

Network Working Group
Internet Draft
Expiration Date: August 2002

A. Corlett
CQOS Inc., Irvine, CA
D.I. Pullin
California Institute of Technology
S. Sargood
Nortel Networks, UK

Statistics of One-Way Internet Packet Delays
<draft-Corlett-Statistics-of-packet-delays-00.txt>

1 Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as “work in progress”.

The list of current Internet-Drafts can be accessed at

<http://www.ietf.org/lid-abstarcts.txt>

The list of Internet-Draft shadow directories can be accessed at

<http://www.ietf.org/shadow.html>

This memo provides information for the internet community. This memo does not specify an internet standard of any kind. Distribution of this memo is unlimited.

2 Abstract

The statistical properties of packet delays for transmission across the Internet are investigated, based on analysis of three datasets obtained using CQOS cNodes, each measured over several days of continuous transmission. Two of these sets comprise high and low bandwidth measurement data for vectors (defined here as a cNode to cNode link) from CQOS headquarters to the Irvine Data Center, while the third is a low-bandwidth dataset obtained from a CQOS-Irvine to London vector. The principal results of this study may be summarized as follows. First, the two local datasets are characterized by quiet periods, where the 300 second mean delay shows little variation, separated by periods of severe volatility in the mean, standard deviation, minimum and maximum delays. In contrast, the international

dataset showed only small variations in these quantities over a four-day measurement period. Second, during the quiescent periods, the probability density function of packet delays is well approximated by a shifted exponential distribution, for all three datasets. This suggests that packet delays tend to be concentrated near the minimum delay. Third, the packet correlation time, defined by the first zero-crossing of the delay autocorrelation function, exhibited a long-tailed distribution, with average correlation of order a few to ten seconds. Fourth, the power spectra of the time series for delay for two of the datasets showed no salient features corresponding to periodic delay variation at any time period smaller than the known daily characteristic time scales for packet delay.

3 Introduction

Knowledge of the detailed statistical properties of one-way packet delay, packet loss, delay variation and other metrics is of paramount importance for understanding the general properties of Internet transmission. This influences the design and construction of both measurement algorithms, the aggregate quality-of-service parameters, and the estimation of statistical errors incurred in measurement strategy. Whilst there have been several studies [e.g. Refs 1-2] that have considered measurement strategies for data collection in the assessment of Internet metrics, we know of only a few quantitative investigations [3-4] of the statistical properties of packet delays across Internet links. Of particular interest for the present work is Mukherjee's [4] analysis of round-trip delay, packet loss and packet orderliness for packet transmission in the range 1 – 60 Hertz. Mukherjee found that, for the low-frequency component of internet transmission, the round-trip delay probability density function (pdf) was well modeled by a shifted Gamma distribution, with shape and scale parameters which varied with load and network segment. He also noted the presence of significant slow oscillation components in the smoothed network delays. Some, but not all of the present findings are in broad agreement with these results.

The layout of this document is as follows. In §4 we briefly describe the present experiments and summarize the parameters of the resulting datasets that comprise the present database. The data analysis, consisting of delay time series, probability density functions and the analysis of the delay autocorrelation function within measurement records is presented in §5. In §6 we consider a metric that measures the degree to which packet delays within individual measurement records tend to be clustered toward the minimum delay for the record. Finally, §7 discuss features of the delay power spectra for two of the delay time series comprising the present database.

4 The data

Datasets were initially obtained on four vectors. Of these, three, which we refer to as datasets # 1, # 2 and # 3 were measured with vectors from CQOS headquarters to the CQOS data

Table 1: Parameters defining datasets # 1, # 3 and # 4. Bandwidth = 1.5×10^6 bps, packet length = 576 bytes, utilization $\rho = 1.0$.

Dataset	Number of records	Measurement period (days)	Pkts per 300 secs (M).	Occupancy fraction ν	Vector
# 1	621	2.2	611	0.0063	Local
# 3	2033	7.1	9556	0.092	Local
# 4	1017	3.5	730	0.0073	London

center in Irvine, across a T1 Internet connection. Transmission was over the Internet, and not on a dedicated link, so that although the two cNodes were physically close, they were not close logically. Three different test-packet bandwidths were used for these measurements. Of these, that of dataset # 2 was later deemed to be too large to meet the specifications of low (1%) total measurement bandwidth utilization, and so only datasets # 1, and # 3 will be discussed presently. We refer to these as low and high-bandwidth *local* datasets, respectively. To complement the physically local data, dataset # 4 was obtained from a vector defined by two cNodes located at CQOS headquarters, and in London, England, respectively. We call this the *London* (low bandwidth) dataset.

For each vector, over sequential 300 second measurement periods, M packets with a fixed length of 576 bytes were dispatched using periodic streaming in which the time between each packet dispatch was fixed at approximately $300/M$ seconds. A separate file, or measurement record, was created for each 300 second measurement period. On each measurement record, the time at which transmission began for that record was recorded, and during the measurement period, the GPS synchronized send and receive times of each packet were measured and recorded. The total number of sent (M) and received ($\leq M$) packets during the period was also recorded to enable later calculation of loss. This procedure produced a large volume of records for each data set, from which detailed statistics of packet delay and loss could be post-processed.

Table 1 shows the main parameters of each data set. The nominal occupancy fraction ν of the T1 line is calculated from

$$\nu = \frac{8 P L}{\rho C}$$

where $L = 576$ bytes is the packet length, $P = M/300$ is the dispatch rate, $C = 1.5 \times 10^6$ bps is the T1 link bandwidth and ρ is the utilization. In Table 1 we have used $\rho = 1$. Note that ν is modest for datasets # 1 and # 4 but is sizable for # 3, even with $\rho = 1.0$.

Table 2: Mean delay, delay standard deviation, and mean correlation time averaged over all records for each of three datasets. Units are ms.

	Dataset # 1	Dataset # 3	Dataset # 4
mean delay	9.622	16.024	108.232
standard deviation	6.216	5.400	3.083
correlation time	7.485	9.720	1.825

5 Results

5.1 Delay time series

Although sufficient data was recorded to study the statistics of one-way packet delay, packet loss, and the correlations between these quantities, the present investigation will henceforth be restricted to a study of one-way packet delay. We denote the one-way packet delay by d , and will take the view that $d(t)$ for live traffic across the internet between two endpoints, represents a random stochastic process [5] in time. Our task will be to attempt to infer some of the low-order statistical properties of $d(t)$ using the present datasets. In what follows, we present a variety of statistical measures for each dataset. In order to maintain confidence in our analysis, almost all statistics presented presently were calculated independently by the second and third author. Calculated values were then compared and recalculated until agreement was achieved. This method gives us a high degree of confidence in the accuracy of the statistical analysis presented herein.

One-way delay time series measured during typical 300 second measurement periods, one from each of the three datasets, are shown in Figure 1. A ‘slice’ from each of these time-series, from within a window $50 \leq t \leq 120$, where t is the measured time in secs since the beginning of the record, are displayed in Figure 2. Over the full 300 second periods, the two local datasets # 1 and # 3, show qualitatively similar behaviour, with periods of very small delay variation punctuated by short bursts of longer delay. The durations of these delay bursts can be seen from the ‘slice’ graphs to be typically of order one to three packets. Note that for dataset #1, the interpacket dispatch time is about $300/611 \approx 0.491$ seconds, while that for dataset # 3 is $300/9556 \approx 0.0314$ seconds. This behaviour has been observed in all previous studies to date, and is the tail component in the Gamma distribution for delay alluded to earlier. However the exact cause for these spikes is difficult to ascertain without knowledge of the ISP networks and the link utilisations over which the packets are routed, and the cause could be simply that alternative routes were taken by these packets, or there was intermittent congestion in an access router. Any network routing update which remained in force would have created a step function in the delay time-series, so these infrequent delay spikes are more likely to be regarded as outliers in the overall dataset, however they remain statistically important and should not be disregarded. Some burstiness is also evident in the time series plots for the London dataset # 4 but the corresponding relative variation in

delay is substantially smaller for this dataset than for the local datasets. This is probably due to the effect of smoothing produced by a relatively large number of router hops over the international route.

Quantities of particular interest are the mean delay $\langle d \rangle$ and the standard deviation of delay s

$$\langle d \rangle = \frac{1}{M_r} \sum_{i=1}^{M_r} d_i, \quad s^2 = \frac{1}{M_r - 1} \sum_{i=1}^{M_r} (d_i - \langle d \rangle)^2, \quad (1)$$

where M_r is the number of received packets, the minimum d_{min} delay and the maximum delay d_{max} , each within a 300 second measurement period. A principal aim of the present study is to quantify the statistics of $\langle d \rangle$ for the three datasets. Figure 3 shows $\langle d \rangle$, d_{min} and d_{max} versus time for the three datasets. The figures consist of plots of these quantities versus the ‘record time’ - that time at which the record begun - for the record sequence comprising each data set. The record time is measured in hours from midnight on the first day on which the measurements started, for that dataset. It is evident that datasets # 1 and # 3 show periods of nearly constant $\langle d \rangle$ separated by periods where there is substantial volatility, with large record to record variations in both $\langle d \rangle$ and d_{max} . There is clearly a diurnal pattern in this cycle, except for a weekend quiet period evident in the data for dataset # 2. Dataset # 3, however, shows much more stable behavior in $\langle d \rangle$, although a small daily variation can nevertheless be observed. Note that for this dataset, the maximum values of delay can be as high as $d_{max} = 440ms$ while the mean of the maximum values is $\langle d_{max} \rangle = 166ms$, which suggests not all packets in a record take the same route or else there is buffering of some appreciable time in access routers/switches, although probably not core routers. If there were major route changes in the network(s) we would again expect step functions in the minimum delay values for periods of time. This is not evident.

The variation of the standard deviation s with record time is shown in Figure 4. Again datasets # 1 and # 3 show volatility on a diurnal cycle, with periods in which s is of the same magnitude as $\langle d \rangle$. For dataset # 4, s is always at least ten times smaller than $\langle d \rangle$. If each record is viewed as a statistical sample of size M drawn from a parent population consisting of laive traffic, then application of the Central Limit Theorem CLT gives estimates of upper and lower bounds on the population mean delay as [5,6]

$$\langle d \rangle_{lower} = \langle d \rangle - \lambda \frac{s}{\sqrt{M}}, \quad \langle d \rangle_{upper} = \langle d \rangle + \lambda \frac{s}{\sqrt{M}}, \quad (2)$$

where λ is a factor that depends on the desired confidence level. For a 95% confidence level, $\lambda = 1.96$ and for 99% confidence level, $\lambda = 2.577$. Figure 5 shows $\langle d \rangle_{lower}$, $\langle d \rangle$ and $\langle d \rangle_{upper}$ for the three data sets using values of M from Table 1, while Figure 6 shows a close-up window of these quantities for two datasets. It is clear that the ‘relative statistical error’, which we define as $\lambda s / (\langle d \rangle \sqrt{M})$ is small for all datasets, owing to a combination of fairly small s and sizable M , particularly for dataset # 3. During quiet periods this relative error is very small.

5.2 Delay pdfs

An ongoing topic of research concerns the precise form of the pdfs of delay over periods of order the 300 second measurement period. Presently, we show in Figure 7, pdfs of delay for typical 300 second records. For datasets # 1 and # 3, these were taken from records during one of the quiet periods of Internet transmission. All three pdfs show very sharp peaked distributions with most values for delay clustered within about 10% of both $\langle d \rangle$ and d_{min} . All show very long and extremely thin tails consistent with the form of the time series discussed earlier. The highly impulse-like shape supports the earlier studies that the mode of the delay distributions here is a more relevant statistic than the mean. These pdfs will be analysed in more detail in future work including methods for efficiently computing the mode. We presently comment that the delay pdfs are well modelled with a shifted exponential distribution, which can be viewed as a special case of a shifted Gamma distribution (see appendix) with parameter $\alpha \approx 1$. This too is consistent with the observations of Mukherjee [4] for round-trip delay pdfs.

5.3 Delay correlations

Correlations in the packet-to-packet delay are of great interest as a measure of time intervals over which individual packet delay can be considered to be independent. The delay correlation time was measured as follows. For a particular 300 record, the delay autocorrelation function $C(m)$ can be defined as

$$C(m) = \frac{\sum_{i=1}^{M_r-m} (d_i - \langle d \rangle) (d_{i+m} - \langle d \rangle)}{\sum_{i=1}^{M_r} (d_i - \langle d \rangle)^2}, \quad (3)$$

where m is the packet separation number. This can be expressed as $C(T)$ where $T = 300m/M_r$, where T is the time separation and $300/M_r$ the (constant) packet dispatch interval. Note that the denominator in (3) is proportional to s^2 (equation 1). $C(T)$ is such that $C(T) = C(-T)$, and $C(T=0) = 1$. Figure 8 shows $C(T)$ for one typical 300 second record, for each of the three datasets. We define the correlation time T_c as the ‘width’ of the $C(T)$ curve. There are various ways to define this width. Presently we use $T_c = 2 \times$ (value of T for the first zero crossing of $C(T)$). In Figure 8, the form of $C(T)$ for the records chosen from datasets # 1 and #4 show little structure, with the first zero crossing occurring at quite small values of T . For these examples $T_c \approx 300/M_r$. This means that the first zero crossing occurs at near $m = 1$, so that the packet delays for these records are essentially uncorrelated. For dataset # 3, the top right-hand graph of Figure 8 shows a very sharp drop from $C(0) = 1$ to about $C(0+) \approx 0.3$, followed by a gradual decline to the first zero crossing at $T_c \approx 30$ seconds. It is clear that there is substantial delay correlation for this record, but that $T_c \ll 300$.

Figure 9 shows plots of both T_c and $\langle d \rangle$ against record time for each of the three datasets. Note that T_c is scaled differently for datasets # 1 and # 3, where we have plotted $T_c/10$, than for dataset # 4, where we have plotted $10 T_c$. For datasets # 1 and # 3, it

appears that large T_c is itself correlated with large variations in the record-to-record values of $\langle d \rangle$.

The data discussed above is summarized in Figure 10, where we show normalized pdfs of $\langle d \rangle$ (top left), s (top right) and T_c (bottom). Note that in these plots, we have not attempted to ‘collapse’ these curves by scaling the abscissa against the mean values of the respective quantities over all records (see Table 2), but this could be done. It will be of relevance to the analysis of the next section, that for the vast majority of records for all three datasets, $T_c \ll 300$ seconds.

6 Minimum delay percentage window

Mean packet delay over 300 second measurement period represents one of many possible delay metrics. Other metrics of interest include d_{min} , d_{max} , delay standard deviation and other moments of the delay probability density function (pdf). It has already been noted that the delay pdfs appear to be well modeled by shifted Gamma distributions, with $\alpha \approx 1$. The shifted exponential distribution has interesting features, notably that the minimum and the mode (most probably delay) are identical. This suggests that, during a single record, individual packet delay will tend to be clustered near d_{min} . In work we will develop an algorithm for estimating the measurement-record delay mode. Presently we describe a simple statistic which quantifies the extend to which measurement-record delay measurements concentrate near d_{min} . Within a single record, we define the $P\%$ *Minimum Delay Window (MDW)* by the delay interval

$$d_{min} \leq d \leq (1 + \frac{P}{100})d_{min}. \quad (4)$$

Equation (4) defines a range of delay values bounded below by d_{min} and above by $(1 + 0.01P)d_{min}$. The 10% MDW is thus defined by $d_{min} \leq d \leq 1.1d_{min}$, a range of delay values of width 10% of d_{min} . Figure 11 shows the fraction of packet delays, averaged over all records, that fall within the $P\%$ MDW, for datasets # 1 and #4. Both curves show steep increase for $P \approx 1 - 2\%$. Averaged over all records, in excess of 90% of packet delay times fall within the 10% MDW for dataset #1, while 98% of delays fall in this same MDW for dataset # 4. Corresponding probability density functions for various percentage MDWs are shown in figure 12. Each curve shows the pdf of the fraction of records within each dataset for each of several values of the $P\%$ MDW, for datasets # 1 and #4.

In figures 13 and 14, we show the fraction of packets with delay within two $P\%$ MDWs over all consecutive 300 second measurement records comprising both datasets #1 and #4. Note that the chosen values of P differ for each dataset, reflecting the rather different characteristics of the time series for the two cases. For comparison, we also show the corresponding variation of the average, minimum and maximum delay for each dataset. Datasets # 1 and # 2 show somewhat different behavior. For both datasets, when the average delay shows little variation from record to record, a fraction near unity (*i.e* nearly 100%) of packets are transmitted with delay within the 10% MDW. This is true for almost all records of dataset

4. Periods of volatility in dataset # 1, where the average and maximum delay within a measurement record show large fluctuations, clearly show a correlation between increasing average/maximum delay with decreasing fraction of packets with delay in the 10% MDW. These periods show little variation in the minimum delay. Even during the volatile periods, however, the fraction of packets with delays that are within the 10% MDW only rarely falls below 0.8. These results suggest that, overall, the network transmits most packets at close to the minimum delay, only failing to do this during periods of heavy demand on local packet routes. This suggests that the percentage of packets transmitted within a $P\%$ MDW (say $P = 10\%$) may provide an interesting diagnostic metric for traffic engineers. This can be done with two sweeps of the delay measurements for each measurement record, the first to determine d_{min} and the second to determine the fraction of packets with delay within a given $P\%$ MDW. This should be straightforward.

7 Power spectra of the delay time series

The sequence of delay measurements, tabulated against the clock time at which packets are dispatched, forms a long time series. It is of interest to study the frequency content of this series. Within each 300 second measurement record, packets are dispatched at time intervals $\Delta t = 288/M$, where M is the number of packets dispatched. For dataset # 1, $\Delta t \approx 0.471389$ seconds, while for dataset # 4, $\Delta t \approx 0.400898$ seconds. Because the measurement period is 300 seconds, there exists a gap of about 12 seconds from the end of one measurement period to the beginning of the next. This is a period where measurement and management housekeeping by the system is performed. In order to obtain a continuous time series over approximately the whole of the measurement periods for both datasets, this gap was ignored for the purposes of computing the power spectrum of the two datasets. Thus, in practice, the records within each dataset were simply concatenated to form a single time series with a total $N = 380030$ entries for dataset # 1 and $N = 729844$ for dataset # 4. In what follows this will lead to a systematic error in the power content of frequencies smaller than $2\pi/300 \approx 0.021 \text{ secs}^{-1}$ and to possible spurious periodic behavior on these time scales. It will be seen that the first of these effects is unimportant and the second is essentially non-existent.

With this approximation, the Fourier series for the delay times series was computed as

$$d(t) = \sum_{k=-N/2}^{N/2-1} \hat{d}_k e^{i\omega_k t}, \quad \omega_k = \frac{2\pi k}{T}, \quad d_{-k} = d_k^*, \quad (5)$$

where \hat{d}_k is the amplitude of mode k with frequency ω_k , t is time, $T = N\Delta t$ is the total length of the (concatenated) dataset and "*" denotes the complex conjugate.

The power spectrum of the delay time series is a plot of $|d_k|^2 = d_k d_k^*$ versus ω_k . If the time series contains dominant periodic behavior, these will appear as local peaks in the power spectra. Figure 15 shows the power spectra of the delay time series for datasets # 1 and #

4. These both show similar features. There is a maximum in the power spectra at hourly to diurnal frequencies, in the range $\omega_k \approx 10^{-5} - 10^{-4} \text{ secs}^{-1}$, followed by noise at higher frequencies. For dataset # 1, there is an indication of ω^{-1} rolloff in the power spectrum, which may indicate a fractal power-law structure of the delay time series. It is notable that there are no features of either spectra that could correspond to inherent periodic behavior on either (i) periods corresponding to the measurement record period of 300 seconds, or (ii) periods corresponding to multiples of the packet dispatch period Δt . This suggests that periodic packet dispatch is adequate for test packet transmission.

8 Concluding remarks

The experiments conducted in this work have focussed on the statistics of Internet delay using three datasets, one of which was over a long-distance international route. The measurement data for packet delays show both quiet and volatile periods over duration of several days and bear out diurnal cycles, and key results regarding Internet behaviour obtained in previous work. A proposal that the probability density function of the delay distribution can be approximated by a shifted exponential distribution is confirmed here, and methods to determine the key scaling and shape paramters are discussed elsewhere. However this observation in turn supports the assertion that computation of the mode is essential as one of the key statistics for assessing Internet delay, due to the impulsive appearance of the delay distribution which extends into a very long tail. This was found to be a characteristic feature of both the local and long-distance data sets, and is indicative of the underlying dynamics of connectionless IP networks.

The data was also found to be subject to varying degrees of non-stationarity and based on the definition for wide-sense stationarity it may be initially concluded that data sets experienced periods where the underlying live traffic was weakly stationary (invariant mean). Other time-periods, however, indicated significant variation with the mean and consequently non-stationarity.

Finally, we find that an interesting metric for packet transmission worthy of further study consists of the percentage of packets within a given measurement record that are transmitted with delay that lies within the 10% minimum delay window. This was found to be near unity for the international dataset and for the local dataset near quiescent periods. An interesting area for further study would be to compare this behavior with that of the record-by record most probable delay (mode).

9 Security Considerations

This document is solely for the purpose of reporting results of a study of empirical data to determine statistical properties of packet delays and describes neither a protocol nor a

protocol's implementation. Therefore, there are no security considerations associated with this document.

10 References

- [1] Chou, P.A and Miao, Z. *Rate-distortion optimized streaming of packet media*. Microsoft Research Corporation, Feb. 2001. <http://www.research.microsoft.com>
- [2] Claffy, K.C., Polyzos, G.C. and Braun, H-W., *Application of sampling methodologies to network traffic characterization*. SIGCOMM 1993:194-203.
- [3] Acharya, A. and Saltz, J., A study of Internet Round-Trip Delay, University of Maryland Report, CS-TR 3736
- [4] Mukherjee, A., *On the dynamics and significance of low frequency components of internet load*. Comp and Inf. Sci. Dept. Tech. Rept. No. MS-CIS-92/83/DSL-12), University of Pennsylvania.
- [5] Papoupolis, A. *Probability, random variables and stochastic processes*. New York. McGraw-Hill. 1984
- [6] Pullin, D.I., Corlett, A. and Mandeville, R. *Statistical accuracy in network quality-of-service measurement*. *CQOS Statistical Papers*, 2001.

11 Authors Addresses

Andrew Corlett
CQOS Inc.,
7 Technology
Irvine, CA 92618

Dale Pullin
Mailstop 105-50
California Institute of Technology
1200 East California Blvd.
Pasadena CA 91125

Stephen Sargood
Nortel Networks UK
Maidenhead Office Park
Westacott Way
Maidenhead, SL63QH

Figures

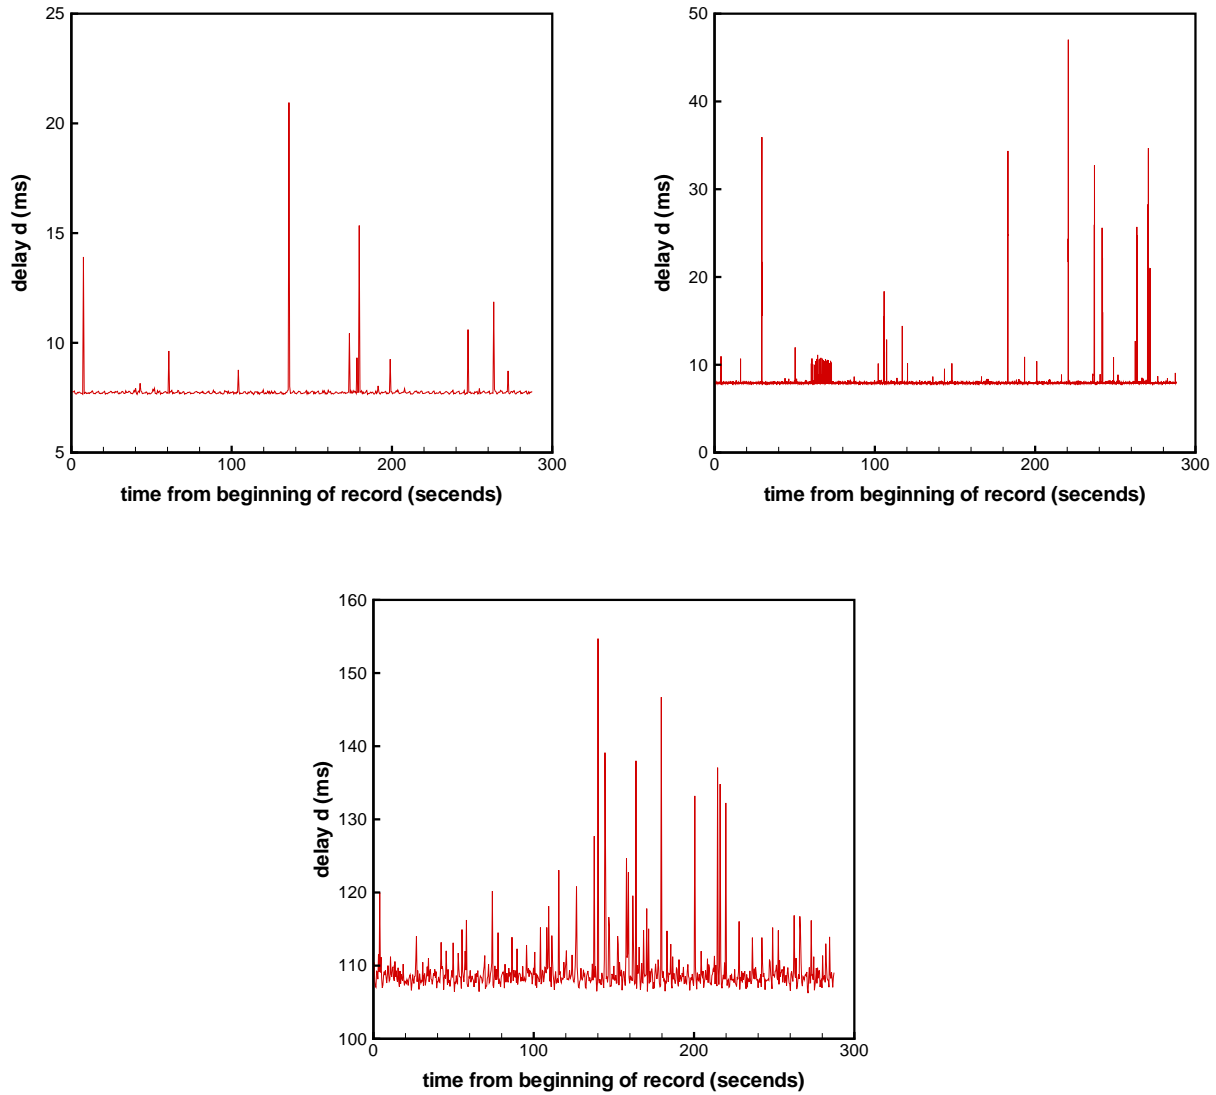


Figure 1: Typical time series of delay over a 300 second measurement record. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

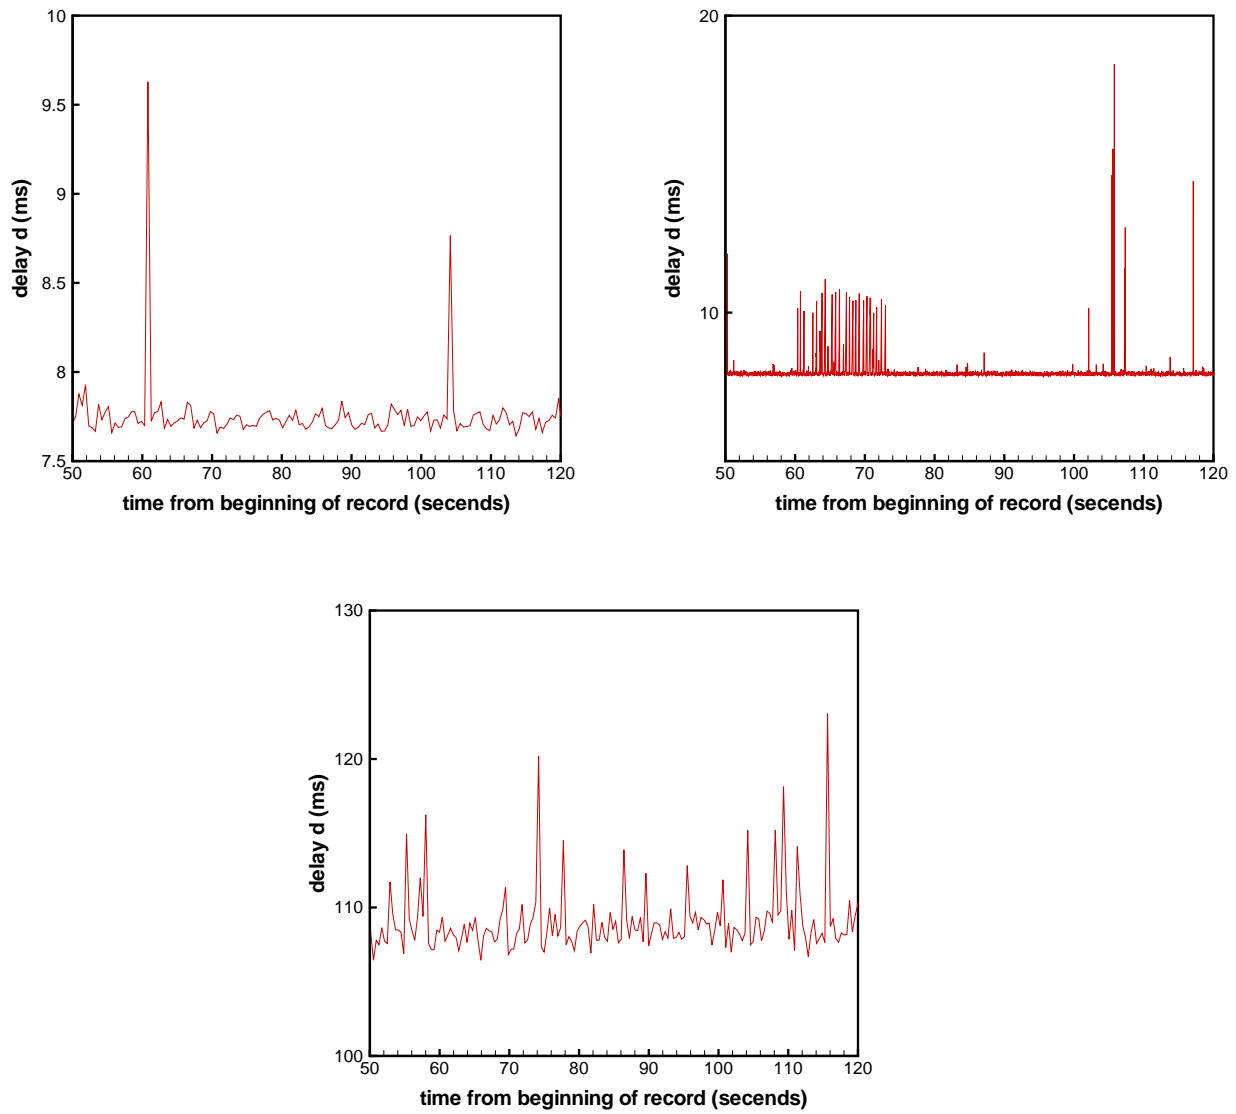


Figure 2: Typical time series of delay over a window inside a 300 second measurement period. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

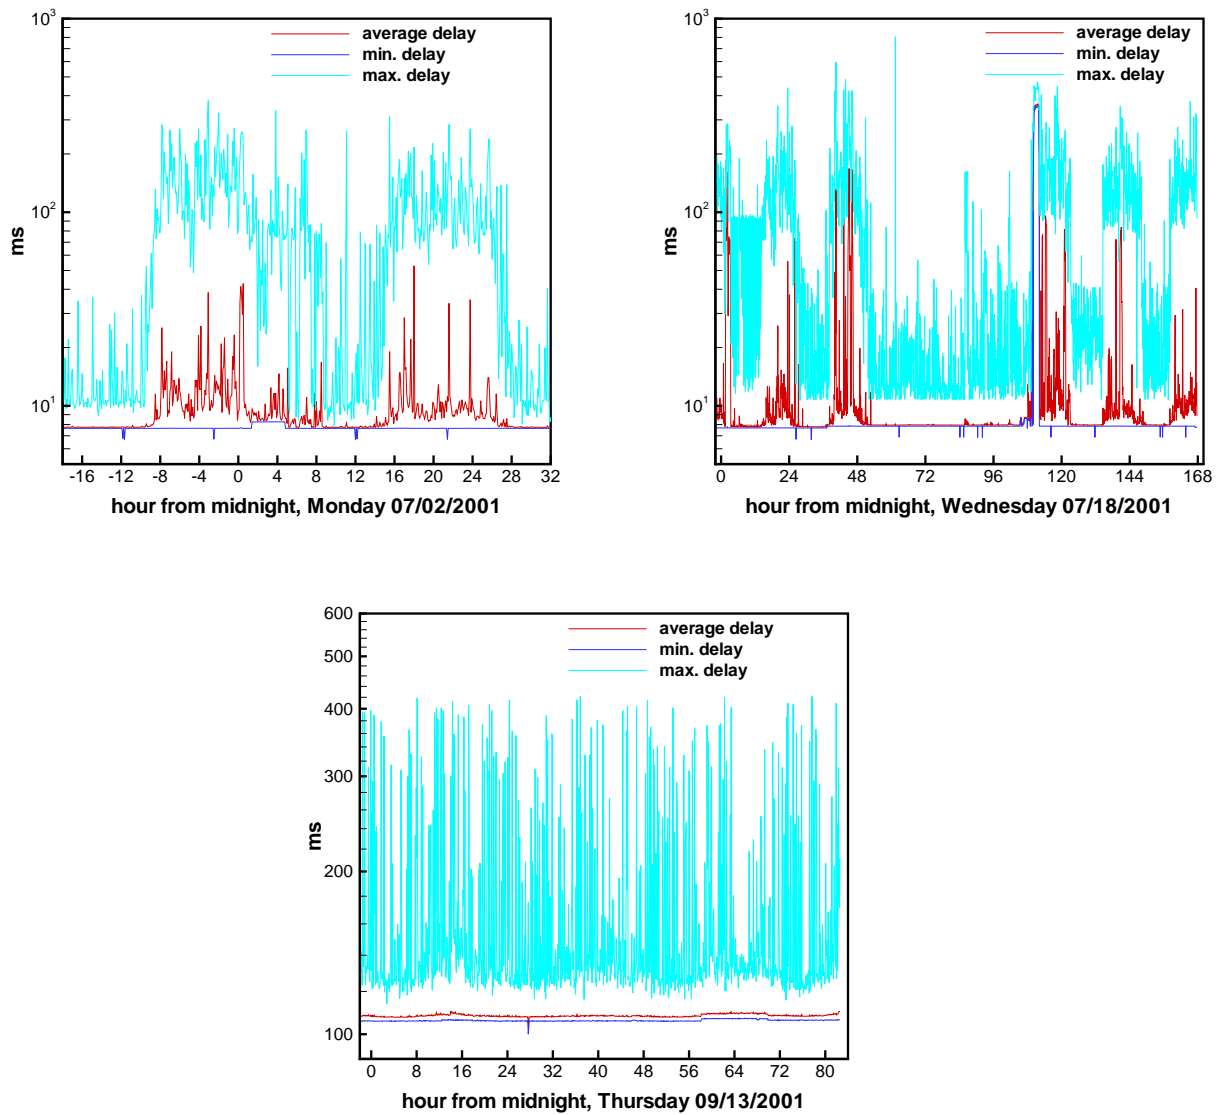


Figure 3: Average, minimum and maximum delay over consecutive 300 second measurement records. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

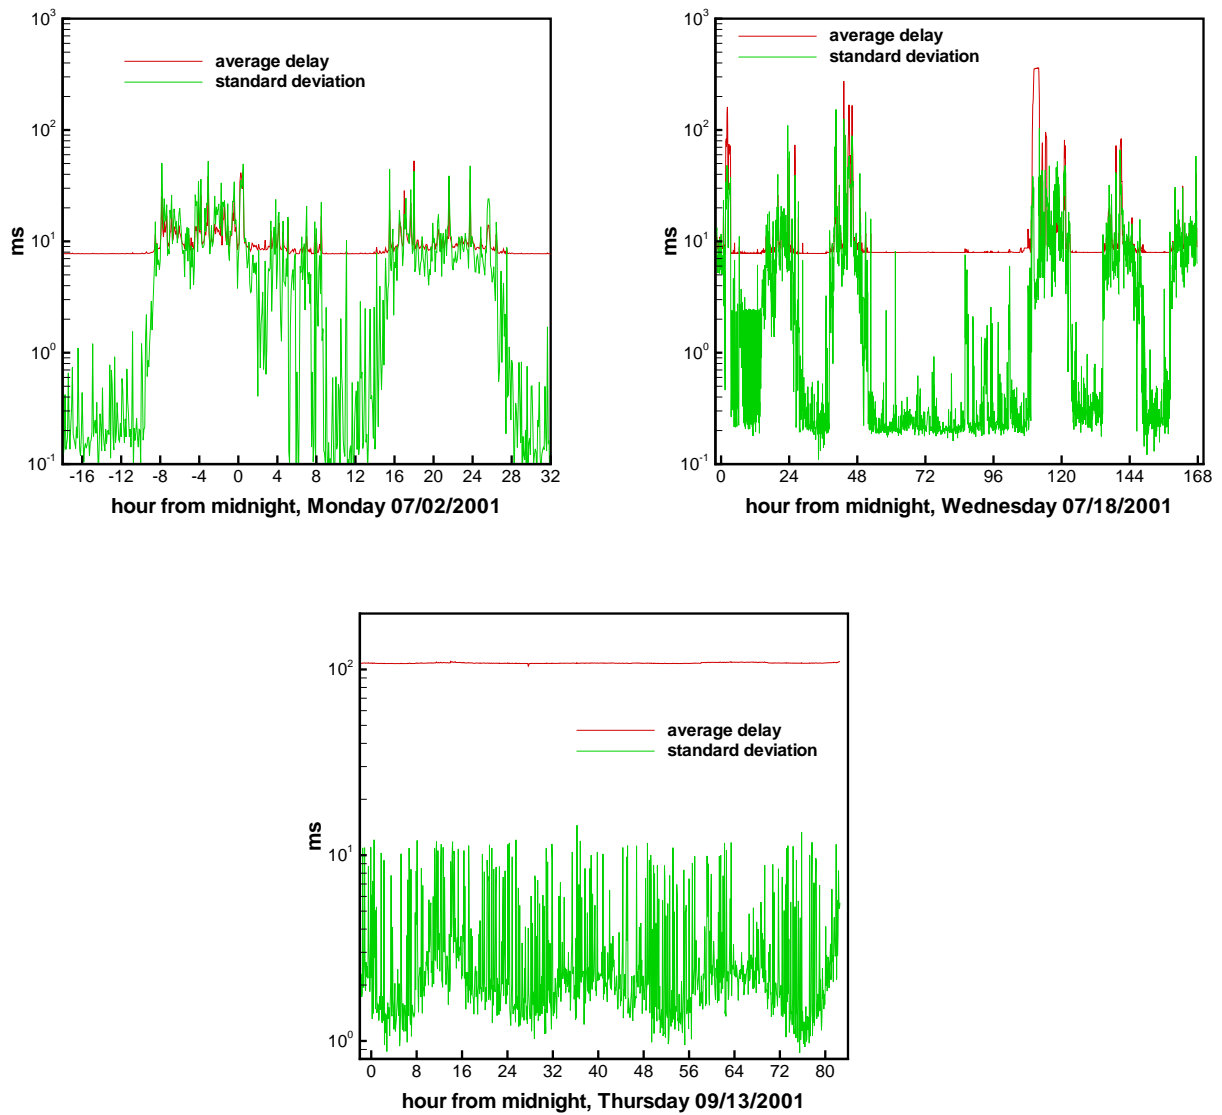


Figure 4: Average delay and standard deviation over consecutive 300 second measurement records. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

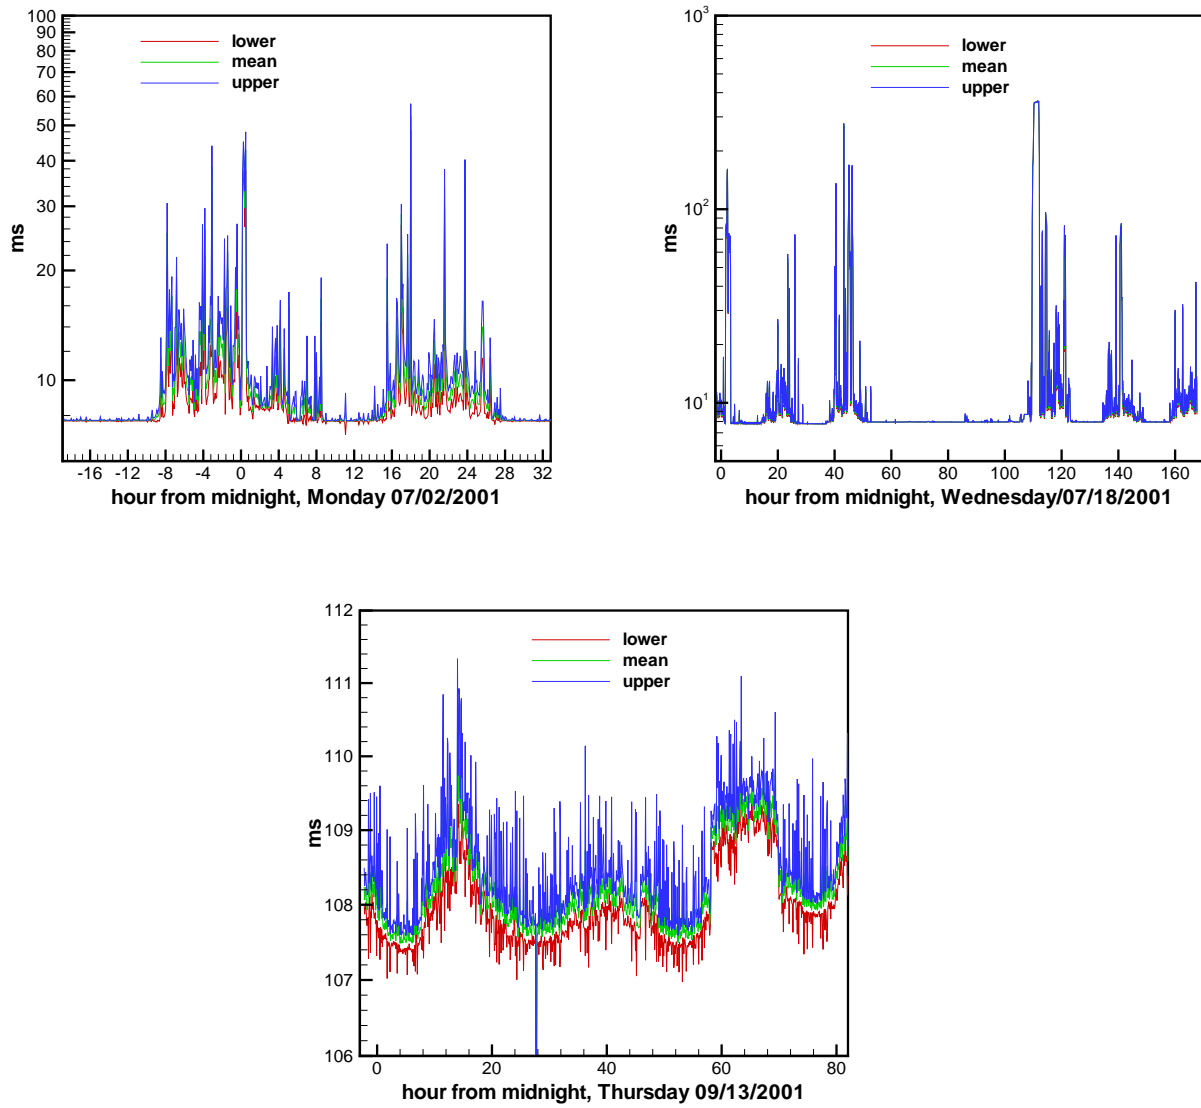


Figure 5: Average delay and upper and lower 99% confidence level. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

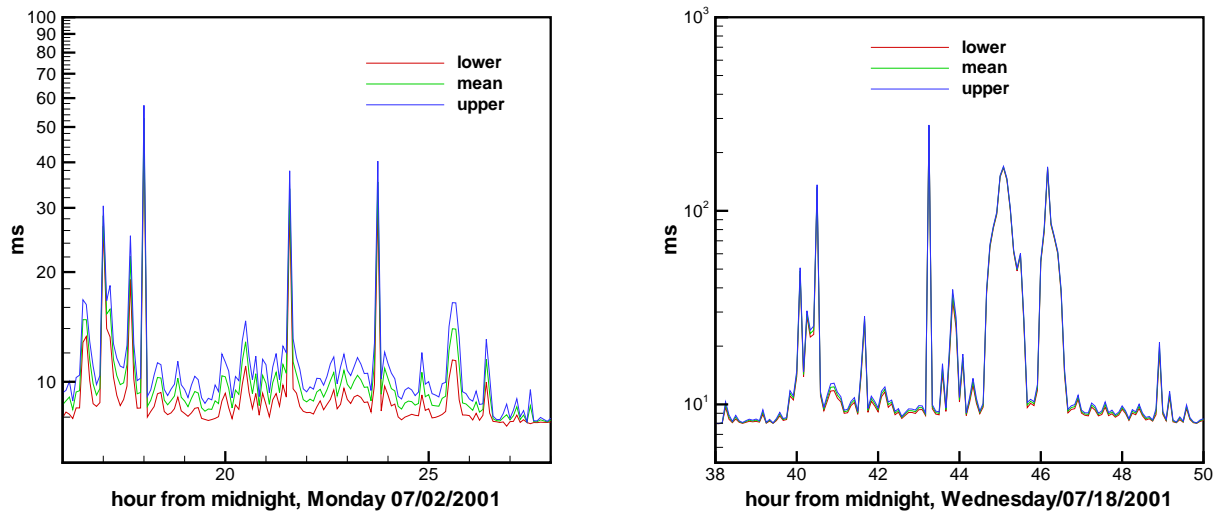


Figure 6: Average delay and upper and lower 99% confidence level inside a window within datasets. Left; dataset # 1. Right; dataset # 3.




Figure 7: Delay pdfs over typical 300 second measurement record. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.




Figure 8: Autocorrelation function over one typical 300 second measurement record for each of three datasets. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4.

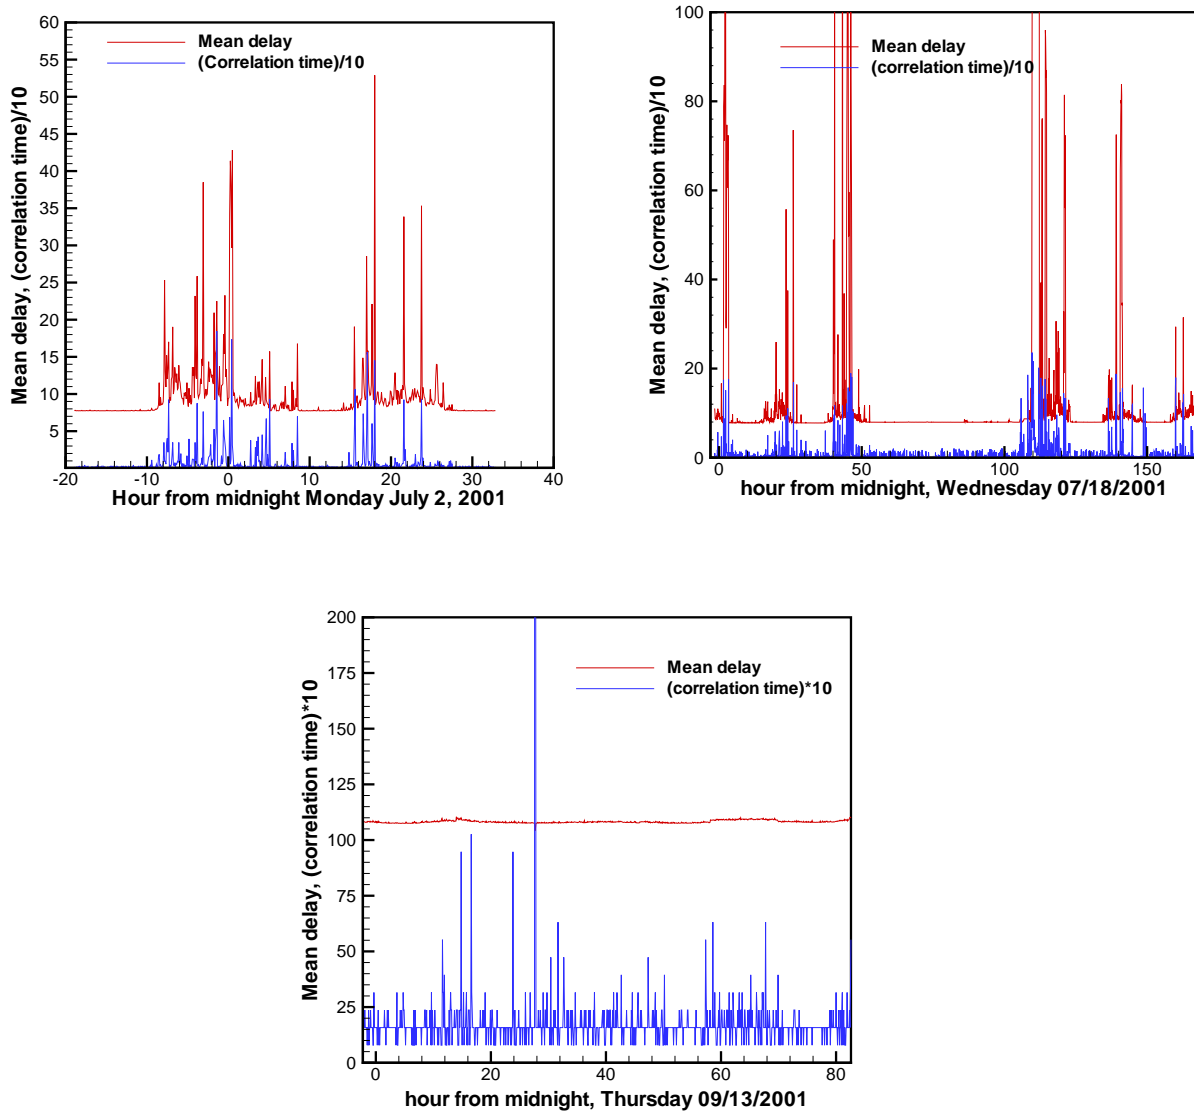


Figure 9: Correlation time and mean delay. Top left; dataset # 1. Top right; dataset # 3. Bottom; dataset # 4. Note the different scaling of the correlation time for # 4 compared with # 1 and # 3

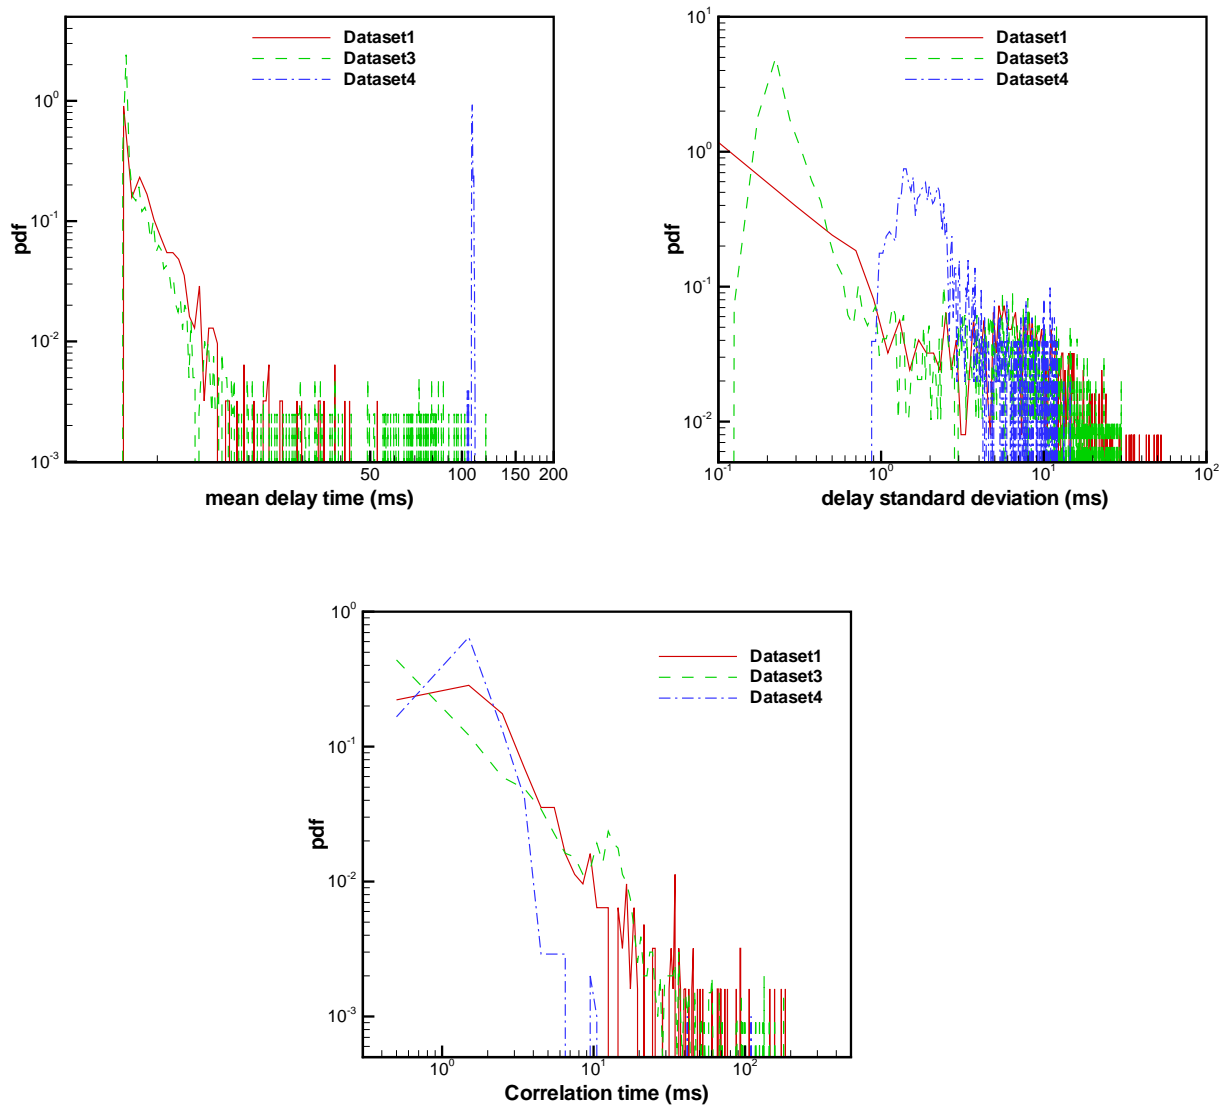


Figure 10: pdfs over typical 300 second measurement records. Top left; mean delay. Top right; standard deviation of delay. Bottom; correlation time.

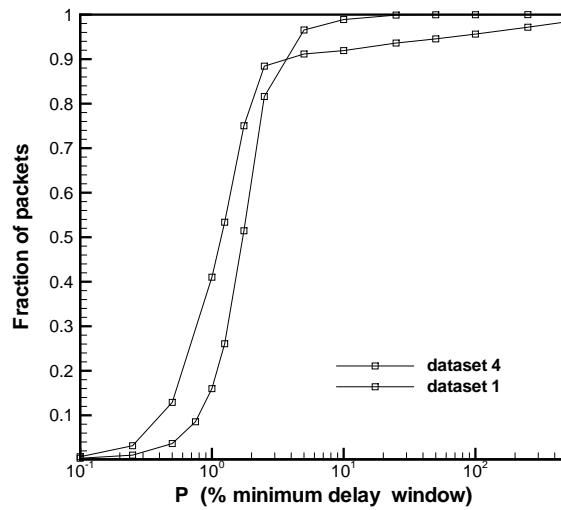


Figure 11: Average fraction of packets with delay within $P\%$ minimum delay window for datasets #1 and # 4

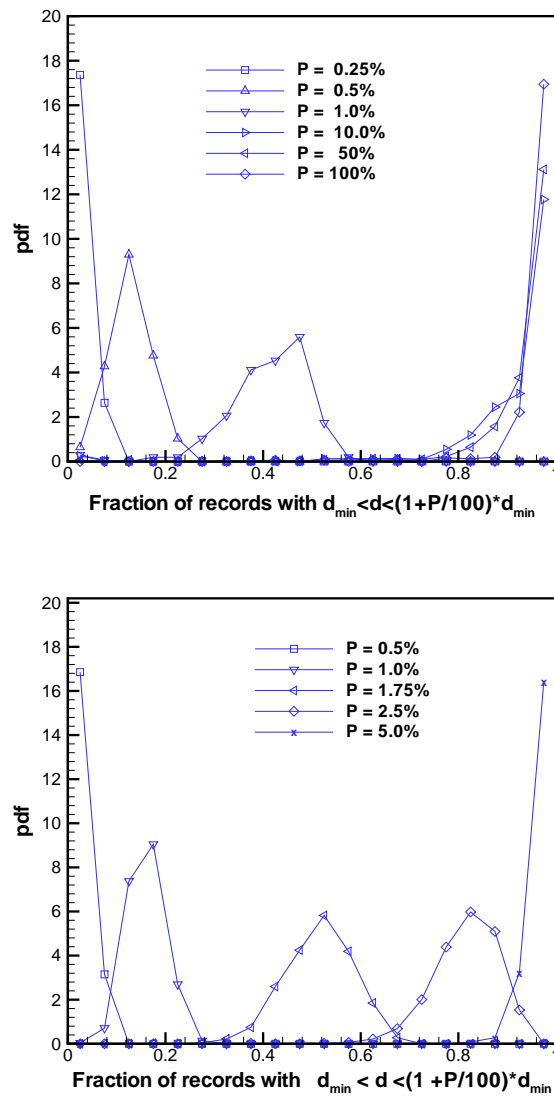


Figure 12: Probability densities for fraction of packets with delay in the $P\%$ minimum delay window. Top, dataset #1. Bottom, dataset #4.

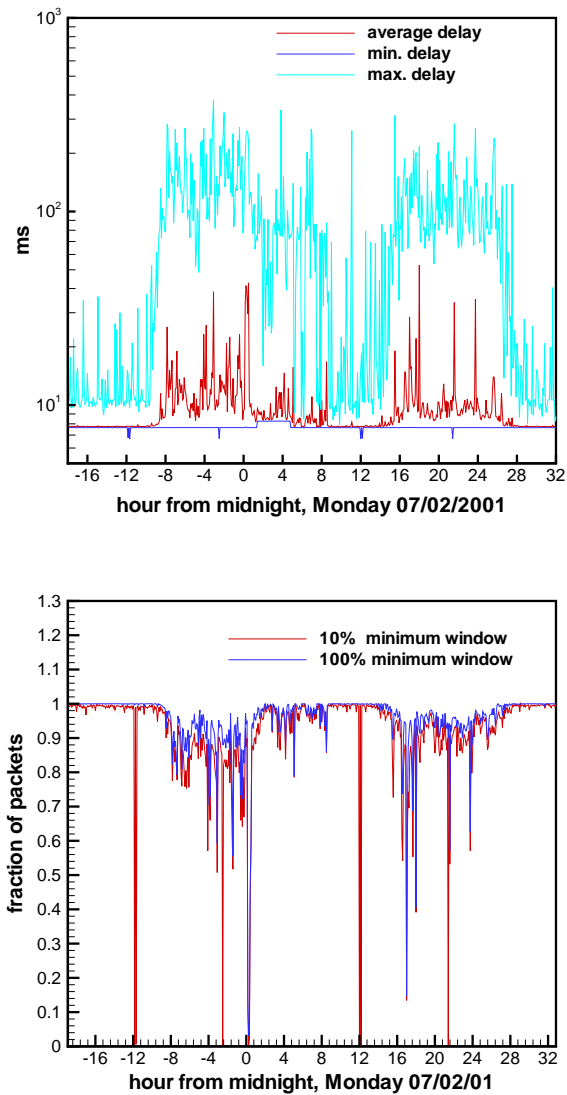


Figure 13: Dataset #1. Top; average, minimum and maximum delay over consecutive 300 second measurement records. Bottom; fraction of packets within each measurement record with delay within $P\%$ minimum delay window, $P = 10\%$, 100% .




Figure 14: Dataset #4. Top; average, minimum and maximum delay over consecutive 300 second measurement records. Bottom; fraction of packets within each measurement record with delay within $P\%$ minimum delay window, $P = 10\%, 20\%$.




Figure 15: Power spectrum of measured delay time series. Top, dataset # 1. Bottom, dataset # 4