in the lossless replays. Additionally, the wavelet spectra in Figure 6.22 did not find better approximations from the lossless replays. Notice for example that the lossless collapsed-epochs replay is clearly the farthest from the original. On the other hand, the tails of the marginal distributions clearly favored the lossless replays, showing lighter tails for the lossy replays. We could argue that the different metrics refer to different measures of burstiness, and conclude that adding artificial losses (using our open-loop method) makes the lossy replays less realistic in terms of Hurst parameter estimates. However, this conclusion seems too simplistic, since it is in contradiction with the Leipzig-II results. Adding losses made the estimated Hurst parameters far closer in that case. Assuming that the observed differences between the estimated Hurst parameters are significant, the reason for these divergent conclusions regarding the impact of losses must necessarily lie in some fundamental difference in the nature of the two network links. The estimated Hurst parameters say little about the difference, since all of the estimates are similarly high (above 0.92).

As discussed in Chapter 4, the Leipzig-II trace is a good example of university traffic dominated by downloading behavior (*i.e.*, inbound traffic is substantially higher than outbound traffic). In contrast, the UNC 1 PM trace is dominated by content downloaded from UNC servers (rather than downloads from UNC clients) due to the presence at UNC of a major Internet repository of software and content, `ibiblio.org`. This made traffic volume and number of connections far higher for UNC. Still, why would these differences make introducing losses beneficial in the Leipzig case and detrimental in the UNC case for the approximation of the original Hurst parameters? We can speculate that the rate-limiting

mechanisms used by `ibiblio.org` create unusual loss patterns that are poorly approximated by our open-loop losses, but we do not have any supporting evidence.

The lessons from the analysis of the scaling in the packet throughput series is quite similar. The plots in Figure 6.24 show reasonably close approximations of the original by all of the replays in the inbound direction, and somewhat worse ones in the outbound direction. The spectrum of the lossless full replay provides the closest approximation to the spectrum of the original in both directions. The spectrum of the lossless collapsed-epochs replay is clearly not as close, showing a higher slope for medium to coarse time scales. As in the case of byte throughput, lossy replays show less energy than the original trace, especially for the fine scales in the outbound direction. Note also the systematic ditch around octave 14 for all four spectra from lossy replays. This suggests some unexpected periodicities at the 1-minute scale. A similar ditch can be found in the outbound direction of the original time series in octave 13, and this ditch is not reproduced by the replays.

Regarding the estimated Hurst parameters and their confidence intervals, Table 6.4 shows different results for the two directions. The estimates for the inbound direction confirm the lossless full replay as an excellent approximation, but here the lossless collapsed-epochs replay is also very close to the original. Both lossy replays are well below the estimated Hurst parameter of the original time series, and outside its confidence interval. The estimates for the outbound direction show again an excellent approximation by the lossless full replay, but here the lossless collapsed-epochs replay is far higher than the original and well within the non-stationarity region. Lossy replays are substantially better in the outbound direction, with the lossy full replay matching the original estimate.

6.3.5 Time Series of Active Connections

As in the case of Leipzig-II, the lossy full replay of UNC 1 PM achieved a perfect match of the original time series of active connections This is clear for both the entire range of the time series shown in the left plot of Figure 6.26 using 1-minute bins, and for the 20-minute region shown in the right plot using 1-second bins. This finer scale view shows several sharp spikes (minutes 24, 28, 30, 35 and 39) that the lossy full replay tracked accurately. The lossless replay has only a slightly lower number of active connections per second, showing similar spikes (but with a negative offset in the y-axis). Collapsed-epochs replays had a far smaller number of active connections. Also, they did not track the features of the original time series so well. Notice for example the absence of the minute 24 spike in the collapsed-epochs replays.

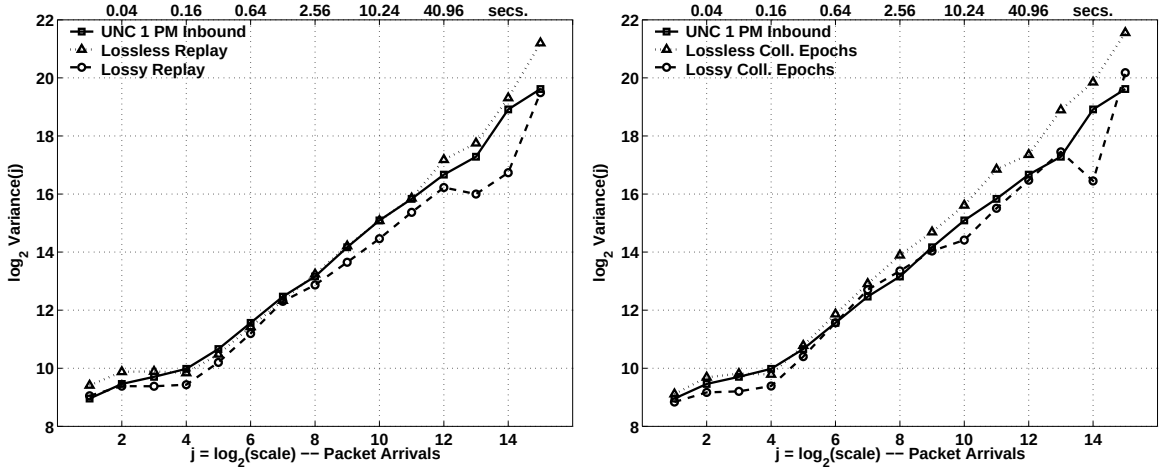

Figure 6.24: Wavelet spectra of the packet throughput time series for UNC 1 PM inbound and its four types of source-level trace replay.

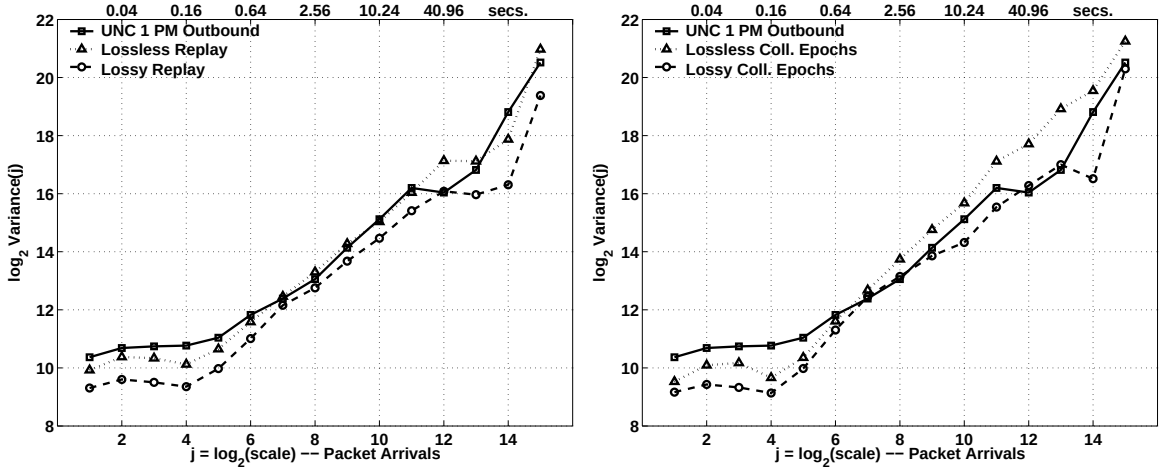

Figure 6.25: Wavelet spectra of the packet throughput time series for UNC 1 PM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 1 PM	H	C. I.	H	C. I.
Original	0.9564	[0.9158, 0.9970]	0.9339	[0.8933, 0.9746]
Lossless Replay	0.9776	[0.9370, 1.0182]	0.9512	[0.9106, 0.9918]
Lossy Replay	0.8719	[0.8313, 0.9125]	0.9512	[0.9106, 0.9919]
Lossless Coll. Epochs	0.9464	[0.9058, 0.9871]	1.0956	[1.0549, 1.1362]
Lossy Coll. Epochs	0.8509	[0.8103, 0.8916]	0.9200	[0.8793, 0.9606]

Table 6.4: Estimated Hurst parameters and their confidence intervals for the packet throughput time series of UNC 1 PM and its four types of source-level trace replay.

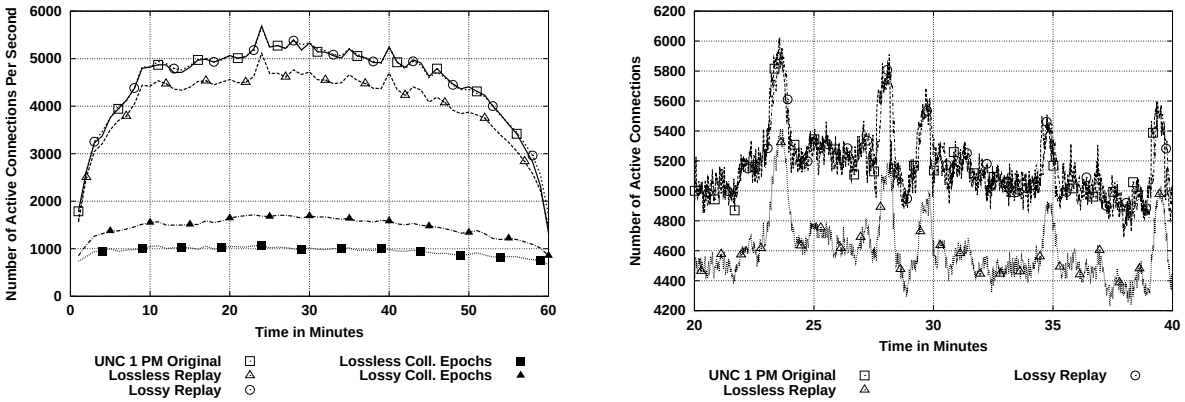

Figure 6.26: Active connection time series for UNC 1 PM and its four types of source-level trace replay.

6.4 Mid-Chapter Review

The present chapter is the longest one in this dissertation, and presents the results of 20 source-level trace replay experiments using 130 plots and 10 tables. The conclusions are not always straight-forward or consistent across traces, so it is difficult to form a coherent picture by simply going through the entire body of results. In this section, we summarize our results so far in order to make the rest of the chapter easier to follow. Our summary is in the form of a list of 18 observations, which report both on findings that were consistent for Leipzig-II and UNC 1 PM, and findings that were inconsistent.

6.4.1 Observations on Byte Throughput

From the analysis of the plots of the time series of byte throughput, their marginal distributions and wavelet spectra, we can make the following observations:

B.1 Both full and collapsed-epochs replays provide a reasonable approximation of the original 1-minute time series of byte throughput and the body of its 10-millisecond marginal. Replays do not track every spike in the original time series, but the similarity is remarkable. The replays achieve a very close approximation of the Leipzig-II time series, but are slightly below the UNC 1 PM time series. For both traces, the approximation of the bodies of the original marginal are somewhat better for the inbound direction than for the outbound one. This observation is not explained by traffic volume asymmetry, since the inbound direction was the dominant direction in terms of byte volume only in the case of Leipzig-II.

B.2 Lossless replays sometimes show substantially more spikes of 1-minute byte throughput above the

original trace than lossy replays. This is clear for UNC 1 PM but not for Leipzig-II. At the finer scales studied by the marginal distributions, we find that the tails of the lossless replays are substantially heavier than those of the lossy replays. However, they are not consistently above the tails of the original distributions. In contrast, the results for every trace show that the bodies of the lossless replays are wider than the bodies of the lossy replays. This reveals higher burstiness in the lossless replays in the sense that they have a higher probability of bins with byte throughput far from the mean (*i.e.*, a larger number of 10-millisecond intervals with have rather low or rather high byte throughput).

B.3 Collapsed-epochs replays show somewhat more bursty 1-minute time series, and track the changes in the shape of the original time series less closely. The extra burstiness may not appear very substantial in the plots, but given the coarse scale, it may have a large impact on experiments sensitive to prolonged byte throughput spikes. We do not find a corresponding phenomenon for the marginal distributions, where collapsed-epochs replays are generally close to the full replays (except for the outbound direction of UNC 1 PM). Together with observation B.5, this shows that the extra burstiness of the collapsed-epochs replays manifests itself in the auto-correlation structure of the byte throughput process, rather than in the set of byte throughputs observed throughout the replays.

B.4 Full replays provide a close approximation of the scaling region (octaves 6 to 15) of the wavelet spectra of the original traces. This does not necessarily translate into similarly good approximations of the estimated Hurst parameters. Only the lossy replays are within confidence intervals for Leipzig-II, while only the lossless ones are within confidence intervals for UNC 1 PM.

B.5 Collapsed-epochs replays tend to show slightly more energy in the scaling region. This is true for the four spectra from lossless replays and for the two spectra from lossy replay of Leipzig-II. However, the energy of the original scaling region is well approximated by the lossy collapsed-epochs replay for the outbound direction of UNC 1 PM. This higher energy in the wavelet spectrum plot does not necessarily translate into higher estimates of the Hurst parameters.

B.6 Both full and collapsed-epochs replays do not consistently match the spectra of the finer scales (octaves 1 to 5). We find higher or slightly higher energy levels for the replays of Leipzig-II, similar levels for the replays of the inbound direction of UNC 1 PM and lower levels for the outbound direction of UNC 1 PM.

B.7 By construction, the most detailed replay is the lossy full replay, so we expect it to achieve the best approximation of the original trace. This was always true for 1-minute time series, the body

of the marginal distribution and the scaling region of the wavelet spectrum. However, it was not consistently true for the tail of the marginal distribution, the energy of the wavelet spectrum at fine scales, and the estimated Hurst parameter.

6.4.2 Observations on Packet Throughput

We can make the following observations regarding packet throughput:

P.1 Full replays achieve a close approximation of the original 1-minute time series of packet throughput, remaining between 2% and 8% below the original for most of the time series. Collapsed-epochs replays result in a substantially worse approximation, being between 20% to 30% below the original for most of the time series. This difference is also present in the bodies of the 10-millisecond marginal distributions. In the best case for full replays, the median of the marginal distribution is equal to the original median for the inbound direction of the UNC 1 PM lossy replay. In the worst case, the median is 7% below the original for the inbound direction of the Leipzig-II lossy replay. Collapsed epochs replays show medians of the marginal distributions that are 20% (UNC 1 PM inbound) and 25% (Leipzig-II outbound) below the original median.

P.2 Incorporating losses into the replays increases packet throughput, reducing the distance to the original time series. While this effect is small for Leipzig-II, it is rather significant for UNC 1 PM inbound. In addition, lossless replays sometimes show more artificial spikes in the 1-minute time series plot than the lossy ones (*e.g.*, UNC 1 PM outbound). This phenomenon seems less prominent for packet throughput than for byte throughput (see observation B.2).

P.3 Unlike the byte throughput case, the tails of the packet throughput from the replays marginals are never significantly heavier than the original tails. Lossless replays provide the best approximations of the original tails, being excellent in some cases (Leipzig-II inbound and UNC 1 PM inbound). Lossy replays show lighter tails than lossless replays, revealing significantly worse approximations of the original tails. We can also observe that the tails of the collapsed-epochs replays are consistently lighter than those of the full replays. However, the impact of detailed modeling on the tails of the marginals is less prominent than the impact of incorporating losses.

P.4 Full replays and lossy collapsed-epochs replays provide good approximations of the original wavelet spectra, while the lossless collapsed-epochs replays show somewhat higher energy. In general, we can say that the best approximation is achieved by the lossless full replay. As in the case of byte throughput, Hurst parameter estimates offer a different picture. Only the estimates for the

lossy full replay are within confidence intervals of the original estimates for Leipzig-II, while the estimates for both lossless and lossy full replays are within confidence intervals for UNC 1 PM.

P.5 Replays do not consistently reproduce the energy levels at the finest scales of the original time series of packet arrivals. We find minor differences for Leipzig-II and UNC 1 PM inbound, and substantially larger ones for UNC 1 PM outbound. Collapsed-epochs replays are significantly worse than full replays only for UNC 1 PM.

6.4.3 Observations on Active Connections

Regarding active connections, we can make the following observations that hold true for both Leipzig-II and UNC 1 PM:

C.1 The number of active connections in the original trace and in the full replays is very similar.

C.2 The lossy full replay provides the best approximation of the active connection time series, being within 1% of the original time series. There is no difference for UNC 1 PM.

C.3 The number of active connections in collapsed-epochs replays is several times smaller than the original (around 3 times smaller for Leipzig-II and UNC 1 PM).

C.4 Adding losses to the replays substantially increases the average number of connections. This increase is of the same magnitude for both full and collapsed-epochs replays.

C.5 Full replays track the features of the original time series very closely. The only difference between lossless and lossy replays is a slowly varying offset. This suggests a homogeneous impact of losses, which lengthens the lifetimes of a stable number of connections throughout the traces.

C.6 Unlike full replays, collapsed-epochs replays do not track the features of the original time series. However, the magnitude of this effect pales in comparison to the much smaller number of active connections.

6.5 Source-level Replay of UNC 1 AM

6.5.1 Time Series of Byte Throughput

The plots of the 1-minute time series of byte throughput for the original UNC 1 AM and its replays are shown in Figure 6.27 (inbound direction) and in Figure 6.28 (outbound direction). For the inbound,

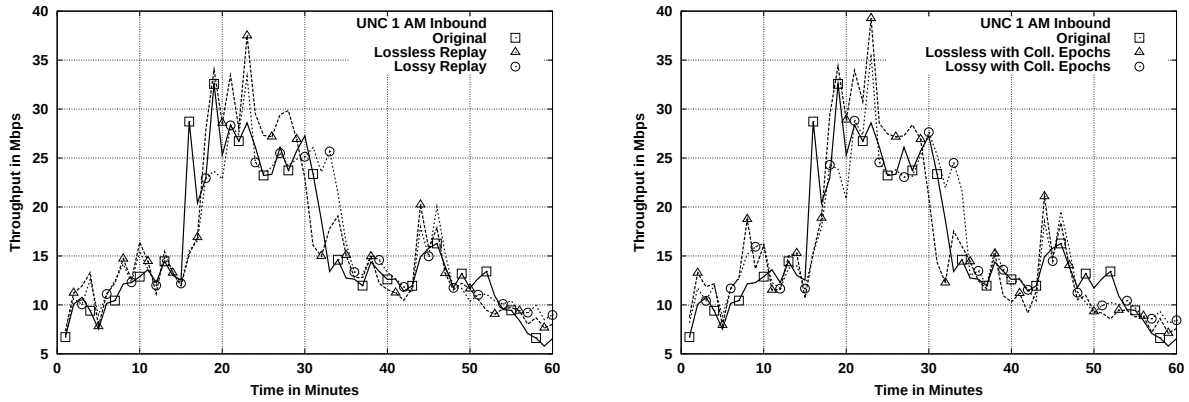

Figure 6.27: Byte throughput time series for UNC 1 AM inbound and its four types of source-level trace replay.

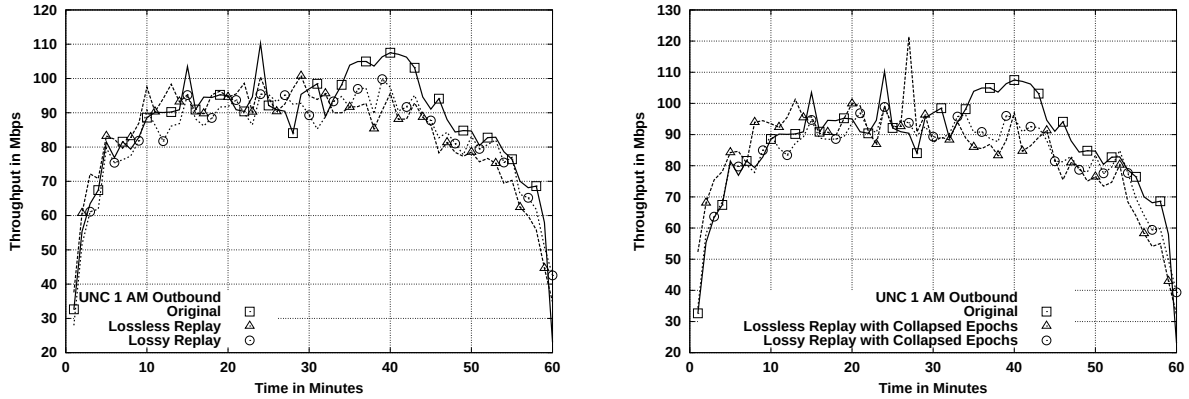

Figure 6.28: Byte throughput time series for UNC 1 AM outbound and its four types of source-level trace replay.

we observe a moderately bursty time series with a large increase in byte throughput between minutes 15 and 32. In good agreement with observation B.1, the replays track the shape of the original time series well. They also approximate some smaller spikes, such as the one in minute 45, and miss others, such as the one in minute 17. The result is similar for the outbound direction, although we again find a slightly lower overall throughput in the replays. There is also an area of higher throughput in the original trace between minutes 35 and 43 that is not properly reproduced by any of the replays. The full lossy replay provide the closest approximation, but there is still a clear difference with respect to the original time series.

The results also support the observation of higher burstiness from lossless replays, B.2, and from collapsed-epochs replays, B.3; especially for the inbound direction. The results are also consistent with observation B.7, since the full lossy replay appears closest to the original.

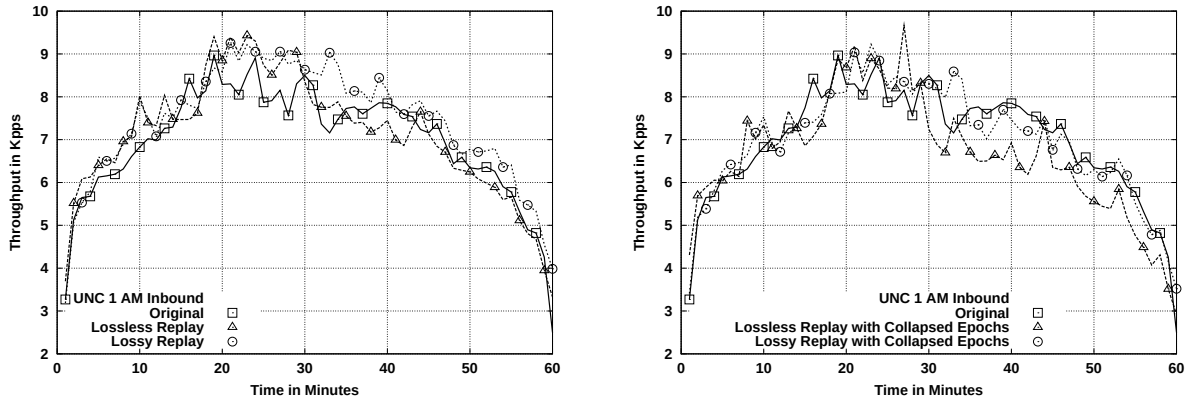

Figure 6.29: Packet throughput time series for UNC 1 AM inbound and its four types of source-level trace replay.

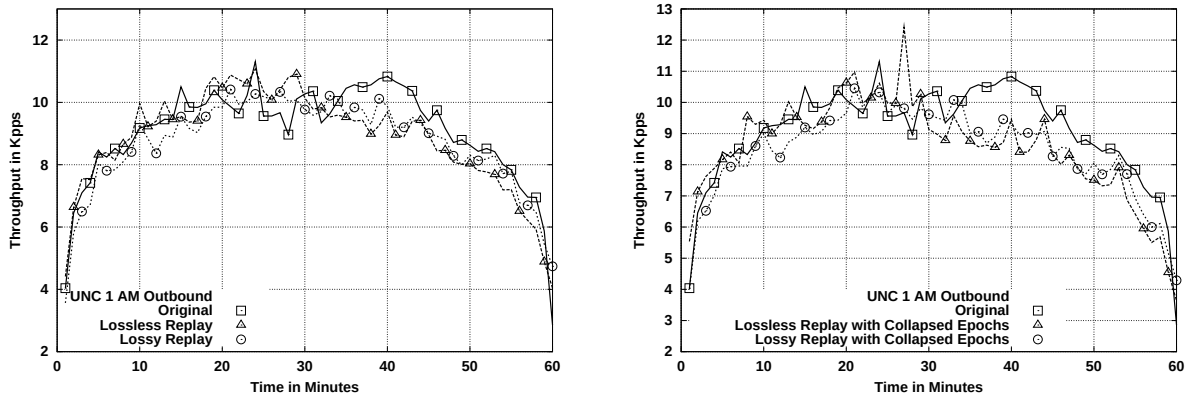

Figure 6.30: Packet throughput time series for UNC 1 AM outbound and its four types of source-level trace replay.

6.5.2 Time Series of Packet Throughput

The time series of packet throughput for UNC 1 AM inbound shown in Figure 6.29 are in sharp contrast to earlier results. As stated in observation P.1, the time series from the replays of the previous traces were generally below the time series of the original trace. However, the full replays of UNC 1 AM are often above the original packet throughput, especially in the case of the lossy full replay. The same is not true for the outbound direction, as shown in Figure 6.30, where the replays are again below the original for a large fraction of the time series. While the replays provide a reasonable approximation of the overall time series, the original packet throughput in the outbound direction is substantially lower between minutes 35 and 43. The difference is most apparent for the collapsed-epochs replays.

Regarding observation P.2, we can see that collapsing epochs substantially reduced packet throughput. Paradoxically, this makes the time series of the lossy collapsed-epochs match the original quite

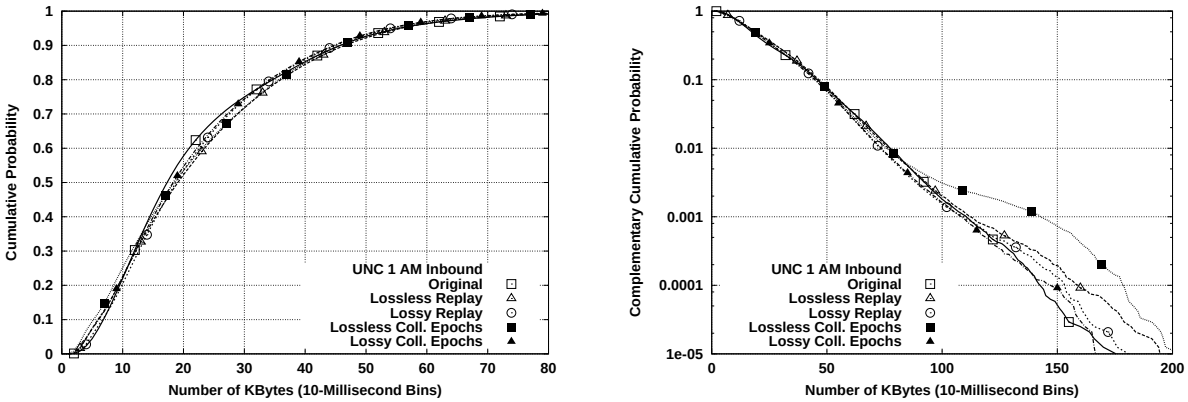

Figure 6.31: Byte throughput marginals for UNC 1 AM inbound and its four types of source-level trace replay.

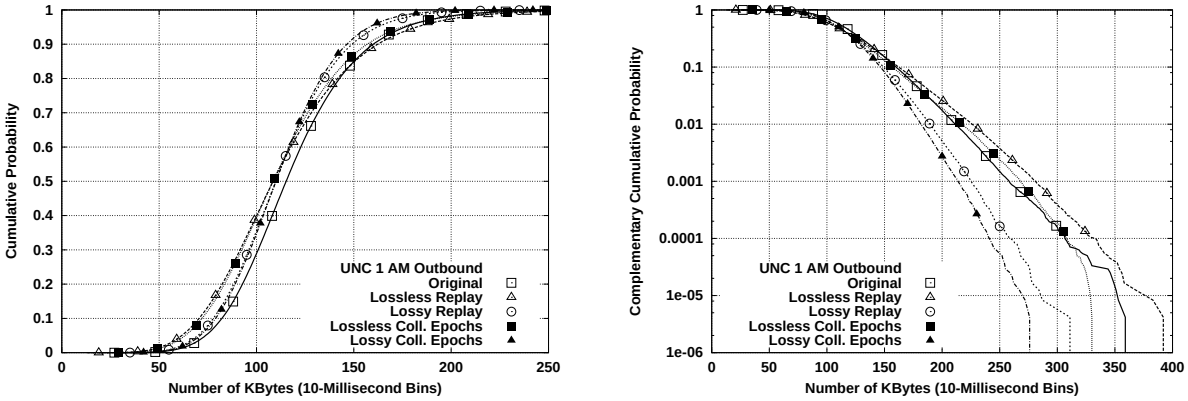

Figure 6.32: Byte throughput marginals for UNC 1 AM outbound and its four types of source-level trace replay.

well, although the same is not true for the lossless collapsed-epochs replay. Note also that it is difficult to argue for this trace that the lossy replays are significantly more bursty than the lossy ones at the 1-minute scale. We only observe one artificial spike in minute 27 for the lossless collapsed-epochs replay.

6.5.3 Marginal Distributions

Figures 6.31 and 6.32 study the marginal distributions of the original 10-millisecond time series UNC 1 AM and those from the source-level trace replays. The bodies of the distributions from the replays are almost identical to the original for the inbound direction, and quite close for the outbound direction, which further supports observation B.1. Observation B.2 is consistent with these results, although the bodies of the lossless replays are very close to those of the lossy ones in this case. We do however observe consistently heavier tails from lossless replays. Note the much heavier tail from the lossless collapsed-

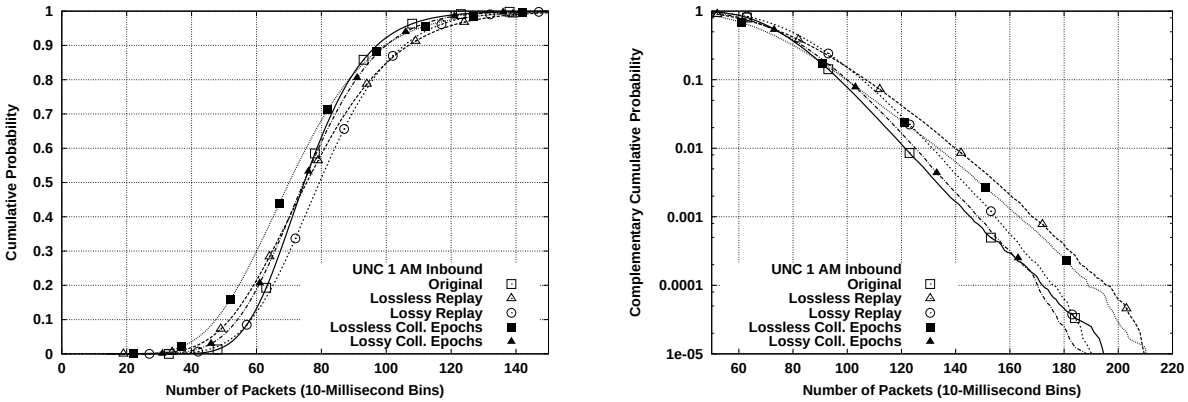

Figure 6.33: Packet throughput marginals for UNC 1 AM inbound and its four types of source-level trace replay.

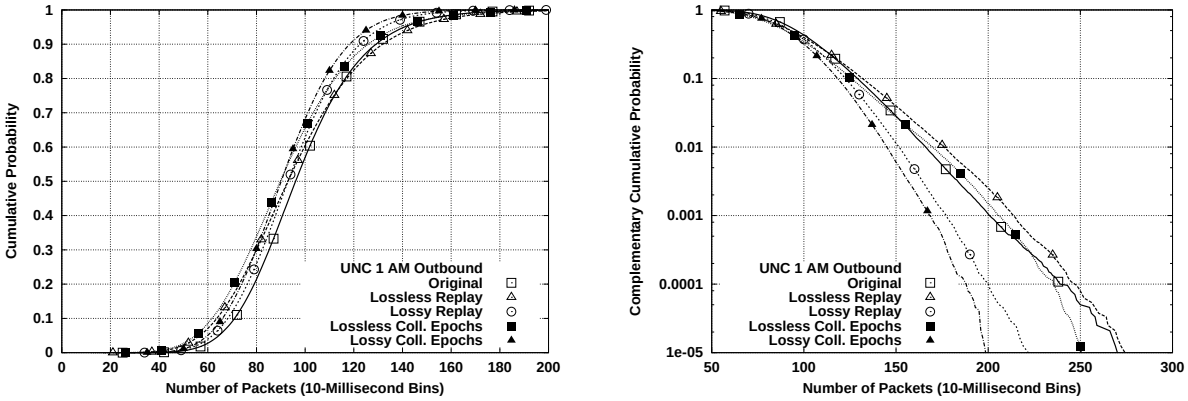

Figure 6.34: Packet throughput marginals for UNC 1 AM outbound and its four types of source-level trace replay.

epochs replay, which reveals an extra burstiness that was not visible in Figure 6.27. In agreement with observation B.3, we do not find consistently wider bodies or heavier tails from the collapsed-epochs replays. Finally, observation B.7 remains valid, with the lossy full replay being best for the tail of the inbound direction, but clearly not for the opposite direction.

The lesson from the plots of the packet throughput marginals shown in Figures 6.33 and 6.34 is similar to the one discussed for the time series in Section 6.5.2. In the inbound direction, the marginal from the lossy full replay is heavier than the original, while the lossless full replay and the lossy collapsed-epochs replays are rather close to the original. In the outbound direction, the results are consistent with the somewhat lower packet throughput for the replay stated in observation P.1.

The tails of the marginals are again surprising for UNC 1 AM, and do not follow observation P.3. The inbound plot shows lossless replays that are significantly heavier than the original, which exhibits the lightest tail. Lossy replays provide far better approximations. The outbound plot appears closer to

the previous observation, with the lossless replays being the closest ones to the original. Note however that they are somewhat heavier, unlike in the Leipzig-II and UNC 1 PM cases.

6.5.4 Long-Range Dependence

While the wavelet spectra for the inbound direction shown in Figure 6.35 are in good agreement with observation B.4, we find substantially higher energy above the original in the spectrum of the lossless full replay for outbound direction. The estimated Hurst parameters shown in Table 6.5 are again difficult to assess, as mentioned in that observation. Lossless replays are the only ones within the confidence interval of the estimate for the original inbound direction, while only the lossless collapsed-epochs replay is outside the confidence interval for the outbound direction. Incidentally, the extremely high estimate for the lossless collapsed-epochs replay is remarkable. It is 0.23 above the lossless full replay, illustrating the major difference that detailed source-level modeling can make on traffic long-range dependence.

In the scaling region, collapsed-epochs replays do show higher energy than full replays, as observed in B.5. This higher energy does not translate into higher Hurst parameter estimates. Notice for example the lower estimates for the inbound direction. For both directions, the lossy collapsed-epochs replay provides a good approximation of the original spectrum, although not as good as the lossy full replay. The results for UNC 1 AM are therefore consistent with observation B.7. At the finest scales, we find that the lossy full replay approximates the energy levels of the inbound direction most closely, while it is the lossy collapsed-epochs replay the best match for the outbound direction. This inconsistency supports observation B.6.

Figures 6.37 and 6.38 reveal that the wavelet spectra from lossless replays do not approximate the original spectra well. For both directions, the full lossless replay shows significantly more energy, while the full collapsed-epochs replay shows higher slope in the scaling region. This poor fit for the lossless full replay contradicts observation P.4. Lossy replays appear closer to the original spectra in the scaling region, especially in the case of the lossy full replay. Replays do not consistently match the energy in the fine-scale region, as stated in observation P.5. Lossy replays are the closest ones in this region.

The estimated Hurst parameters shown in Table 6.6 do not follow observation P.4 very clearly. The estimates from the lossless collapsed-epochs replay are far larger than the original estimates. The estimates from the lossless full replays are far lower, but they are still above the upper ends of the confidence intervals. Finally, both lossy replays are within confidence intervals, although the actual

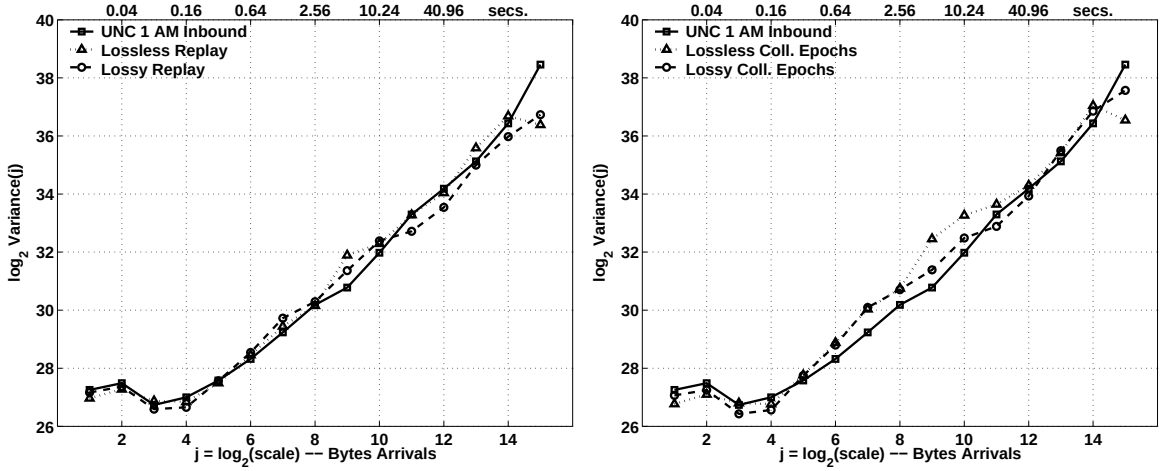

Figure 6.35: Wavelet spectra of the byte throughput time series for UNC 1 AM inbound and its four types of source-level trace replay.

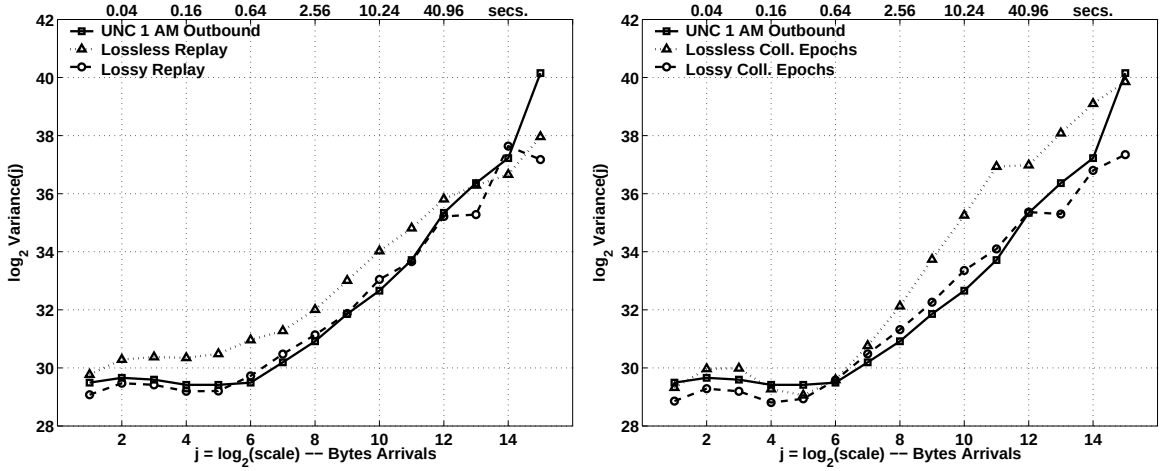

Figure 6.36: Wavelet spectra of the byte throughput time series for UNC 1 AM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 1 AM	H	C. I.	H	C. I.
Original	0.9885	[0.9479, 1.0292]	0.9990	[0.9584, 1.0397]
Lossless Replay	1.0275	[0.9868, 1.0681]	0.9705	[0.9299, 1.0111]
Lossy Replay	0.9465	[0.9058, 0.9871]	0.9546	[0.9140, 0.9953]
Lossless Coll. Epochs	1.0089	[0.9683, 1.0495]	1.2036	[1.1630, 1.2443]
Lossy Coll. Epochs	0.9136	[0.8730, 0.9542]	0.9720	[0.9313, 1.0126]

Table 6.5: Estimated Hurst parameters and their confidence intervals for the byte throughput time series of UNC 1 AM and its four types of source-level trace replay.

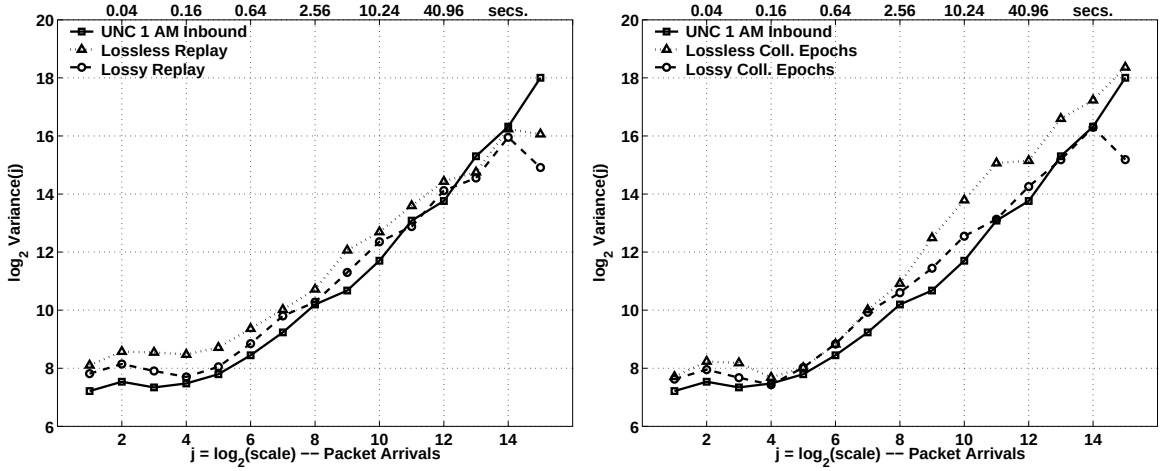

Figure 6.37: Wavelet spectra of the packet throughput time series for UNC 1 AM inbound and its four types of source-level trace replay.

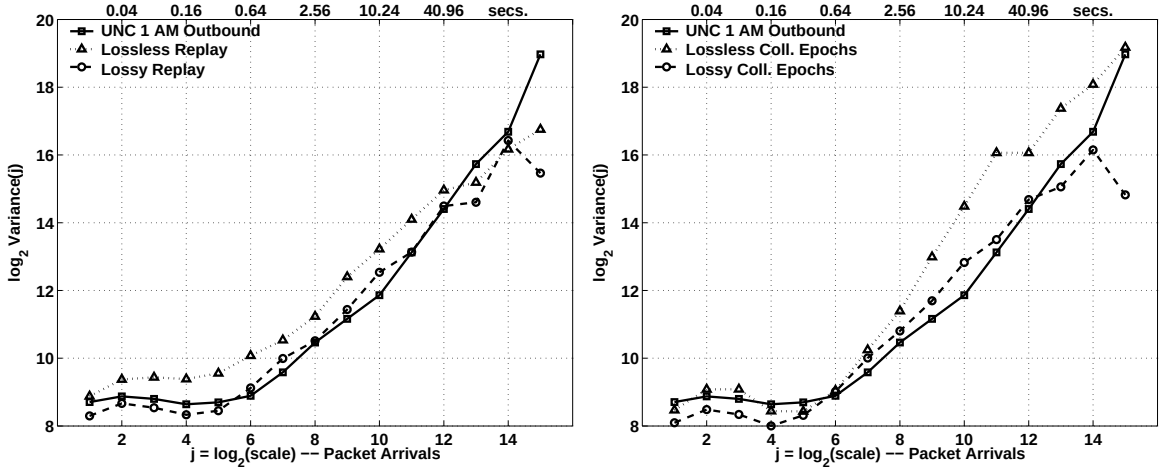

Figure 6.38: Wavelet spectra of the packet throughput time series for UNC 1 AM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 1 AM	H	C. I.	H	C. I.
Original	0.9316	[0.8871, 0.9760]	0.9309	[0.8864, 0.9753]
Lossless Replay	0.9860	[0.9416, 1.0305]	0.9830	[0.9385, 1.0274]
Lossy Replay	0.9749	[0.9304, 1.0193]	0.9759	[0.9315, 1.0204]
Lossless Coll. Epochs	1.1478	[1.1034, 1.1923]	1.2128	[1.1683, 1.2572]
Lossy Coll. Epochs	0.9504	[0.9059, 0.9948]	0.9757	[0.9313, 1.0202]

Table 6.6: Estimated Hurst parameters and their confidence intervals for the packet throughput time series of UNC 1 AM and its four types of source-level trace replay.

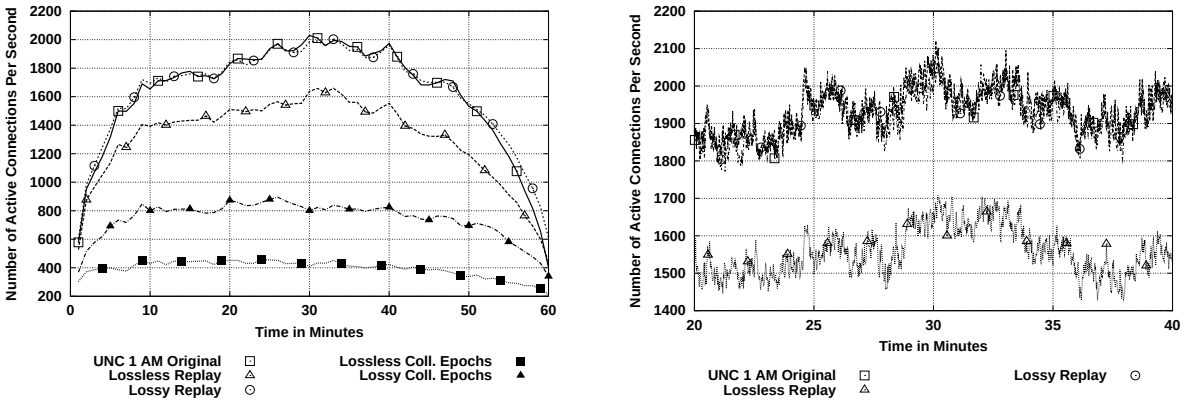

Figure 6.39: Active connection time series for UNC 1 AM and its four types of source-level trace replay.

estimates are higher.

6.5.5 Time Series of Active Connections

The time series of active connections shown in Figure 6.39 confirm the list of observations in Section 6.4. It is clear that observations C.1 and C.2 hold, being the lossy full replay a perfect match of the original time series. Observation C.3 is also true, although the relative gap between the number of connections in full and collapsed-epochs replays is smaller for this trace. The impact of losses is somewhat more significant for the collapsed-epochs replays, which is not completely in agreement with observation C.4. Observations C.5 and C.6 are consistent with the results for UNC 1 AM.

6.6 Source-level Replay of UNC 7:30 PM

6.6.1 Time Series of Byte Throughput

The time series of byte throughput for the inbound and outbound directions of UNC 7:30 PM are shown in Figures 6.40 and 6.41 respectively. Observation B.1 is clearly applicable to these results. For the inbound direction, note the very good approximation of the time series features between minutes 20 and 60, and the accurately reproduced spikes in minutes 46 and 53. In contrast, the replay seems out of phase for the initial spike in minute 1, and the large spike in minute 32. This could be explained by one or a few fast connections in the original trace that could not be replayed fast enough. In the outbound direction, we find replays with somewhat lower byte throughput, which was also observed in the other

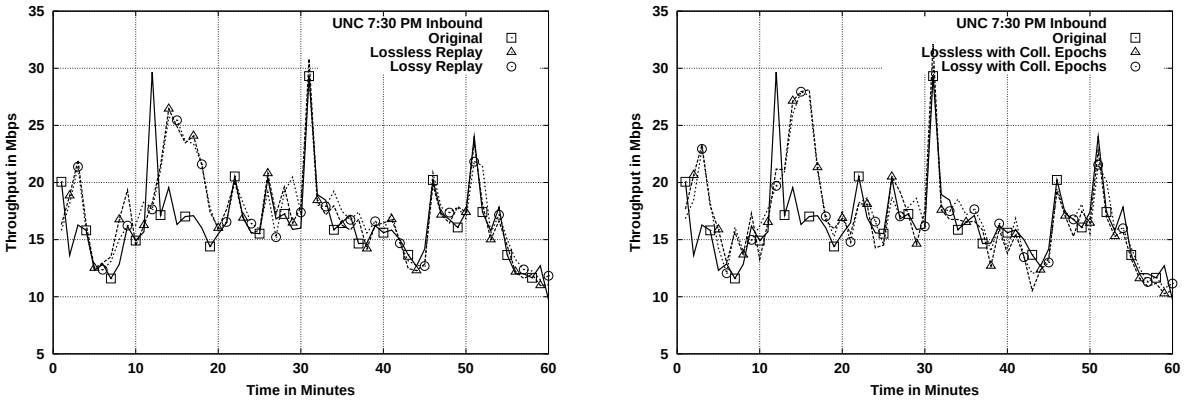

Figure 6.40: Byte throughput time series for UNC 7:30 PM inbound and its four types of source-level trace replay.

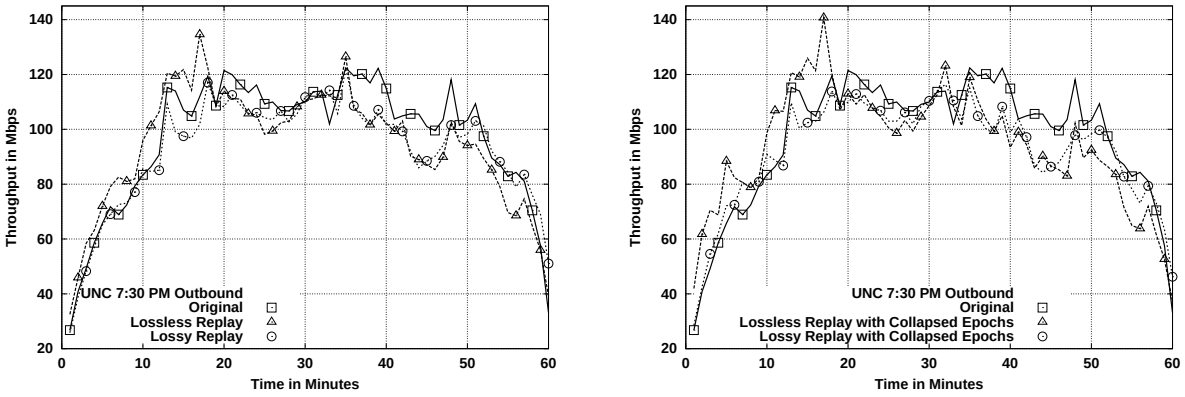

Figure 6.41: Byte throughput time series for UNC 7:30 PM outbound and its four types of source-level trace replay.

two UNC traces.

The possible extra burstiness in lossless replays mentioned in observation B.2 is not present in the inbound direction, and the last 40 minutes in the outbound direction. We do however observe substantially higher throughputs in the outbound direction for the first 20 minutes, especially in the case of the lossless collapsed-epochs replay. Regarding observation B.3, we do observe slightly more bursty time series from the collapsed-epochs replay in both directions, although the difference seems minor in this case. A few of the (smaller) features in the inbound direction are more closely approximated by the full replays, such as the spike in minute 22 and the ditch in minute 43.

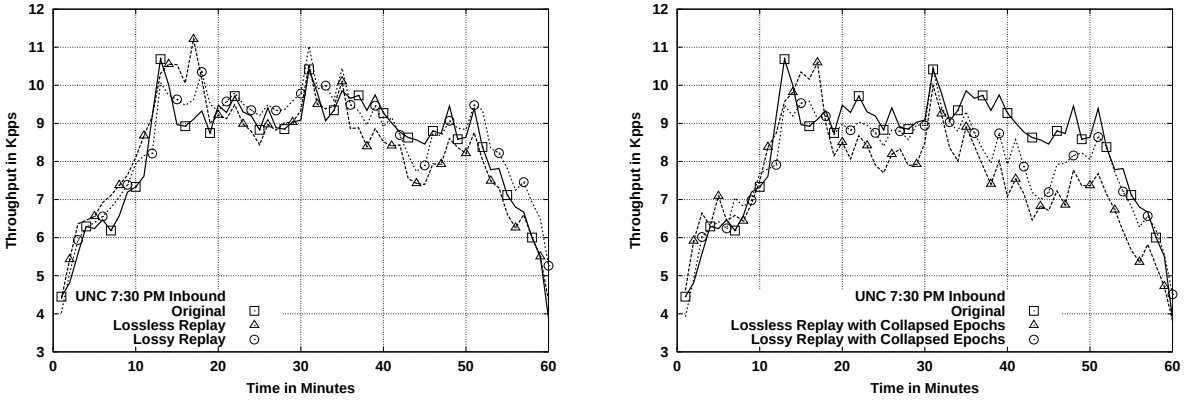

Figure 6.42: Packet throughput time series for UNC 7:30 PM inbound and its four types of source-level trace replay.

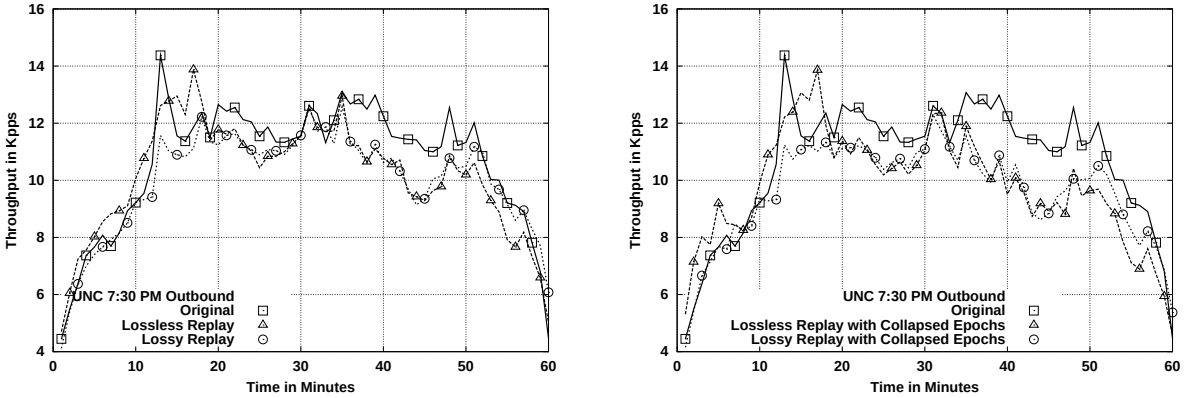

Figure 6.43: Packet throughput time series for UNC 7:30 PM outbound and its four types of source-level trace replay.

6.6.2 Time Series of Packet Throughput

The analysis of the packet throughput results shows again some interesting differences with respect to earlier results and observations P.1 and P.2, but only in the outbound direction. The results for the inbound direction presented in Figure 6.42 are in good agreement with observation B.1, since we observe substantially lower packet throughput for collapsed-epochs replays. The result is also consistent with observation B.2, showing higher packet throughput in the lossy replay. However, the result for the outbound direction is more surprising. Unlike previous cases, losses have a minimal impact on the replays, as shown in Figure 6.43. We could argue that the lossy replay provides a better fit between minutes 30 and 35 and after minute 52, but the rest of the time series for this replay is very similar to the one for the lossless replay. We can also say that the lossless replay makes a better attempt to match the spike in minute 14, although the replay spike seems shifted to minute 16.

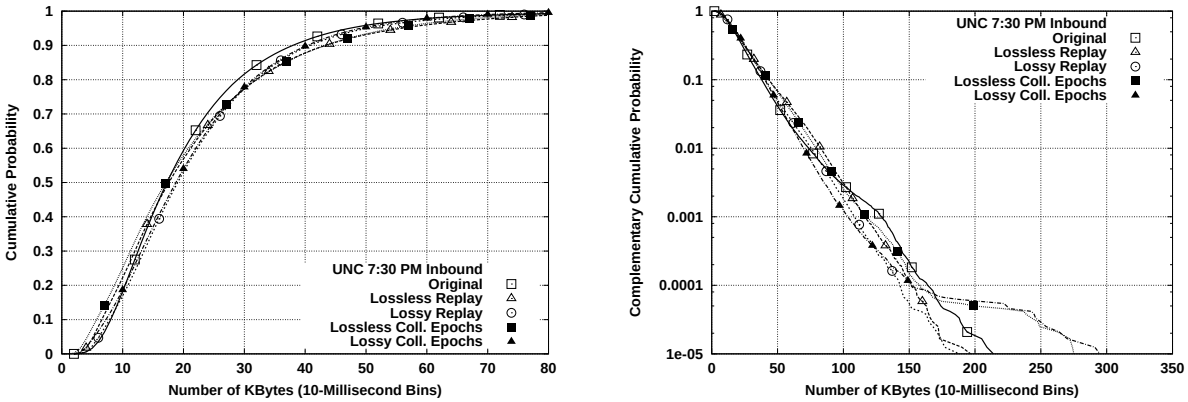

Figure 6.44: Byte throughput marginals for UNC 7:30 PM inbound and its four types of source-level trace replay.

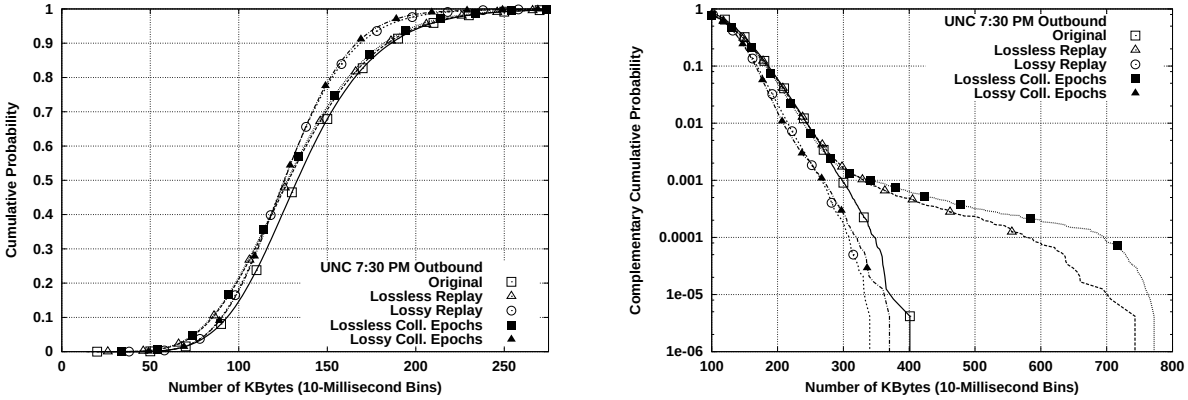

Figure 6.45: Byte throughput marginals for UNC 7:30 PM outbound and its four types of source-level trace replay.

6.6.3 Marginal Distributions

The marginal distributions of the 10-millisecond time series of byte throughput for the inbound direction are shown in Figure 6.44, while the ones for the outbound direction are shown in Figure 6.45. In agreement with observation B.1, the body of the marginals are closely approximated by the replays, especially in the case of the inbound direction. For the outbound, it is interesting to note a better approximation by the lossless replays for the upper part of the body and the first half of the tail.

As observation B.2 and B.3 pointed out, it is difficult to make general a statement about the approximation of the tails. For UNC 7:30 PM inbound, collapsed epoch replays match the original as closely as the full replays that match the original tail below 10^{-4}, but they are substantially heavier above that probability. Note that this heaviness did not manifest itself in the plots of 1-minute byte throughput. For UNC 7:30 PM outbound, we however have that the lossless replays are the ones showing an excellent

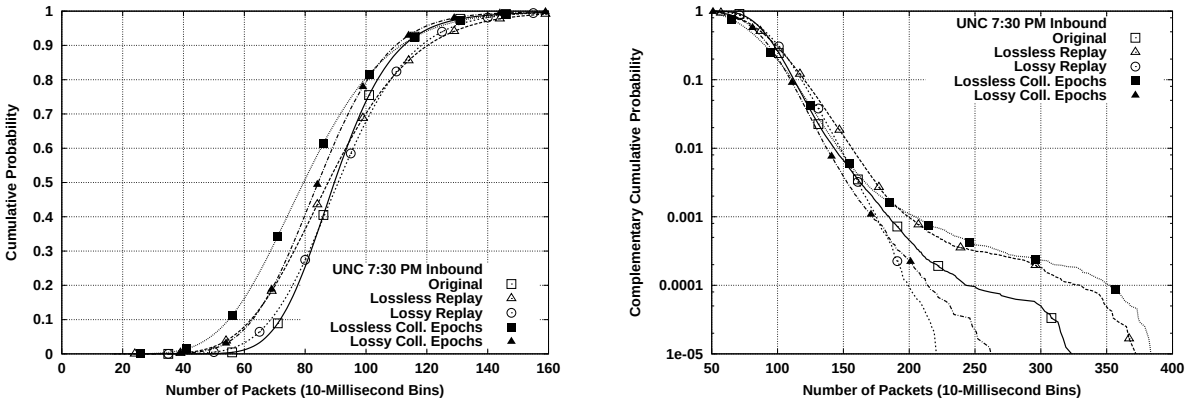

Figure 6.46: Packet throughput marginals for UNC 7:30 PM inbound and its four types of source-level trace replay.

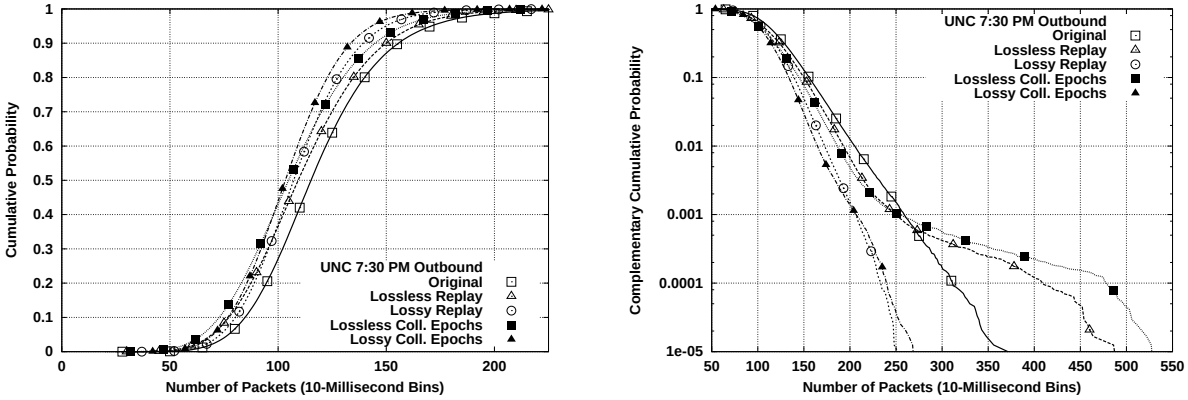

Figure 6.47: Packet throughput marginals for UNC 7:30 PM outbound and its four types of source-level trace replay.

match below 10^{-3}, but a far heavier tail above that probability. Overall, the only type of replay that did not show an overly heavy tail was the lossy full replay, which supports observation B.7.

The body of the marginal distributions for the 10-millisecond packet throughputs shown in Figure 6.46 and 6.47 do not clearly follow observation P.1. In the inbound direction, the distributions from the collapsed-epochs replays are clearly lighter than those from the full replays and the original trace for most of the distribution. However, they provide a better approximation above 100 packets. Interestingly, the distributions from the lossless replays exhibit similar shapes, but a clear offset, and the same is true for the lossy replays. For this direction, we find that the impact of detailed source-level modeling and per-connection losses is of the same order. The lesson is similar for the outbound direction, although all of the replay distributions are lighter than the original distribution in this case.

Regarding the tails of the marginal distributions, we observe similar conditions in both directions. Lossless replays exhibit substantially heavier tails than the original, while lossy replays exhibit substan-

tially lighter tails. This is in sharp contrast to the results for the 1-minute time series studied in Section 6.6.2, where losses had a very small impact. We can argue that lossless replays are artificially more bursty only at fine-time scales for UNC 7:30 PM. The results clearly support observation P.3, showing that the tails are far more sensitive to losses than to detailed source-level modeling for this trace.

6.6.4 Long-Range Dependence

The spectra of the byte throughput in the full replays are close to the original spectrum, as shown in Figure 6.48. However, the spectra of the outbound byte throughput shown in Figure 6.49 reveals a lossless replay with substantially more energy in the scaling region (which starts at octave 6). As in the case of UNC 1 AM outbound, this finding does not support observation B.4 regarding lossless full replays. The lossy full replay provides however a closer approximation to the original spectrum. Estimated Hurst parameters in Table 6.7 show similar results. Estimates from the replays are within the confidence interval of the inbound estimate, but somewhat lower. However, they are above the upper end of the confidence interval of the outbound estimate. The estimate from the lossy collapsed-epochs replay is specially high in this case, probably driven by the spike in octave 11. It is again difficult to draw any strong conclusion other than the general finding of inconsistent results already stated in observation B.4.

Collapsed-epochs replays show substantially more energy in the lossless case, but the difference in not so substantial for the lossy replay, especially in the inbound direction. This is in agreement with observation B.5.

The lossy full replay provides the best approximation again, as pointed out in observation B.7, However, some regions, such as the one between octaves 9 to 12 in the inbound direction, are more closely reproduced by the lossy collapsed-epochs replay. Interestingly, the four replays track the fine-scale energy profile of the original spectrum for the inbound direction, but only the lossless ones do so for the outbound direction. Observation B.6 already reflected this type of inconsistency in the results.

The lessons from the wavelet spectra in Figures 6.50 and 6.51 are surprisingly similar to those from the analysis of the packet throughput spectra for UNC 1 AM, discussed in Section 6.5.4. The full lossless replay shows higher energy, and the full collapsed-epochs replay shows substantially higher slope in the scaling region. Lossy replays provide far better approximations of the original spectra. As we observed for UNC 1 AM, these findings are inconsistent with observation P.4. Regarding fine scale energy levels

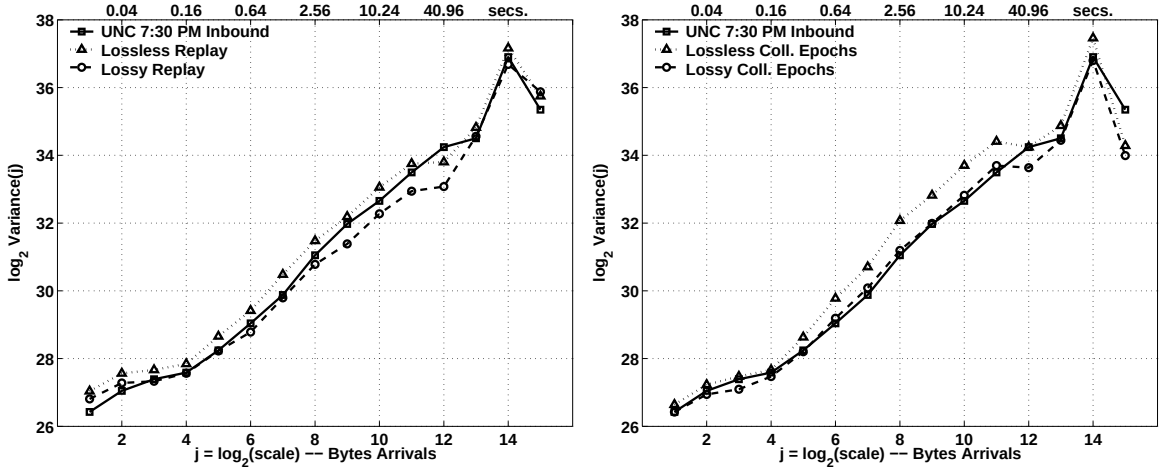

Figure 6.48: Wavelet spectra of the byte throughput time series for UNC 7:30 PM inbound and its four types of source-level trace replay.

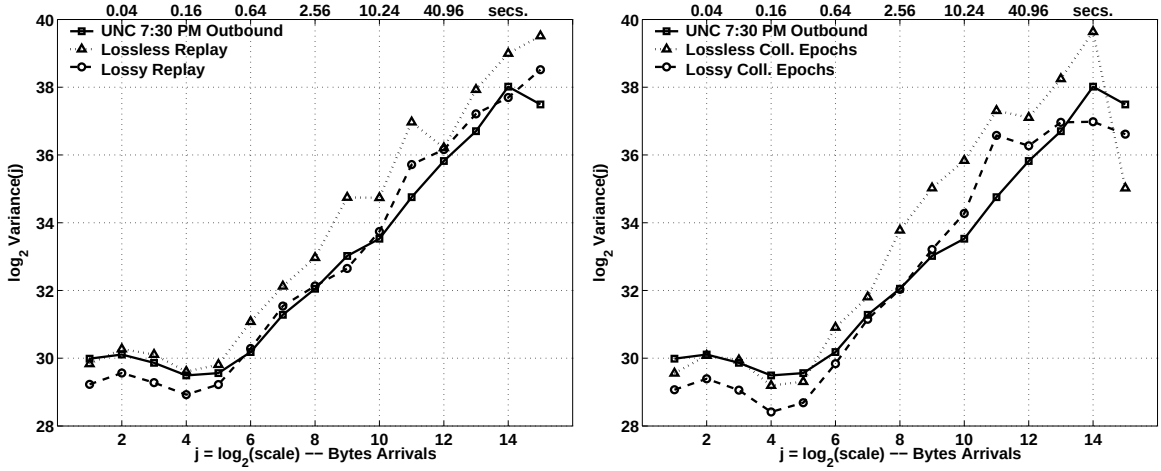

Figure 6.49: Wavelet spectra of the byte throughput time series for UNC 7:30 PM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 7:30 PM	H	C. I.	H	C. I.
Original	0.8927	[0.8520, 0.9333]	0.9424	[0.9018, 0.9830]
Lossless Replay	0.8490	[0.8083, 0.8896]	1.0191	[0.9784, 1.0597]
Lossy Replay	0.8449	[0.8043, 0.8856]	1.0044	[0.9637, 1.0450]
Lossless Coll. Epochs	0.8392	[0.7985, 0.8798]	0.9984	[0.9578, 1.0390]
Lossy Coll. Epochs	0.8655	[0.8249, 0.9062]	1.0971	[1.0564, 1.1377]

Table 6.7: Estimated Hurst parameters and their confidence intervals for the byte throughput time series of UNC 7:30 PM and its four types of source-level trace replay.

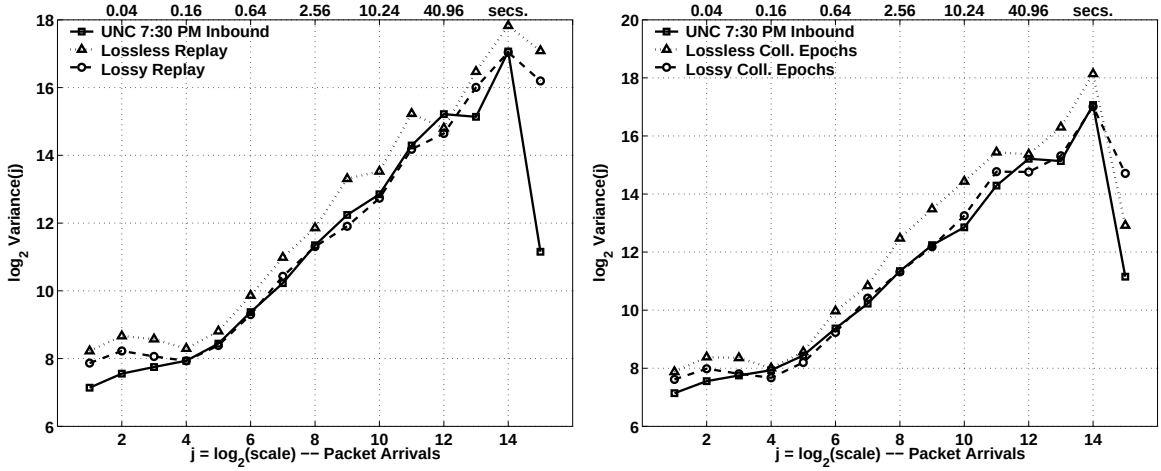

Figure 6.50: Wavelet spectra of the packet throughput time series for UNC 7:30 PM inbound and its four types of source-level trace replay.

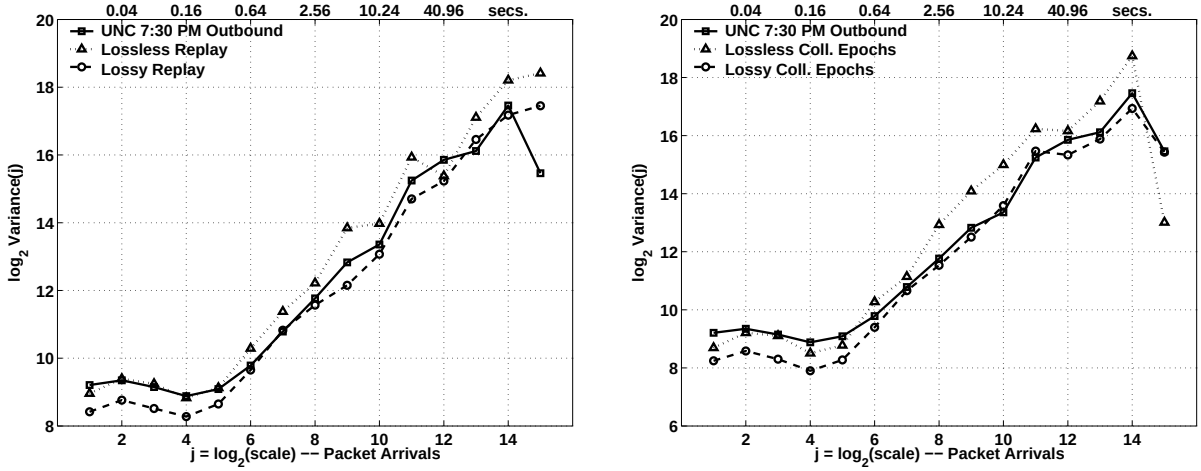

Figure 6.51: Wavelet spectra of the packet throughput time series for UNC 7:30 PM outbound and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
UNC 7:30 PM	H	C. I.	H	C. I.
Original	0.9560	[0.9116, 1.0005]	1.0061	[0.9617, 1.0506]
Lossless Replay	0.9655	[0.9210, 1.0099]	1.0043	[0.9599, 1.0488]
Lossy Replay	0.9186	[0.8742, 0.9631]	0.9524	[0.9080, 0.9969]
Lossless Coll. Epochs	0.9491	[0.9047, 0.9936]	0.9931	[0.9487, 1.0375]
Lossy Coll. Epochs	0.9967	[0.9523, 1.0411]	1.0508	[1.0064, 1.0953]

Table 6.8: Estimated Hurst parameters and their confidence intervals for the packet throughput time series of UNC 7:30 PM and its four types of source-level trace replay.

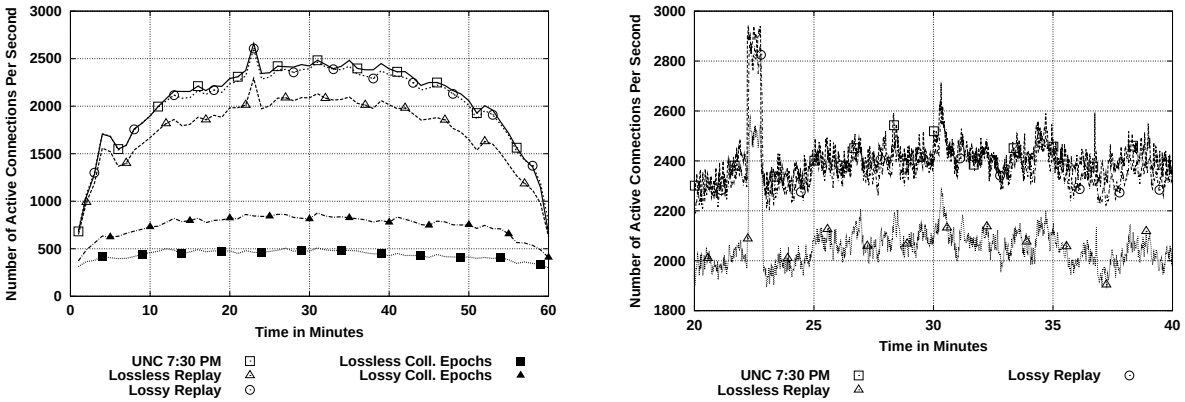

Figure 6.52: Active connection time series for UNC 7:30 PM and its four types of source-level trace replay.

in the outbound direction, lossless replays show higher energy than lossy ones, which are still above the original. In the outbound direction, lossless replays are very close to the original, while lossy ones are below it. Once again, the difficulties for matching fine scale energies mentioned in observation B.6 are present in this trace.

The estimates of the Hurst parameters are not so consistent with the results for UNC 1 AM, and are in better agreement with observation P.4. Here lossless replays approximate the original spectra closely, while lossy replays appear lower (full case) or higher (collapsed-epochs case) than the original.

6.6.5 Time Series of Active Connections

The time series of active connections shown in Figure 6.52 are in good agreement with earlier traces. Every observation listed in Section 6.4.3 is confirmed by the UNC 7:30 PM results. Unlike the two previous UNC traces, the lossy full replay is not a perfect fit of the original time series, but it still provides a very close approximation, well within the 1% bound mentioned in observation C.2. The large spike around minute 23 does not appear in the collapsed-epochs replay, providing another clear illustration of observation C.6. Note that whatever the cause of this spike, it is not due to a difference in the number of connections started, since they are identical in the four replays. The spike is necessarily explained by a set of connections with substantially longer lifespans in the full replays than in the collapsed-epochs replays.

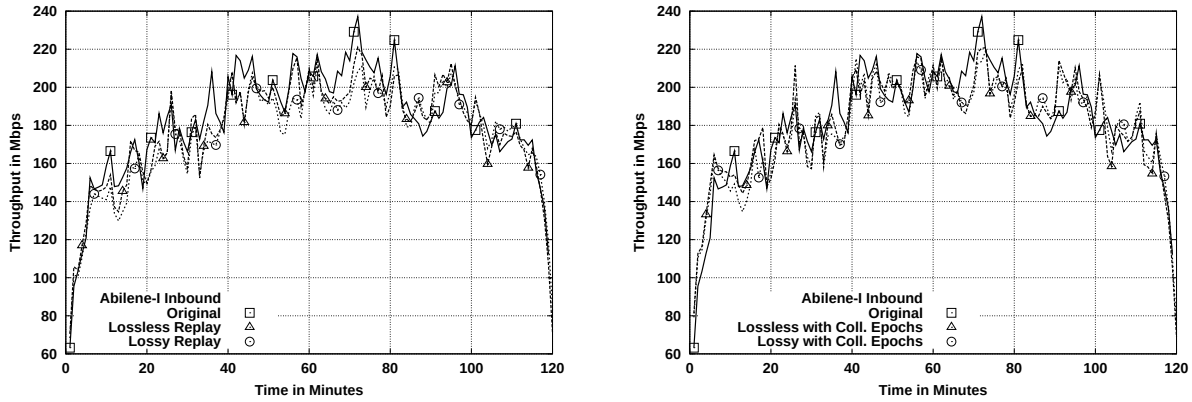

Figure 6.53: Byte throughput time series for Abilene-I Clev/Ipls and its four types of source-level trace replay.

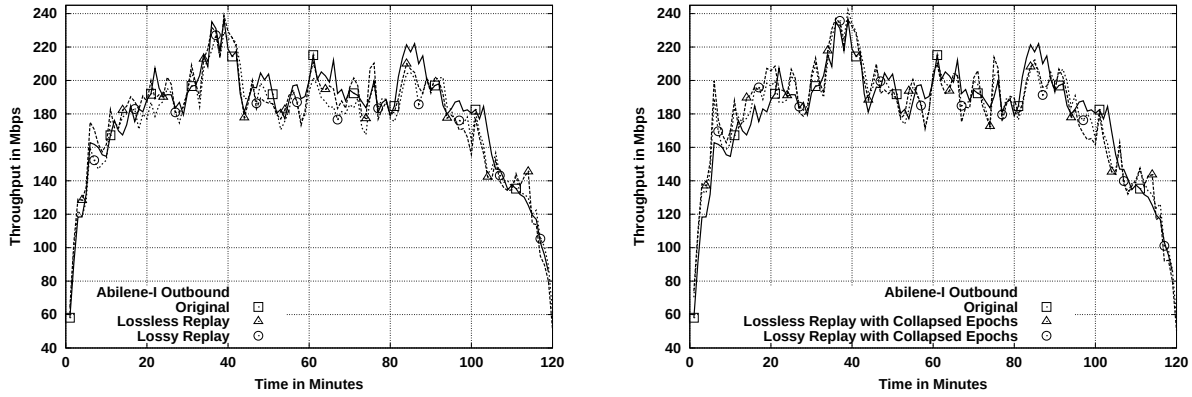

Figure 6.54: Byte throughput time series for Abilene-I Ipls/Clev and its four types of source-level trace replay.

6.7 Source-level Replay of Abilene-I

6.7.1 Time Series of Byte Throughput

The two directions of Abilene-I show the highest throughput of the five traces considered in this chapter. Combined byte throughput is often above 400 Mbps, creating the most challenging traffic generation scenario in terms of traffic volume. The excellent agreement between original and replay data shown in Figures 6.53 and 6.54 provide convincing evidence in favor of observation B.1. The replays closely track the general shape of the time series, even reproducing major changes such as the one between minutes 30 and 42. In general, we observe some spikes that appear in both original and replay time series, while others do not.

Lossless replays and collapsed-epochs replays do not seem to add any significant burstiness for this

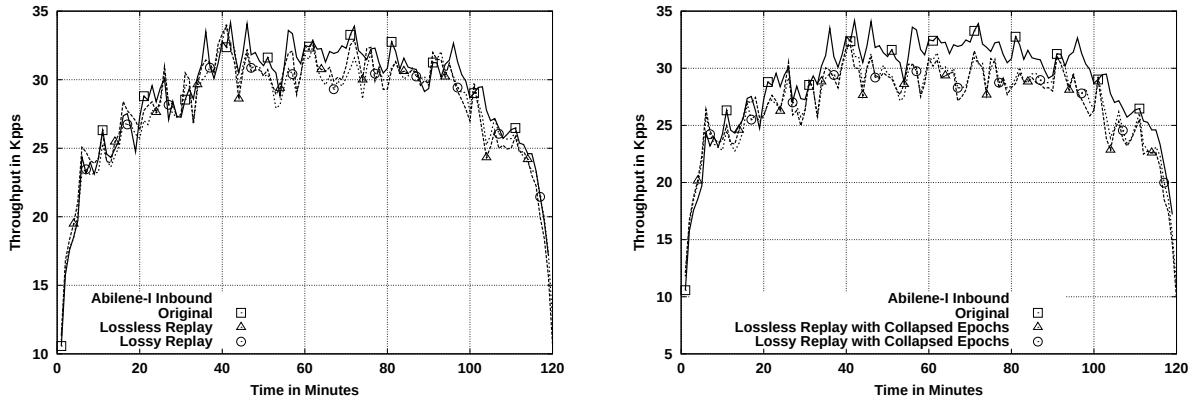

Figure 6.55: Packet throughput time series for Abilene-I Clev/Ipls and its four types of source-level trace replay.

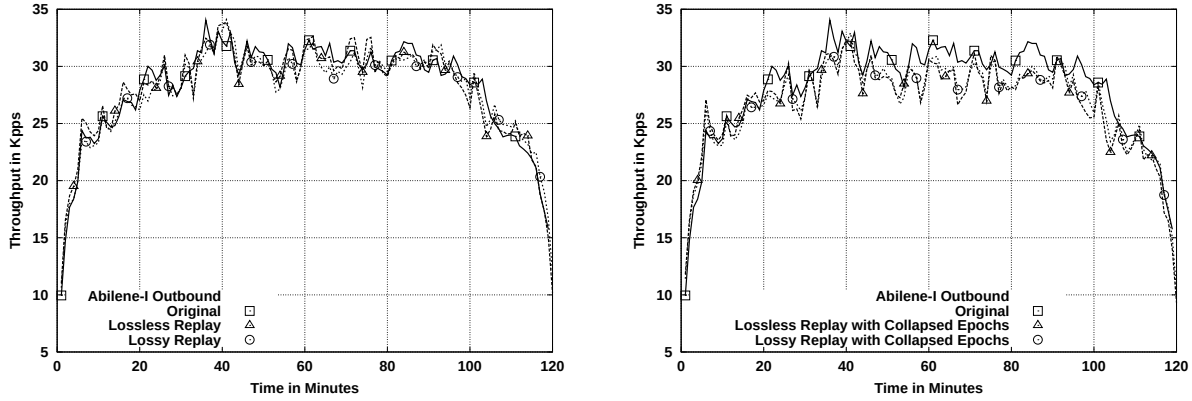

Figure 6.56: Packet throughput time series for Abilene-I Ipls/Clev and its four types of source-level trace replay.

trace, which agrees with the weak statements in observations B.2 and B.3. Note however that the high aggregate throughput could easily be hiding extra burstiness of the magnitude observed for previous traces. For example, careful examination uncovers higher throughput above the original in collapsed-epochs replay, for the spike in minute 7 and for the region between minutes 15 and 30.

6.7.2 Time Series of Packet Throughput

The time series of packet throughput in Figures 6.55 and 6.56 are consistent with observation P.1, showing an excellent match between original and full replays. Given that Abilene-I is the trace with the lowest loss level (see Section 4.1.3), this could suggest that the difficulties with the last two UNC traces were probably due to the complexity of their loss characteristics. Collapsed-epochs replays show a lower packet throughput, generally 2,000 to 3,000 packets below the original. In relative terms, the difference

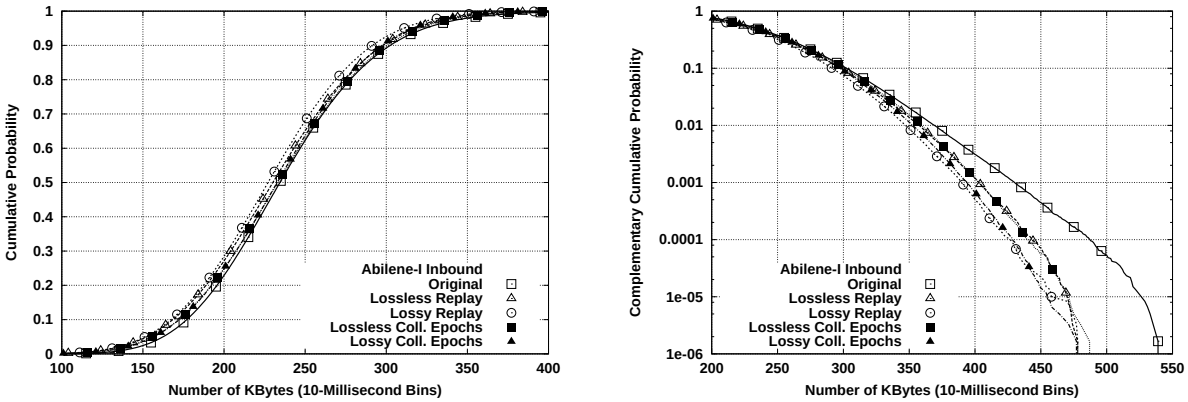

Figure 6.57: Byte throughput marginals for Abilene-I Clev/Ipls and its four types of source-level trace replay.

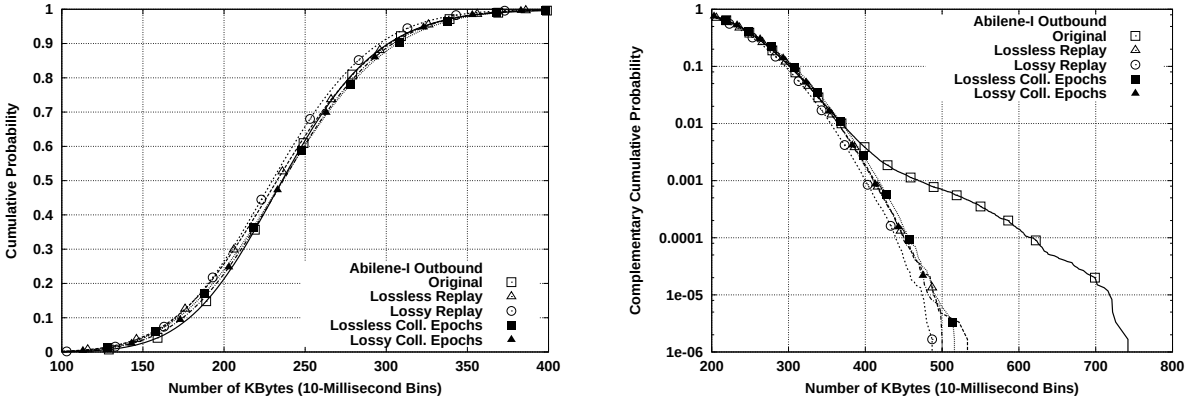

Figure 6.58: Byte throughput marginals for Abilene-I Ipls/Clev and its four types of source-level trace replay.

is between 8% and 10%, which is smaller than for previous traces. This could easily be explained by a larger percentage of bulk transfers in Abilene-I, where a single ADU carrying a single file constitutes the only payload of the TCP connection. This is for example the case in FTP-DATA connections.

6.7.3 Marginal Distributions

The marginal distributions from Abilene-I presented in Figure 6.57 and 6.58 show very similar bodies for original and replay traces. This further confirms observation B.1. Unlike previous traces, we find remarkably similar tails for all four replay traces that are consistently lighter than the original tail. The difference is specially striking in the outbound direction. One possible explanation for this intriguing result for the Abilene-I trace comes from the type of monitored link. Abilene-I is the only trace in this chapter collected in a link technology (OC-48, 2.5 Gbps) different from the one used in the replay

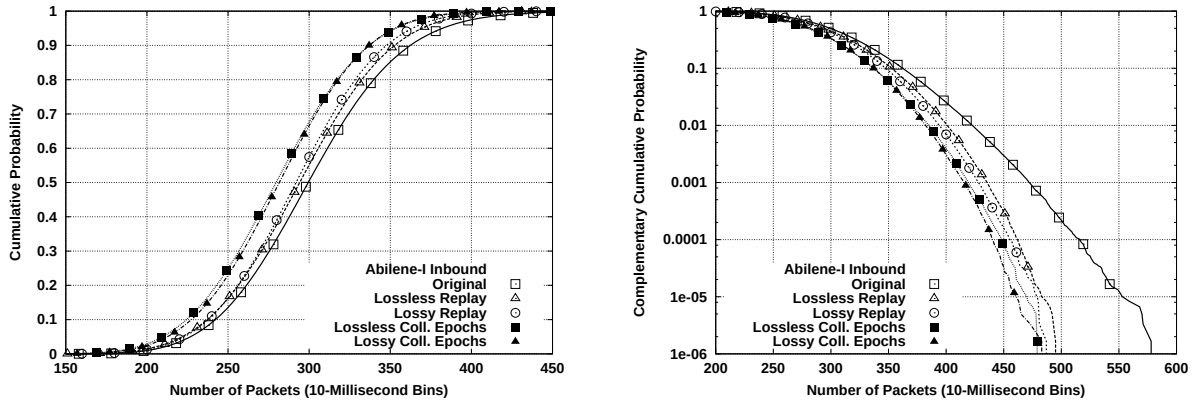

Figure 6.59: Packet throughput marginals for Abilene-I Clev/Ipls and its four types of source-level trace replay.

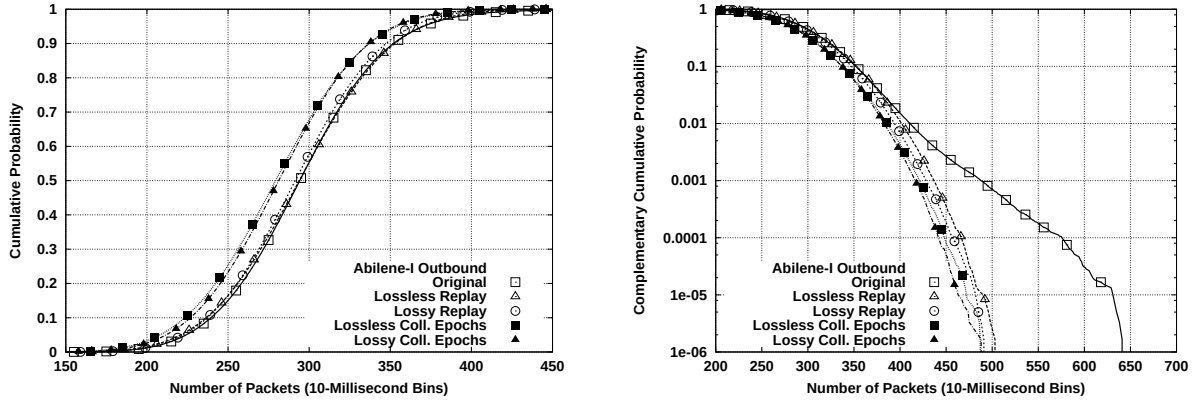

Figure 6.60: Packet throughput marginals for Abilene-I Ipls/Clev and its four types of source-level trace replay.

(Gigabit Ethernet, 1 Gbps). While the plots in Figures 6.53 and 6.54 showed no single minute with more than 500 Mbps, it is perfectly possible to have shorter (*e.g.*, 10 millisecond) intervals with far higher byte throughput. An alternative explanation is the presence of some possible limit in the forwarding capacity of our software routers, which is not far above 500 Mbps.

The bodies of the marginal distributions for packet throughput in Figures 6.59 and 6.60 are consistent with observation P.1. Collapsed-epochs replays show substantially lighter distributions, while full replays are closer to the original. The approximation in the outbound direction is remarkably good. As a consequence of the low loss in this trace, observation P.3 does not apply to Abilene-I. The impact of losses is smaller than the impact of source-level modeling. This effect is however dwarfed by the large difference between the tails of the replays and the original one. As discussed for byte throughput, differences in link technology between Abilene-I and the testbed could explain the far lighter tails in the replays.

6.7.4 Long-Range Dependence

The wavelet spectra for the inbound direction, shown in Figure 6.61, support observation B.4. However, the difference between original and full replays is substantial in the outbound direction, shown in Figure 6.62. Given the major change in slope after octave 11, it is difficult to draw any conclusions from this finding. Regarding observation B.5, we do observe worse approximations by the collapsed-epochs replays, which exhibit substantially deeper ditches around octave 5 (notice the lower smallest value in the y-axis of the outbound plot). In any case, the replays do not closely track the fine scale shape of the spectra, which is in agreement with observation B.6.

Hurst parameters, shown in Table 6.9, are remarkably high for this trace. All of them are above 1, suggesting significant non-stationarity, which is clearly preserved in the replays. The estimate for the lossy full replay is the closest one for both directions. Together with the wavelet spectra, this supports observation B.7.

As for byte throughput, the wavelet spectra of the packet throughput for Abilene-I shown in Figures 6.63 and 6.64 are not comparable to previous cases, and inconsistent with observation P.4. The difference between the replays and the original follows the same pattern in all of the cases, with a large ditch in octave 5. This ditch is much more pronounced for the collapsed-epochs replays. In any case, the result is a poor match between the original spectra and the replays, both at fine scales and at the the scaling region. Estimated Hurst parameters for the replays are lower for the inbound replays and higher for the outbound ones, and mostly outside confidence intervals. Incorporating losses had a minimal impact on the wavelet spectra of the Abilene-I replays, resulting only in a small decrease of the slope in the scaling region. This decrease translated into a slightly smaller Hurst parameters estimates for the lossy replays.

6.7.5 Time Series of Active Connections

Unlike the results for previous traces, Figure 6.65 shows a substantial difference between the lossy full replay and the original time series. This weakens observations C.1 and C.2 from Section 6.4.3, being the replay around 15% below the original. The rest of the observations clearly hold. The relative magnitude of the gap between the full replays and the collapsed-epochs ones is largest for Abilene-I. The reason is unclear, especially given the excellent approximations for the other traces. It is hard to imagine a larger fraction of bandwidth-constrained connections in this trace, and round-trip time estimation should be as accurate as for the other traces. We are more inclined to think that the mix of applications in

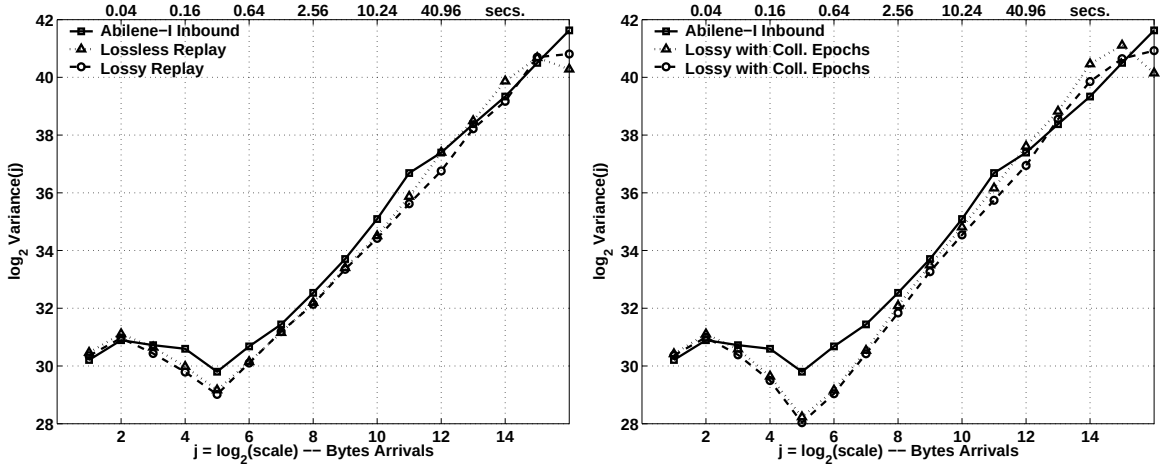

Figure 6.61: Wavelet spectra of the byte throughput time series for Abilene-I Clev/Ipls and its four types of source-level trace replay.

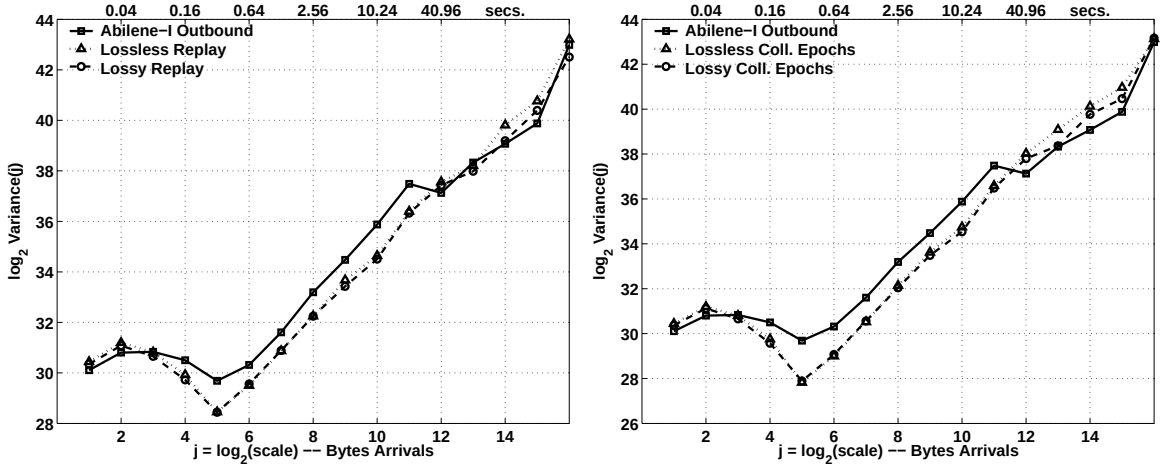

Figure 6.62: Wavelet spectra of the byte throughput time series for Abilene-I Ipls/Clev and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
Abilene-I	H	C. I.	H	C. I.
Original	1.0597	[1.0320, 1.0874]	1.0604	[1.0327, 1.0881]
Lossless Replay	1.1170	[1.0893, 1.1447]	1.1356	[1.1079, 1.1633]
Lossy Replay	1.0814	[1.0537, 1.1091]	1.1079	[1.0802, 1.1356]
Lossless Coll. Epochs	1.1824	[1.1573, 1.2075]	1.2111	[1.1860, 1.2362]
Lossy Coll. Epochs	1.1580	[1.1329, 1.1831]	1.1874	[1.1623, 1.2125]

Table 6.9: Estimated Hurst parameters and their confidence intervals for the byte throughput time series of Abilene-I and its four types of source-level trace replay.

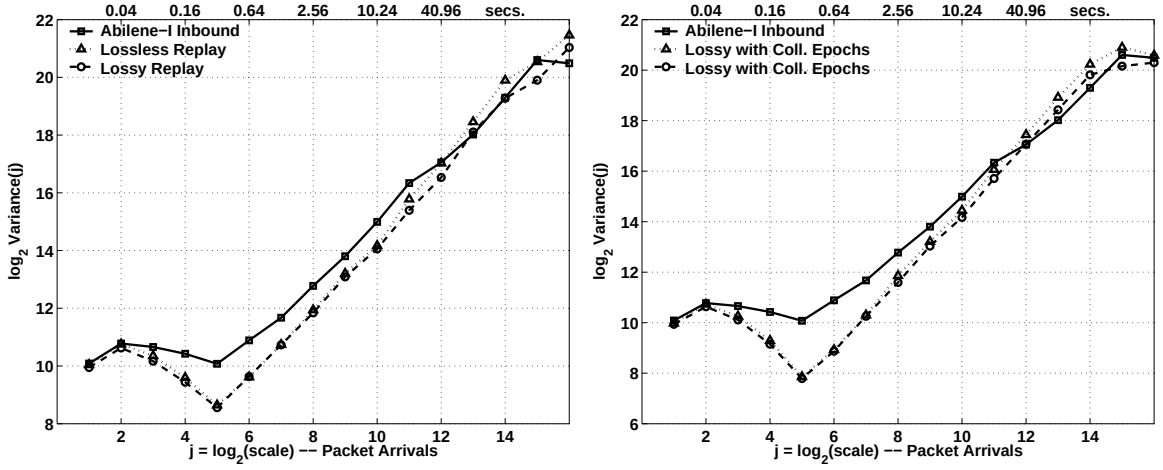

Figure 6.63: Wavelet spectra of the packet throughput time series for Abilene-I Clev/Ipls and its four types of source-level trace replay.

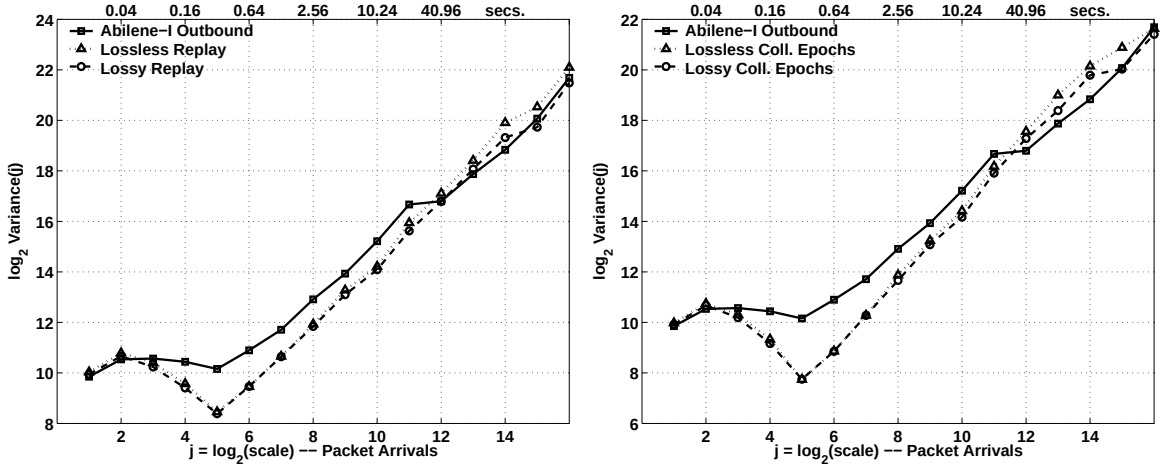

Figure 6.64: Wavelet spectra of the packet throughput time series for Abilene-I Ipls/Clev and its four types of source-level trace replay.

Trace	**Inbound**		**Outbound**	
Abilene-I	H	C. I.	H	C. I.
Original	1.1326	[1.1075, 1.1577]	1.0996	[1.0745, 1.1247]
Lossless Replay	1.1191	[1.0941, 1.1442]	1.1443	[1.1192, 1.1694]
Lossy Replay	1.0849	[1.0598, 1.1100]	1.1232	[1.0981, 1.1483]
Lossless Coll. Epochs	1.1841	[1.1563, 1.2118]	1.1923	[1.1646, 1.2200]
Lossy Coll. Epochs	1.1757	[1.1480, 1.2034]	1.1850	[1.1573, 1.2127]

Table 6.10: Estimated Hurst parameters and their confidence intervals for the packet throughput time series of Abilene-I and its four types of source-level trace replay.

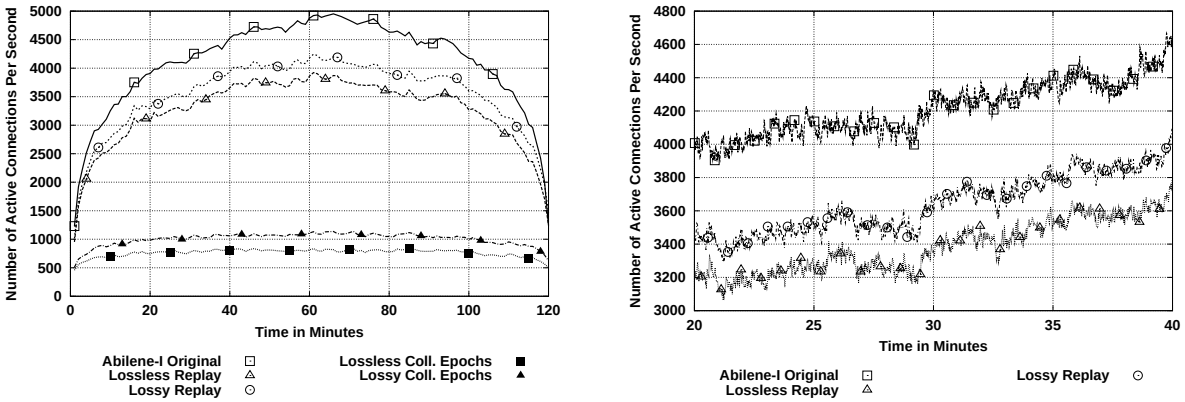

Figure 6.65: Active connection time series for Abilene-I and its four types of source-level trace replay.

Abilene-I includes a substantial number of (probably long) connections whose driving application is not well-described by our source-level model.

6.8 Summary

The results in this chapter demonstrated that source-level trace replay can closely approximate the characteristics of real traffic traces. We have also shown that full source-level replays are closer or far closer to original traces than collapsed-epochs replays for several metrics. In particular, the largest difference is observed for the time series of packet throughput, the body of the packet throughput marginal and the time series of active connections. Byte throughput is similar for full and collapsed-epochs replays. The latter exhibits somewhat more bursty time series, but the bodies of the marginals do not change significantly.

The rest of the metrics cannot be clearly interpreted, since losses have a much more significant impact on them than the use of full or collapsed-epochs replays. Lossy full replays are clearly better than lossy collapsed-epochs replays in terms of wavelet spectra, estimated Hurst parameters and tails of the marginals for some traces, but this is not consistent for the five traces. Our analysis clearly demonstrated the need to carefully consider the impact of losses on evaluating the quality of synthetic traffic. Without our direct comparison of lossless and lossy replays, the results for certain metrics could have mislead our conclusions regarding source-level modeling. In contrast, other metrics are less affected by the loss model. This is the case for the time series of packet throughput, the body of the packet throughput marginal and the time series of active connections, where full replays are clearly better

approximations than collapsed-epochs replays.