Statistical Analysis of Time Series in IP networks

Thesis

Written by:
Éltető Tamás

Supervisor:
Pröhle Tamás

Eötvös Loránd University of Budapest
Faculty of Science

Department of Probability Theory and Statistics

1998.

Contents

List of Figures

Chapter 1

Introduction

Abstract

The purpose of this paper is to examine certain statistical aspects of time series that can be measured in IP networks. In recent times the word *internet* has become very popular. Many statistics tell us about the exponentially growing number of participants, the "citizens of the internet", as they called. Perhaps it is not completely useless to find some way of defining and controling the state of the networks, and at least to be able to give notice when something is going wrong.

1.1 About TCP

TCP (Transmission Control Protocol) is the most widely used tool (protocol) for exchanging data via the Internet. The protocol suite allows computer of all sizes, from many different computer vendors, running totally different operating systems, to communicate with each other. In the late 1960's a US government-financed research project started dealing with packet switching networks. In the 1990's it turned into the most widely used form of networking between computers. For a historical perspective on the early development of TCP/IP se [1].

TCP can be considered a complex algorithm which has several states, and the processing moves around these states according to the demands of the environment. There is one particularly important part of this protocol from our point of view: Bulk Data Flow.

Because my aim is to give a description and model for the interarrival time of packets travelling through the Internet in the special case of a continuous data stream, it is quite obvious to examine the practical background of the processes. I shall neglect TCP's connection and breaking it off, and concentrate directly on the exchange of data. First of all we have to define the three determinant parts of the communication:

- Sender - this will be the sender TCP
- Receiver - this will be the receiver TCP
- Channel - this is the Internet

TCP is a self controlled system, which means that the current pace of the network is controlled by the previous pace, and by the state of the network. The sender TCP emits a packet into the network and when the receiver TCP accepts this packet, it replies with an acknowledgement packet. The sender TCP is allowed to emit one more packet into the network, when it gets this acknowlegdement. This is done by sliding windows, as follows.

The sender TCP administers the number of the next byte to be sent and the last received, and acknowledged byte. The number of the last acknowledged byte is increased when the sender gets an acknowledgement - by the amount indicated in this packet. The number of the next byte to be sent shows the first unsent byte. The "*next byte to be sent*" number is always greater than the last acknowledged byte number. These two values determine an important concept, the *Congestion Window*. The former number is the right (upper) edge of this window, the latter is the left (lower) edge. This window slides along the data set. Its right side can move both left and right, its left side only to the right.

How does TCP control itself? When TCP gets an acknowledgement, it increases the lower end of the *Congestion Window*, as described before. Now the window itself doesn't deflate, so the upper end is increased by the same value. This is the *next byte to be sent*, that is, the TCP sends as many packets as can be fitted into the occurring gap without fragmentation.

There are two main phases of TCP under Bulk Data Flow: the Slow Start, and the Congestion Avoidance.

1.1.1 Slow Start

When a new connection is established, the TCPs may not know anything about the capacity of the line between them. The first aim is to discover the capacity of the line, and Slow Start is intended to do this. The method is the following:

- The sender sends a packet
- When the sender accepts the acknowledgement, it opens the congestion window for one packet.
- Each time the sender gets an acknowledgement, it opens the congestion window for one packet.

The effect of the opening is that the sender sends two packets for each acknowledged packet. One because of the sliding window, and one because of the opening of the window. This protocol exponentially increases the number of the packets flowing between the two parties, that is, the number of packets in the *pipe*. The time scale here is the round trip time. The growth of the number of packets in the pipe implies the growth of the sending rate.

The sender follows this protocol until packet loss occurs. In case of packet loss it sets a value called *ssthresh* to half of the current value of the congestion window. Then it reduces the congestion window to one packet, and returns to the begining of Slow Start.

After the first Slow Start phase, when *ssthresh* is already set, the sender performs Slow Start until packet loss occurs, or until the size of the congestion windowhas reached the value of *ssthresh*. In this case it ends Slow Start and enters into Congestion Avoidance.

1.1.2 Congestion Avoidance

The main purpose of Congestion Avoidance is expressed in its name: we would like to avoid packet congestion in the network. On the other hand, it is possible that there is free capacity in the line, and we would like to utilise this capacity; therefore it's no good to remain at a fixed sending rate. To satisfy both demands, in Congestion Avoidance we increase the sending rate, but only linearly (approximately). Certainly this is done by increasing the congestion window:

$$cwnd_{\mathrm{new}} = cwnd_{\mathrm{old}} + \frac{\mathrm{packet\_size}^2}{cwnd_{\mathrm{old}}}$$

The result of this formula is that the number of packets in the pipe increases by one packet per round trip time. Sooner or later we reach the limit capacity of the connection and packet loss occurs.

How does the sender detect packet loss? When the receiver TCP gets packets of upper part of the data stream, it concurrently notifies, puts this packet into a buffer, and replies with an acknowledgement of the last packet which was received continuously. Thus, in case of packet loss, the sender gets repeated acknowledgements. The receiver puts each out-of-order packet into its buffer until the buffer is full, or until it receives the missing packet. When this latter happens, the receiver empties its buffer and replies, an acknowledgement of the highest number of the received bytes in the data stream. It is possible that a reordering of packets occurs along the line, so repeated acknowledgement doesn't necessearly mean packet loss. But it is not frequent, and it usually involves neighbouring packets. The conclusion is that we detect packet loss in case of three repeated acknowlegdements.

In this case we have to reduce the congestion window to one packet, and enter into slow start phase. However, it is possible that we have packet loss because of something else. Let us assume, for example, that there is another TCP connection concurring through a common buffer. This connection may fill the queue of this computer for a while, and therefore it drops one of our packets. After the other TCP detects its own packet loss, it reduces its filling pace, and we are able to use that buffer's queue again, and don't have to enter into slow start, because that would cause an unnecessary reduction of the connection's speed. *Fast Retransmit* is intended to handle this case.

When the sender gets the third repeated acknowledgement, it retransmits the missing packet, and although it gets further repeated acknowlegements (about the size of the congestion window), it keeps on sending the data stream as if it has not been interrupted. The assumption is that we have occasional packet loss, and the receiver misses only one packet. According to the rules of the sliding window protocol, the sender has to increase the congestion window, because the repeated serial number doesn't allow the congestion window to slide. In this case we can't say that the size of the congestion window equals the number of packets in the pipe, but the rule of thumb that we must not increase the load of the pipe is not violated, because of the assumption on the repeated acknowledgements.

1.1.3 The advertised window

The most important concept in the spirit of my paper is the *Advertised Window*. As we have seen in the previous section, sometimes the receiver is required to buffer the data stream. The space in the buffer is bounded, so the sender must not count on buffering for a long time. In an acknowledgement, there is a field which expresses the size of the remaining buffer, i.e. the size of buffer space at the moment of emitting that acknowledgement.

With fast retransmit, the congestion window increases, and the advertised window decreases. The sender is not allowed to keep more packets in the pipe than the advertised window. If the advertised window becomes zero, the sender stops sending more packets until it becomes non-zero again.

In my investigations, the most common case was that after slow start and congestion avoidance, the congestion window reaches the size of the advertised window, and for a shorter or longer time interval, it remains fixed.

1.1.4 A bit more on TCP

There is one very important rule on Internet. We must reduce the number of packets in the network as much as possible. We face this problem quite often in the case of interactive data flow, but this is beyond our scope here. The acknowledgements take part in this effort as well. The receiver does not reply to a packet immediately, rather it starts a timer. When this timer goes off, the receiver sends a *delayed acknowledgement*. The meaning of this is that the receiver may get more packets before the expiration, and therefore it is able to acknowledge more packets at the same time. Usually the TCP does not wait for a third packet before its timer expires, but sends an acknowledgement when the second packet arrives. There are several implementations where the TCP is even more economical, and waits for a third packet to acknowledge before the expiration. For a detailed description of the TCP/IP protocol see [2].

The delayed acknowledgements cause a very strong time pattern in inter-arrival time series, and make the traffic become more burst-like.

Chapter 2

Delay Series

I examined this series on two different way. One was an attempt to fit a more traditional model to the series, namely an ARMA model and see the attributes of a fitted model. The conclusion of this is, that it is possible to find a fairly good ARMA model to the series, however there are two strong counter-arguements against this kind of modeling. One is that an ARMA model of this kind - we shall see that it has many parameteres - hides the details of the series itself, and even if it is a simple model, we may find a simpler model. The other problem is that the process that control the delay - see the previous and the final chapter - is not dependent from the past several seconds, only from several hundreds of milliseconds.

2.1 Fitting of an ARMA model

My first aim was to fit an ARMA model to a smaller part of the series. I chosen a 500 data long series that we can see on Figure 2.1. The interval starts around 4 pm, and 500 seconds long. This part of the day is not so busy than the previous (peak) hours.

The first thing to examine is the autocorrelation and partial autocorrelation function of the series. The estimations are shown on Figure 2.2. The initial guess for the parameters of the autoregressive (p), and moving average (q) parameter of the model was 7 and 8 respectively, so I started with ARMA(7,8).

The parameter estimation was done by the program SPSS 7.5, and the brief results are shown on Table 2.3. There we can see, that the constant

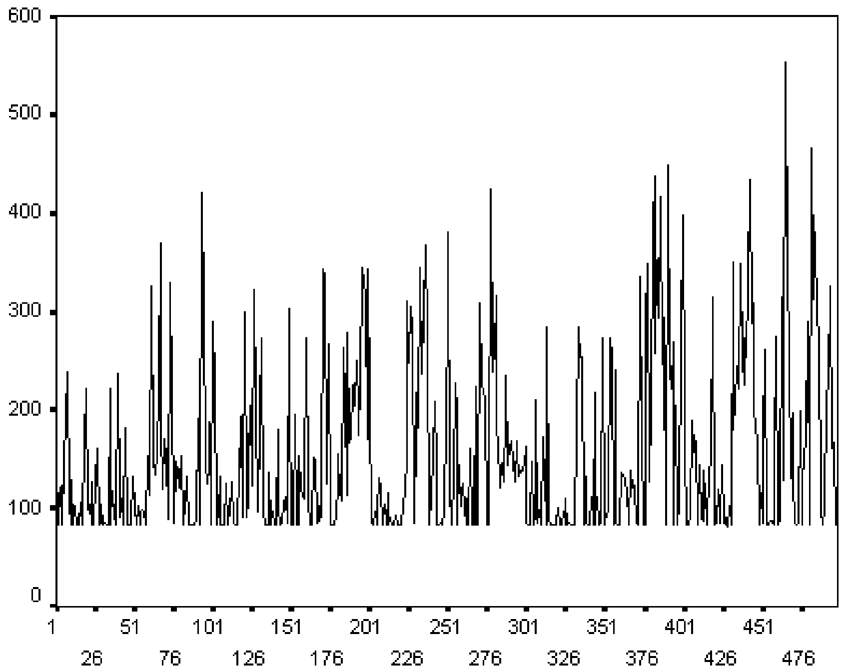

Figure 2.1: The 500 second long interval

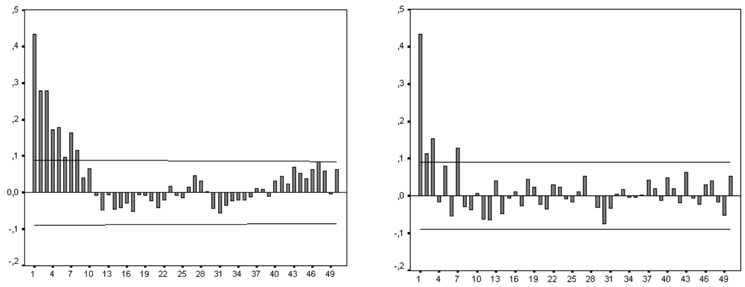

Figure 2.2: The estimated ACF, and PACF of the series

	B	SEB	APPROX. PROB.
AR1	.88163	.1501666	.00000000
AR2	.18024	.2117400	.39504649
AR3	.33469	.1995433	.09413056
AR4	-.66033	.1389176	.00000264
MA1	.50813	.1513635	.00084949
MA2	.34643	.1656694	.03703683
MA3	.33397	.1756222	.05781014
MA4	-.39394	.0939374	.00003261
MA5	-.20570	.0565317	.00030336
MA6	.09842	.0560762	.07985995
MA7	-.23452	.0489499	.00000221
CONSTANT	155.26124	7.0518565	.00000000
Residual Variance	5792.2123		

Table 2.1: The parameter estimation of ARMA(4,7)

	B	SEB	APPROX. PROB.
AR1	.89085	.2462230	.00032792
AR2	.18177	.2301219	.42998978
AR3	.28260	.2468839	.25291650
AR4	-.67979	.2098339	.00127869
AR5	.08885	.2123568	.67582901
MA1	.50817	.2423256	.03650481
MA2	.34796	.1771354	.05005487
MA3	.27893	.2244039	.21448097
MA4	-.43759	.2057320	.03392254
MA5	-.13685	.1457087	.34808090
MA6	.12539	.0823178	.12833773
MA7	-.21212	.0534513	.00008322
CONSTANT	155.00767	7.6231436	.00000000
Residual Variance	5854.3079		

Table 2.2: The parameter estimations of ARMA(5,7)

value is significantly nonzero and also the parameters AR5 and MA7 differ from zero significantly. This implies, that it is better to use a simpler model, namely ARMA(5,7).

	B	SEB	APPROX. PROB.
AR1	.60065	.8398467	.47483782
AR2	-.19563	.4418463	.65814263
AR3	.14490	.4078446	.72254169
AR4	-.22673	.1579408	.15177932
AR5	.69212	.2816353	.01434042
AR6	-.09127	.5381483	.86539502
AR7	-.27225	.2602182	.29597792
MA1	.21649	.8445944	.79780651
MA2	-.14115	.3561480	.69203885
MA3	-.05950	.4546398	.89593285
MA4	-.17937	.3150430	.56937369
MA5	.48069	.4097054	.24126643
MA6	.21260	.2859640	.45756816
MA7	-.37072	.1931372	.05551238
MA8	-.00193	.1727638	.99108310
CONSTANT	154.64048	8.2345099	.00000000
Residual Variance	5820.9730		

Table 2.3: The parameter estimations of the model ARMA(7,8)

This time (Table 2.2) we can see that the constant has significant nonzero value. The parameter MA7 is nonzero also, but this time AR5 is neglectabe and AR4 is significant, therefore the I fitted ARMA(4,7) to the series, and Table 2.1 shows the parameter estimations of this model.

On Table 2.1 almost all the parameters have significant values, therefore I accepted it as the final model for the series. We can see that the residual variances are not significantly different in the three model, but the ARMA(4,7) estimation has the smallest value. The significant values of the estimation suggest that it is worth examining some other aspects of the estimation.

The program computed the fitted values by the estimated model, and also computed the error that is the difference of the original and the fitted data. Figure 2.3 shows the error series, while Figure 2.4 shows the estimated autocorrelation and partial autocorrelation functions.

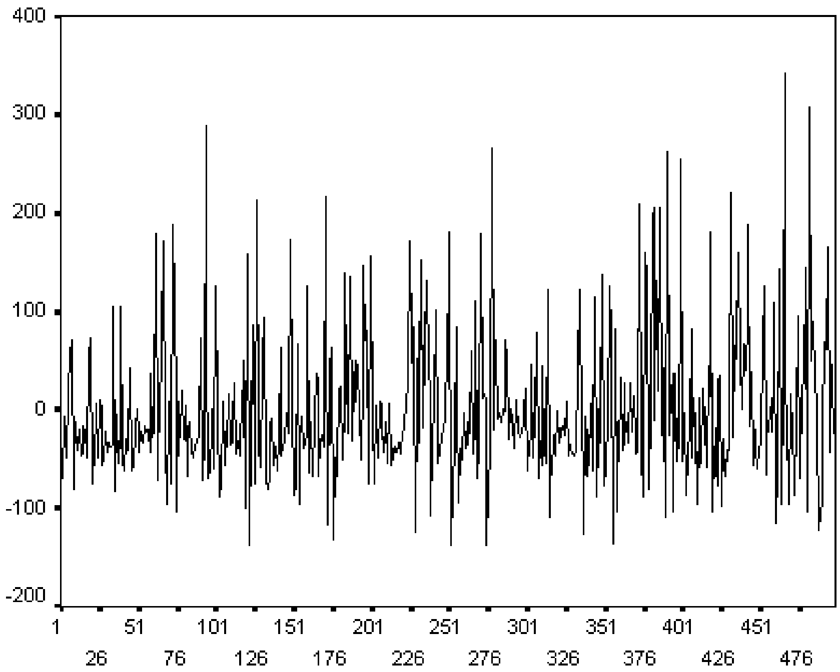

Figure 2.3: The error series

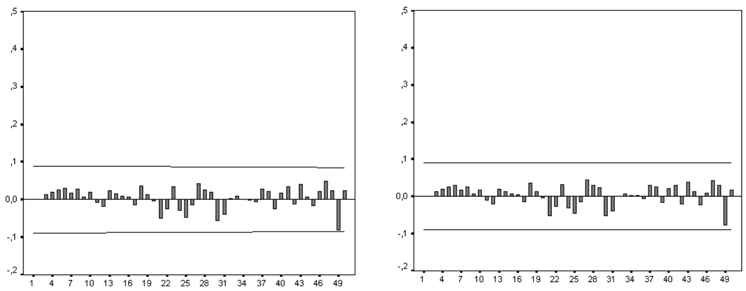

Figure 2.4: The estimated ACF and PACF of the error series

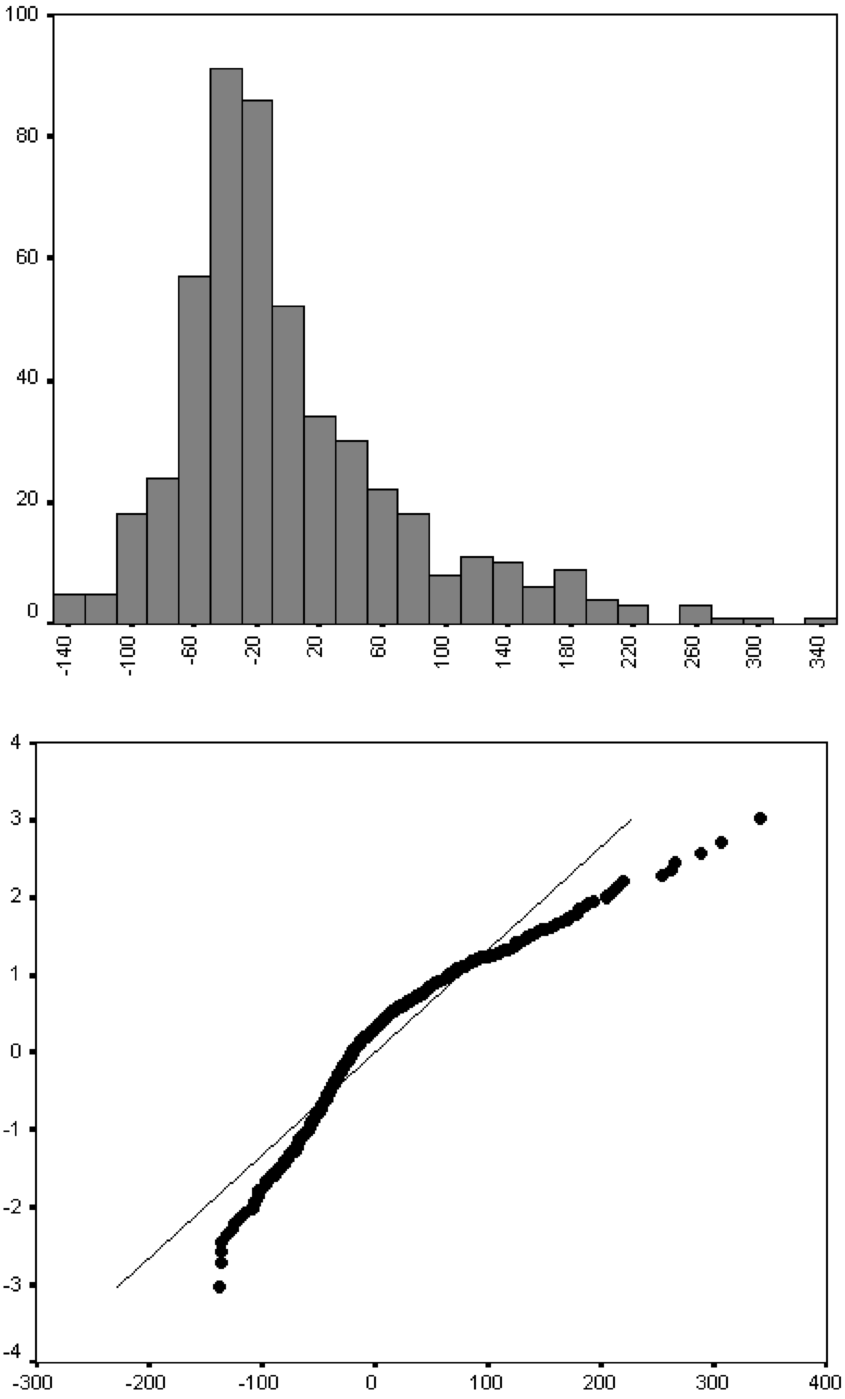

Figure 2.5: The histogram, and the normal quantile-quantile plot

We can say that the errors are independent and identically distributed, therefore it is white noise. The question arises that the distribution of the error is normal or something else.

Figure 2.5 shows the histogram of the series, which is unfortunately not symmetrical. Indeed, if we fit a normal distribution to the error series, and compute the quantile-quantile plot of the two functions we can see, that the two distributions are considerably different. (The quantile function will be discussed later.)

2.2 Finding Appropriate Time Scales

I wrote at the begining of the chapter that I tried to find an other model, which is more simple than ARMA. My model states, that the series is simply concateation of intervals, that are iid processes. Are there any interval of an iid process in the series? If there are, then how can we detect these intervals? Finally, what kind of distribution these intervals have?

First I tried to find the edges of the intervals by assuming that the state of the network changes suddenly. (By studying the TCP connection we can see that it reaches the limit of the buffer very quickly.) Therefore we used nonparametric change-point detecion for finding these points. This is done in two parts.

- Use a hypothesis test to find an interval, where is a change in distribution. The H_0 hypothesis states that the appropriate interval is independently and identically distributed. In the literature many authors examine different estimators under H_0 or H_1.
- Use change-point detection to find the appropriate location of the change point within an interval.

When we have designated the interval, we should check the trends in the interval. This is done by moving averages, where we present bounds on the average of a smaller subinterval.

2.3 The two sample Kolmogorov-Smirnov test

The most important tool that theoretical change-point detection uses, and I shall also use in practice, is the (two sample) Kolmogorov-Smirnov test.

This section describes how this test works, and also tells about the special case of equal sample series length.

Let

$$X_1,\ X_2,\ \ldots X_N$$

be the series of the first sample series (of length N), and

$$Y_1,\ Y_2,\ \ldots Y_M$$

be the series of the second. Now $F_N(x)$ and $G_M(x)$ are the corresponding empirical distributions. Then the statistics is

$$D_{N,M} = \max_x |F_N(x) - G_M(x)|$$

The limit distribution theorem says

$$\lim_{N\to\infty} P\left(\sqrt{\frac{NM}{N+M}} D_{N,M} < y\right) = K(y), \quad 0 < y < \infty,$$

where $K(y)$ is the Kolmogorov distribution.

Assume that $N = M$. In this case we count $D_{N,N}$ practically in the following way. We order the united sample in an increasing order

$$Z_1^{(2N)} \leq Z_2^{(2N)} \leq \ldots \leq Z_{2N}^{(2N)}.$$

Now, let

$$\delta_i = \begin{cases} 1, \text{if } Z_i^{(2N)} \text{ comes from the } X_k \text{ series} \\ -1, \text{else} \end{cases}$$

and

$$S_j = \delta_1 + \delta_2 + \ldots + \delta_j.$$

Here we get $D_{N,N}$ so that

$$D_{N,N} = \frac{1}{2N} \max_{1\leq j\leq 2N} |S_j|.$$

2.4 Different hypotheses on change-points

Let $X_1, \ldots, X_n$ be independent random variables. Suppose we want to test the null hypothesis

H_0 $X_i, 1 \le i \le n$, have the same distribution

versus the alternative hypothesis that there is a change-point in the sequence $X_1, \ldots, X_n$ namely, that we have

H_1 there is a $\lambda \in (0,1)$ such that $P(X_1 \le t) = P(X_2 \le t) = \cdots = P(X_{[n\lambda]} \le t)$, $P(X_{[n\lambda]+1} \le t) = \cdots = P(X_n \le t)$, $-\infty < t < \infty$, and $P(X_{[n\lambda]} \le t_0) \neq P(X_{[n\lambda]+1} \le t_0)$ for some t_0.

Csörgő and Horváth study tests for the change-point problem which is based on processes of U-statistics. Let $h(x, y)$ be a symmetric function and consider

$$Z_k = \sum_{1 \le i < k} \sum_{k+1 \le j \le n} h(X_i, X_j), \quad 1 \le k < n \tag{2.1}$$

Given $H_0, X_1, \ldots, X_n$ are i.i.d. random variables we assume

$$Eh^2(X_1, X_2) < \infty \tag{2.2}$$

and let $Eh(X_1, X_2) = \Theta$, $\tilde{h}(t) = E(h(X_1, t) - \Theta)$. Condition 2.2 implies that $E\tilde{h}^2(X_1) < \infty$ and we assume

$$0 < \sigma^2 = E\tilde{h}^2(X_1). \tag{2.3}$$

Here they investigate

$$U_k = Z_k - k(n-k)\Theta, \quad 1 \le k < n.$$

While U_k is not a U-statistic, it can be expressed as a linear combination of U-statistics.

In order to state their results they define the Gaussian process Γ by

$$\Gamma(t) = (1-t)W(t) + t(W(1) - W(t)), \quad 0 \le t \le 1, \tag{2.4}$$

where $\{W(t),\ 0 \le t < \infty\}$ is a Wiener process.

Theorem 1 *Assume that H_0 holds, and 2.2 and 2.3 are satisfied. Then we can define a sequence of Gaussian processes $\{\Gamma_n(t),\, 0 \le t \le 1\}$ such that, as $n \to \infty$,*

$$\sup_{0 \le t < 1} \left| \frac{n^{-\frac{3}{2}}}{\sigma} U_{[(n+1)t]} - \Gamma_n(t) \right| = o_P(1), \tag{2.5}$$

where for each $n \ge 1$

$$\{\Gamma_n(t),\ 0 \le t \le 1\} \overset{D}{=} \{\Gamma(t)\}. \tag{2.6}$$

Proof. Theorem 2.1 in [3].

Theorem 2 *Assume that H_0 holds,*

$$E|h(X_1, X_2)|^{\nu} < \infty \quad \textit{for some} \quad \nu > 2, \tag{2.7}$$

and 2.3 is satisfied. Then we can define a sequence of Gaussian processes $\{\Gamma_n(t), 0 \le t \le 1\}$ such that 2.5 holds,

$$\sup_{1/(n+1) \le t \le n/(n+1)} \left| \frac{n^{-\frac{3}{2}}}{\sigma} U_{[(n+1)t]} - \Gamma_n(t) \right| \Big/ (t(1-t))^{1/2} = O_P(1), \tag{2.8}$$

and we have 2.6 for each $n \ge 1$.

Proof. Theorem 2.2 in [3].

Let $a(y, \log n) = (y + 2\log\log n + \frac{1}{2}\log\log\log n - \frac{1}{2}\log\pi)(2\log\log n)^{-1/2}$, $-\infty < y < \infty$.

Theorem 3 *Assume that H_0 holds, and 2.3, 2.7 are satisfied. Then*

$$\lim_{n\to\infty} P\left\{\sigma^{-1} \max_{1 \le k \le n} \frac{U_k}{(k(n-k+1)n)^{1/2}} \le a(y, \log n)\right\} = \exp\left(-\exp\left(-y\right)\right) \tag{2.9}$$

and

$$\lim_{n\to\infty} P\left\{\sigma^{-1} \max_{1 \le k \le n} \frac{|U_k|}{(k(n-k+1)n)^{1/2}} \le a(y, \log n)\right\} = \exp\left(-2\exp\left(-y\right)\right) \tag{2.10}$$

Proof. Theorem 2.3 in [3].

Typical choices of h are xy, $(x-y)^2/2$ (sample variance in U-statistics), $|x-y|$, $\mathrm{sign}(x+y)$. The results cannot be applied in the case of $h(x,y) = \mathrm{sign}(x-y)$, because this is not a symmetric function. However Csörgő and Horváth in [3] show that there are similar results in the case of antisymmetric kernels. Assume that

$$h(x,y) = -h(y,x). \tag{2.11}$$

In this case $Eh(X_1,X_2) = 0$ and similarly to the symmetric case let $\tilde{h}(t) = Eh(t,X_1)$. Assume

$$Eh^2(X_1,X_2) < \infty \quad \text{and} \quad 0 < \sigma^2 = E\tilde{h}^2(X_1). \tag{2.12}$$

Now $U_k = Z_k$, where Z_k is defined by 2.1.

Theorem 4 *Assume that H_0 holds, and 2.11 and 2.12 are satisfied. Then we can define a sequence of Brownian bridges $/B_n(t), 0 \le t \le 1/$ such that, as $n \to \infty$,*

$$\sup_{0\le t\le 1}\left|\frac{n^{-\frac{3}{2}}}{\sigma}U_{[(n+1)t]} - B_n(t)\right| = o_P(1). \tag{2.13}$$

Proof. Theorem 4.1 in [3]

Theorem 5 *Assume that H_0 holds, 2.11 and 2.12 are satisfied, and*

$$E|h(X_1,X_2)|^\nu < \infty \quad for\ some \quad \nu > 2, \tag{2.14}$$

Then we can define a sequence of Brownian bridges $/B_n(t), 0 \le t \le 1/$ such that 2.13 holds and

$$\sup_{1/(n+1)\le t\le n/(n+1)}\left|\frac{n^{-\frac{3}{2}}}{\sigma}U_{[(n+1)t]} - B_n(t)\right| \Big/ (t(1-t))^{1/2} = O_P(1). \tag{2.15}$$

Proof. Theorem 4.2 in [3].

Theorem 3 holds in this case too, according to Theorem 4.3 in [3].

In [3] they consider the asymptotics in case of H_1. Let

$$\Theta = Eh(X_{[n\lambda]-1}, X_{[n\lambda]}), \quad \mu = Eh(X_{[n\lambda]+1}, X_{[n\lambda]+2})$$
$$\tau = Eh(X_{[n\lambda]}, X_{[n\lambda]+1}),$$

and write $\log^+ x = \log(x \vee 1)$.

Theorem 6 *Assume that H_1 holds and*

$$E|h(X_{[n\lambda]-1}, X_{[n\lambda]})| < \infty, \quad E|h(X_{[n\lambda]+1}, X_{[n\lambda]+2})| < \infty,$$
$$E|h(X_{[n\lambda]}, X_{[n\lambda]+1})| \log^+ |h(X_{[n\lambda]}, X_{[n\lambda]+1})| < \infty \tag{2.16}$$

are satisfied. Then

$$\lim_{n\to\infty} Z_{[(n+1)t]}/n^2 = \begin{cases} \theta t(\lambda - t) + t\tau(1-\lambda), & 0 < t \le \lambda, \\ \mu(t-\lambda)(1-t) + \tau\lambda(1-t), & \lambda \le t < 1, \end{cases} \tag{2.17}$$

in probability.

Proof. Theorem 3.1 in [3].

The conclusion of the current subsection is that when the want to test the hypothesis that there is a change-pont in the current interval (H_1) or not (H_0), we reject H_0 for the large values of

$$\sigma^{-1} \max_{1\le k\le n} \frac{U_k}{(k(n-k+1)n)^{1/2}}$$

where U_k is defined as previously.

2.5 Change-point estimation in a fixed interval

Assume that we already have an interval, which contains a change-point. Recall that $\lambda \in (0,1)$ divides the interval into two different subsamples. Dümbgen in [4] considers the following change-point model: for $n = 2, 3, 4, \ldots$ let P_n and Q_n be two probability distributions on a measurable space E, and let $X_{1,n}, X_{2,n}, \ldots, X_{n,n}$ be independent random variables. The variables $X_{i,n}$ have distribution P_n for $i \le n\lambda_n$ and Q_n otherwise. The change-point λ_n is some unknown number in $T_n = \{1/n, 2/n, \ldots, (n-1)/n\}$. For increasing n λ_n remains bounded away from 0 and 1, and for simplicity assume that $\lambda_n \to \lambda \in (0,1)$ as $n \to \infty$.

We describe a general procedure for estimating λ_n. For every hypothetical change-point $t \in T_n$, consider the empirical distributions

$${}^tP_n = \frac{1}{nt}\sum_{i=1}^{nt} \delta_{X_{i,n}} \quad \text{and} \quad P_n^t = \frac{1}{n(1-t)} \sum_{i=nt+1}^{n} \delta_{X_{i,n}}.$$

The (pointwise) expectation of these empirical measures ${}^tP_n$ and P_n^t is denoted by ${}^t\Pi_n$ and Π_n^t respectively. Then the difference $P_n^t - {}^tP_n$ estimates the signed measure

$$\Pi_n^t - {}^t\Pi_n = [(\lambda_n/t) \wedge ((1-\lambda_n)/(1-t))]\,(Q_n - P_n).$$

If t is near 0 or 1, this is a poor estimator. Introduce the weights

$$w(t) \;=\; t^{1/2}(1-t)^{1/2}$$

and consider the measures

$$D_n^t = w(t)(P_n^t - {}^tP_n), \quad t \in T_n.$$

These signed measures D_n^t estimate

$$\Delta_n^t = w(t)(\Pi_n^t - {}^t\Pi_n) = \rho_n(t)(Q_n - P_n),$$

where

$$\rho_n(t) = (1-\lambda_n)[t/(1-t)]^{1/2} \wedge \lambda_n[(1-t)/t]^{1/2}.$$

On the interval $[0, \lambda_n]$ this function ρ_n is strictly increasing, and on $[\lambda_n, 1]$ it is strictly decreasing. It attains its maximum in λ_n with $\rho(\lambda_n) = w(\lambda_n)$. Now we choose a seminorm N_n on the space $\mathcal{M}$ of all finite signed measures and compute $N_n(D_n^t)$ as a measure of difference between ${}^tP_n$ and P_n^t. In particular consider the estimator

$$\hat{\lambda}_n = \arg\max\,(N_n(D_n^t) :\ t \in T_n).$$

If some prior information about λ_n is available, one can use a truncated estimator

$$\hat{\lambda}_n = \arg\max\,(N_n(D_n^t) :\ t \in T_n, t_0 \le t \le t_1),$$

where $0 < t_0 < \lambda < t_1 < 1$. The seminorm N_n may be deterministic or random. It is useful to compare the observed process $(N_n(D_n^t) : t \in T_n)$ with the theoretical process $(N_n(\Delta_n^t) :\ t \in T_n)$. The latter has a simple structure

$$N_n(\Delta_n^t) = \rho_n(t) N_n(Q_n - P_n), \quad t \in T_n.$$

Hence $\lambda_n = \arg\max\,(N_n(\Delta_n^t) :\ t \in T_n)$ and $\max\,(N_n(\Delta_n^t) :\ t \in T_n) = w(\lambda_n)$ $N_n(Q_n - P_n)$. That is $\hat{\lambda}_n$ is a reasonable estimator for λ_n, if the maximal difference of these two processes $N_n(D_n^{(\cdot)})$ and $N_n(\Delta_n^{(\cdot)})$ is small compared to $N_n(Q_n - P_n)$.

Theorem 7 *If the seminorms N_n and the distributions P_n and Q_n satisfies the requirements of Dümbgen in [4], particularly*

$$\begin{aligned} P(N_n(Q_n - P_n) \geq C_0\gamma_n^{-1}) \to 1 \quad \text{and} \\ \gamma_n(\log\log n)^{1/2} n^{-1/2} \to 0 \quad \text{as} \quad n \to \infty, \end{aligned} \tag{2.18}$$

where $C_0 > 0$ is a constant, and $(\gamma_n) \in \mathbf{R}^+$ is a sequence, then

$$\hat{\lambda}_n - \lambda_n = O_P(\gamma_n^2 n^{-1}).$$

Proof. Theorem 1 in [4].

We shall consider the seminorm $\|\cdot\|$ where

$$\|\nu\| = \sup\{|\nu(x)| : x \in \mathbf{R}\}$$

is the Kolmogorov-Smirnov norm. Further, assume the model $P_n = P_0$ and $Q_n = Q_0$ is fixed for all $n \geq 2$ with different continuous cumulative density functions P_0 and Q_0 on $\mathbf{R}$. Here one may take $\gamma_n = 1$ and $H = Q - P$.

The key for understanding the limiting behavior of estimators, confidence sets, etc. is the limiting behavior of the process $\tilde{W}_n(t) : t \in T_n$, where

$$\tilde{W}_n(t) = n\left(N_n(D_n^t)^2 - N_n(D_n^{\lambda_n})^2\right).$$

For example one may write

$$\hat{\lambda}_n = \arg\max(\tilde{W}_n(t) : \ t \in T_n).$$

Proposition The limiting behavior of $\tilde{W}_n(\cdot)$. *Under regularity conditions below, there is a stochastic process $(W(z) : z \in \mathbf{Z})$ such that for arbitrary integers $d > 0$ the vector*

$$(\tilde{W}_n(\lambda_n + n^{-1}z) : -d \leq z \leq d)$$

converges in distribution to

$$(2\mu W(z) - \mu^2|z| : -d \leq z \leq d)$$

as $n \to \infty$. Here $\mu = \|H\|$, $W(0) = 0$ and

$$W(z) = \begin{cases} \sum_{i=z+1}^{0} [g(X_i) - E_{X\sim P_0} g(X)], & (X_i \sim P_0) \quad for\ z < 0, \\ -\sum_{i=z+1}^{0} [g(X_i) - E_{X\sim Q_0} g(X)], & (X_i \sim Q_0) \quad for\ z > 0, \end{cases}$$

where $g : \mathbf{R} \to \mathbf{R}$ is a certain bounded function.

Regularity conditions and definitions: Assume that there is a unique $x_0 \in \mathbf{R}$ such that

$$|H(x_0)| = \|H\| > 0.$$

The function g is given by

$$g(x) = \mathrm{sign}(H(x_0))\chi(x \leq x_0).$$

Proof. Proposition 2 in [4].

Now the Proposition, and Theorem 7 imply the following theorem on the limiting distribution of the estimators $\hat{\lambda}_n$.

Theorem 8 *The variable $n(\hat{\lambda}_n - \lambda_n)$ converges in distribution to*

$$\arg\max\,(2\mu W(z) - \mu^2|z| : \; z \in \mathbf{Z}),$$

provided the latter is unique almost surely.

Proof. Theorem 2 in [4].

The distribution of the discrete-time process $W(\cdot)$ depends on the (infinite-dimensional) parameter (λ, P_0, Q_0), and it cannot be approximated by a continuous-time Gaussian process. For that reason Dümbgen proposes a bootstrap procedure. For every hypothetical change-point $t \in T_n$, we test the hypothesis "$\lambda_n = t$" on the nominal level α using the statistic

$$M_n(t) = max_{s \in T_n} \tilde{W}_n(s,t),$$

where

$$\tilde{W}_n(s,t) = n\left(N_n(D_n^s)^2 - N_n(D_n^t)^2\right).$$

Here we estimate the hypothetical (unconditional) distribution $L_n(t)$ of $M_n(t)$ by

$$\hat{L}_n(t) = \mathcal{L}(M_n(t, X_1^*(t), X_2^*(t), \ldots, X_n^*(t))|X_{1,n}, X_{2,n}, \ldots, X_{n,n}),$$

where $M_n(t, X_1^*(t), X_2^*(t), \ldots, X_n^*(t))$ is computed as $M_n(t)$ using resampling variables $X_1^*(t), X_2^*(t), \ldots, X_n^*(t)$. Given $(X_{1,n}, X_{2,n}, \ldots, X_{n,n})$, these variables $X_i^*(t)$ are independent, havin distribution ${}^tP_n$ for $i \leq nt$ and P_n^t otherwise. Then we reject the hypothesis "$\lambda_n = t$", if

$$\hat{L}_n(t)[M_n(t), +\infty) \leq \alpha.$$

The confidence set $K_n(\alpha)$ consists of all $t \in T_n$, where the bootstrap test did not reject. According to Dümbgen [4], this confidence set has asymptotic coverage probability not less than $1 - \alpha$ if

$$\limsup_{n\to\infty} P(\hat{L}_n(\lambda_n)[M_n(\lambda_n), +\infty) \leq \alpha) \leq \alpha.$$

Theorem 9 *Suppose that the assumptions of the Proposition hold. Then we have for any $\alpha \in (0,1)$:*

$$\max_{t\in K_n(\alpha)} |t - \lambda_n| = O_P(\gamma_n^2 n^{-1}).$$

Proof. Theorem 5 in [4].

Ferger in [5] examines hypothesis tests

$$H_0: \lambda \in \Lambda_0 \quad \text{versus} \quad H_1: \lambda \in (0,1)/\Lambda_0$$

because sometimes it is more natural to ask whether the change occurred in a certain time period or not. Like Dümbgen, he also starts out from ${}^tP_n$ and P_n^t, but instead of these, he uses the following modified empirical measures of the two subsamples generated by $t \in (0,1)$:

$${}^t\mu_n = n^{-1} \sum_{1\leq i\leq nt} \delta_{X_{in}} \quad \text{and} \quad \mu_n^t = n^{-1} \sum_{nt\leq i\leq n} \delta_{X_{in}}.$$

Then he uses integrating by product measure for the distance between the distributions (like Csörgő and Horváth in [3]):

$$r(t) = \int K \, d\mu^t \otimes {}^t\mu.$$

Here $K(x,y)$ is a measurable kernel. If we define the weight function as

$$w(t) = t^{-a}(1-t)^{-b} \quad 0 \leq a, b < 1,$$

and assume that K is antisymmetric. In that case λ is the unique maximizer(, or minimizer) of $r(t)$. The weighted function

$$\rho(t) = w(t)r(t), \quad 0 \leq t \leq 1,$$

has roughly the same shape as $r(t)$. Consequently Ferger defines the estimator λ_n of λ as:

$$\hat{\lambda}_n = \arg\max w(t)|r_n(t)| : \ t \in T_n$$

with

$$r_n(t) = \int K \, d\mu_n^t \otimes {}^t\mu_n \quad 0 \le t \le 1,$$

and T_n is the same as previously. Under some assumptions on $\rho(t)$ (in [5]) H_0 is rejected if the distance of $\hat{\lambda}_n$ and Λ_0 is too large:

$$\tau_n = \chi(\hat{\lambda}_n \in (0,1)/\Lambda_0^{\epsilon/n}),$$

where for $A \subseteq [0,1]$ and $\epsilon > 0$, A^ϵ denotes the ϵ-neighborhood of A : $A^\epsilon = \{x \in [0,1] : d(x,A) < \epsilon\}$ with $d(x,A) = \inf\{|x-a| : a \in A\}$.

Theorem 10 *Under the assumptions of Theorem 1.1 in [5]* $\alpha \in (0,1)$,

$$\limsup_{n\to\infty} P_\lambda(\tau_n reject H_0) \le \alpha \quad \forall\lambda \in \Lambda_0$$

and

$$\lim_{n\to\infty} P_\lambda(\tau_n reject H_0) = 1 \quad \forall\lambda \in (0,1)/\overline{\Lambda}_0,$$

where $\overline{\Lambda}_0$ *denotes the topological closure of* Λ_0.

2.6 Sliding intervals for change-point detection

As a conclusion of the previous chapter, we can state that the Kolmogorov-Smirnov distance of the empirical distributions of two subsamples is suitable for finding the location of the change-point. The drawback of the estimators, and hypothesis tests that were outlined previously is that we need a fixed interval with an only change-point to investigate. I used a slightly different approach to find the change-point in the series. Here I compared the empirical distribution of two neighbouring intervals and let the two intervals slide together along the series. In this way, we get another series, lets call it test-series. (The comparsion was done by Kolmogorov-Smirnov test of the subsamples.)

A previous problem remains, that is, how to find the length of the intervals. On the other hand, this method is very sensitive to the changes of the empirical distribution, so instead of searching a “stationary period”, I prefer to search the edges of possible periods. This is done by observing values that are much more significant than their neighbourhood.

We had some difficulties in using the inverse Kolmogorov function, so we used the plot of $D_{N,N}$'s, not $K^{-1}(D_{N,N})$. In fact this serves as good as the other, because now we are interested in local maximums, and not in exact possibilities, and any distribution function is monotone non-decreasing. On the other hand, we would not have to have significant results proving the similarity of two empirical distributions, because it is the network which changes its state and so the possible parameters of the distribution.

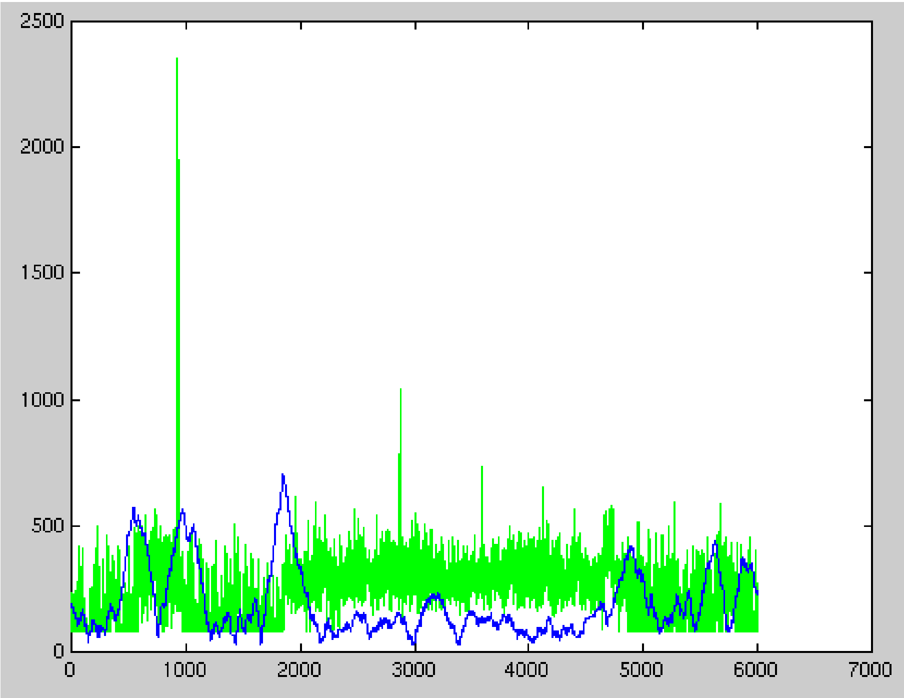

Figure 2.6: Result of the test together with the original data when it's a longer period

On Figure 2.6 you can see a period that can be thought to be appropriate for fitting parametrical distribution, and the two strong local maximums at the edges of the period. The time scale is valid for the original data. In this case the length of the sampling intervals was 250 data. The resulting series has been rescaled in order to be able to view the two series at the same plot. The conclusion is that if a period is much longer than the sampling length, one can observe the edges quite clearly.

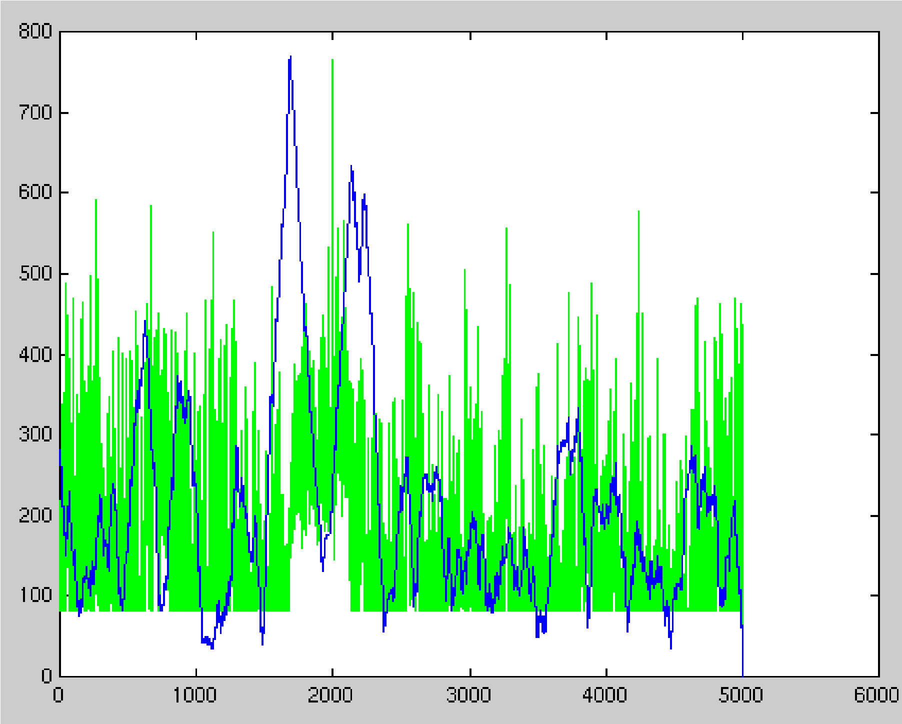

Figure 2.7: Test of a shorter period

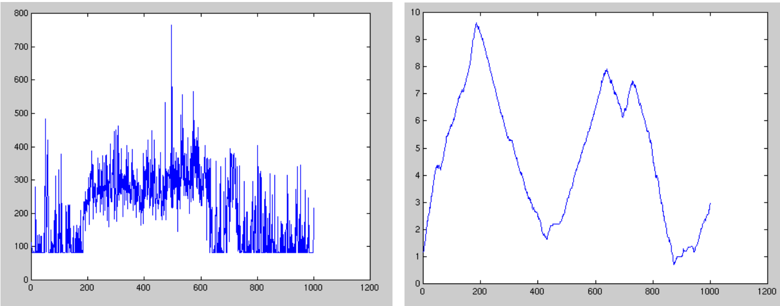

Figure 2.8: The interval on a smaller time scale

Figure 2.7 shows a shorter period, which has less than twice the length of the sampling interval's length, while Figure 2.8 has larger detail. Fortunately we can recognize the sharp edges here. The reason is, I think, that the distribution changes suddenly, and the distributions remains approximately constant on both sides of the edges.

If the length of a "calm" period is approximately equal to the length of the sampling intervals, there are problems with this approach.

Despite the difficulties mentioned, a possible change-point detection method is quite straightforward. We can consider a point to be a change-point if the test-series

- is monotone increasing before the point at least for the length of one-tenth the window size except for several percent,
- is monotone decreasing after the point also tenth the window size long except several percent of this length,
- has local maximum in that point.
- if there are more than one maximum close to each other, we take the average of them

Two change-points may bound an interval, if the length of the interval is greater or equal to the length of the window size.

I used this method to find possible intervals that are appropriate to examine for a probability distribution, and I used moving averages to avoid the longer trends that may remain between the change-points.

2.7 Fitting Parametric Distribution To Data

2.7.1 Remarks about the expected goodness of fit

In the previous section I tried to find intervals that are suitable in that distributions have been fitted to them. Unfortunately this can't be performed automatically, so the results may suffer from the conditions of choosing. The larger part of a daily series is not examined for parametric distributions, because the distribution (the average, empirical variance etc.) changed rather quickly.

I also tried to estimate the autocorrelation function of the "calm" intervals. As one can see, the estimation shows independence at longer lags, but suggests dependence before four or five lags, that is, the packet delays are correlated within five seconds. I have doubts concerning this dependence for two main reasons:

- *Theoretically* there must be no recognizable dependence of several seconds of queue length, because the size of the congestion window that is approximately equal to the number of the packets of the connection in the queue changes very quckly, and they may depend on 100-200 ms having passed.

- *Practically* if I estimated the autocorrelation function of a smaller period of data, I should observe the same correlation pattern, and obviously on Figure 2.10 that is not the case.

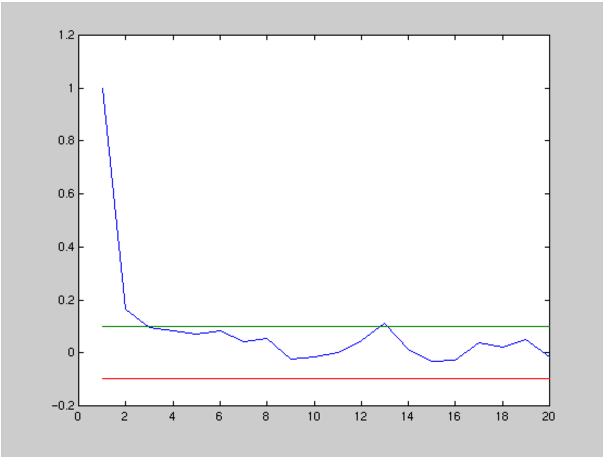

Figure 2.9: The autocorrelation function of a calm period

The estimated autocorrelation function of a series where the parameters have small changes and the series has small trend may exhibit phenomena like this. I think the problem is that we cannot find *clear* periods. This raises further doubts about the goodness of fit here. This was the reason that I did not try to use the statistical hypothesis tests to validate my estimations.

2.7.2 Descriptives, histograms, and suggestions

In order to set up a parametric model to the distributions of the chosen periods, we have to find good methods for the exploratory data analysis. Parzen in [6] advises quantile functions and density quantile functions in addition to the empirical distribution function. The quantile function $Q(u)$,

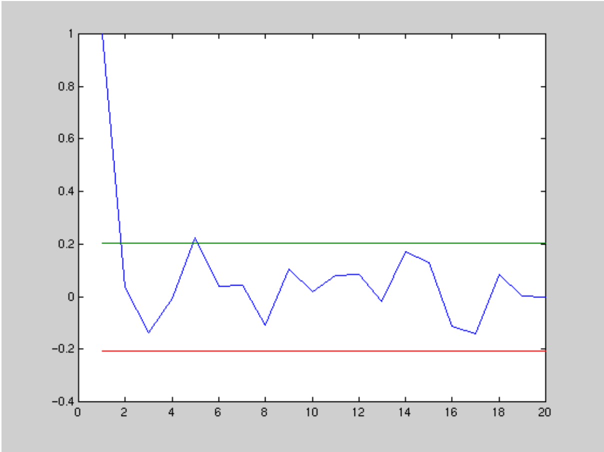

Figure 2.10: The autocorrelation function of a short interval of a calm period

$0 \leq u \leq 1$ is defined as follows: for a general distribution function $F(x)$ which is continuous from the right,

$$Q(u) = F^{-1}(u) = \inf\,(x: \; F(x) \geq u).$$

For $-\infty < x < \infty$ and $0 < u < 1$

$$F(x) \geq u \quad \leftrightarrow \quad Q(u) \leq x.$$

Since

$$P(Q(U) \leq x) = P(U \leq F(x)) = F(x),$$

where U has uniform distribution in $[0, 1]$,

$$Q(U) \sim F, \tag{2.19}$$

that is, $Q(U)$ has the distribution function F. When F continuous, Q satisfies

$$Q(u) = \inf\,(x: \; F(x) = u)$$

$$FQ(u) = F(Q(u)) = u, \quad 0 \leq u \leq 1.$$

Assume that F is differentiable in its range, and define

$$fQ(u) = f(Q(u)) \quad and \quad q(u) = Q'(u)$$

which Parzen calls density-quantile and quantile-density function. If we differentiate FQ, we obtain

$$fQ(u)q(u) = 1.$$

If we know Q, we can simulate (Equation 2.19) and compute $Eg(X)$:

$$Eg(X) = Eg(Q(U)) = \int_0^1 g(Q(u))du. \tag{2.20}$$

An interesting consequence of Equation 2.20 is

$$\int_0^1 \log fQ(u)du = \int_{-\infty}^{\infty} f(x)\log f(x)dx.$$

Denote $\hat{F}$ the empirical distribution function. Then denote $\hat{Q}$ the empirical quantile function, which is:

$$\hat{Q}(u) = \inf(x:\ \hat{F}(x \geq u)).$$

It is a piecewise constant function whose values are the order statistics:

$$\hat{Q}(u) = X_{(j)} \quad \text{for} \quad (j-1)/n < u \leq j/n, j = 1, \ldots, n$$

where $X_{(i)}$ is the i-th value:

$$X_{(1)} < X_{(2)} < \ldots < X_{(n)}.$$

For $u = 0$ define $\hat{Q}(0)$ be the sample minimum.

On Figure 2.11 there are three different empirical quantile functions corresponding to three periods with different traffic in each. Later I try to draw the appropriate histograms to these periods.

The minimal value of the daily series is 64 ms. The series takes this value quite often during the night. This is the propagation delay, that is, it takes physically 64 milliseconds to have the packet sent, for it to arrive, and "turned back", and delivered back to us. We have three main types of histograms.

The main difference between the histograms is the number and location of local maximums. As Figure 2.12 shows, there is one local maximum in the slack period, namely the propagation delay, and there is also one local maximum in the busy period, which is approximately equal to the average delay,

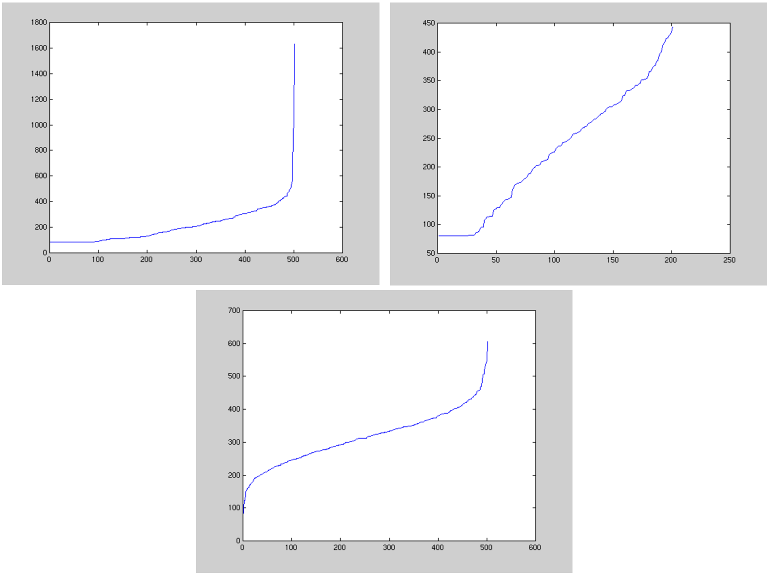

Figure 2.11: The quantile function of the rare, medium, and busy period

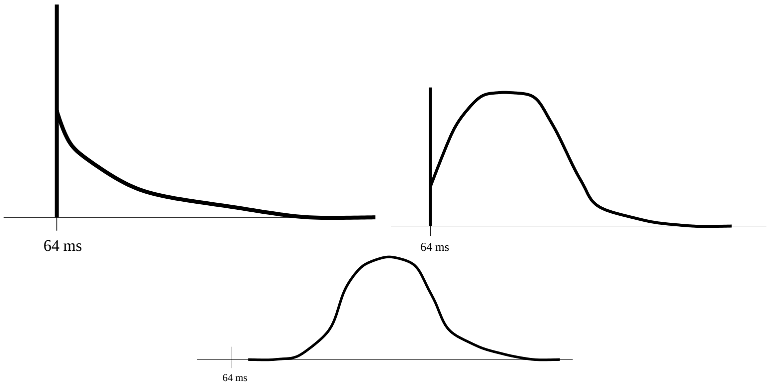

Figure 2.12: Roughly outlined histograms of the rare, medium, and busy period

because the histogram is nearly symmetrical in the busy period. However, there are two local maximums in the medium period.

The symmetrical, bell-shaped curvature of the histogram in busy periods suggested normality. The quickest way of checking this is the quantile-

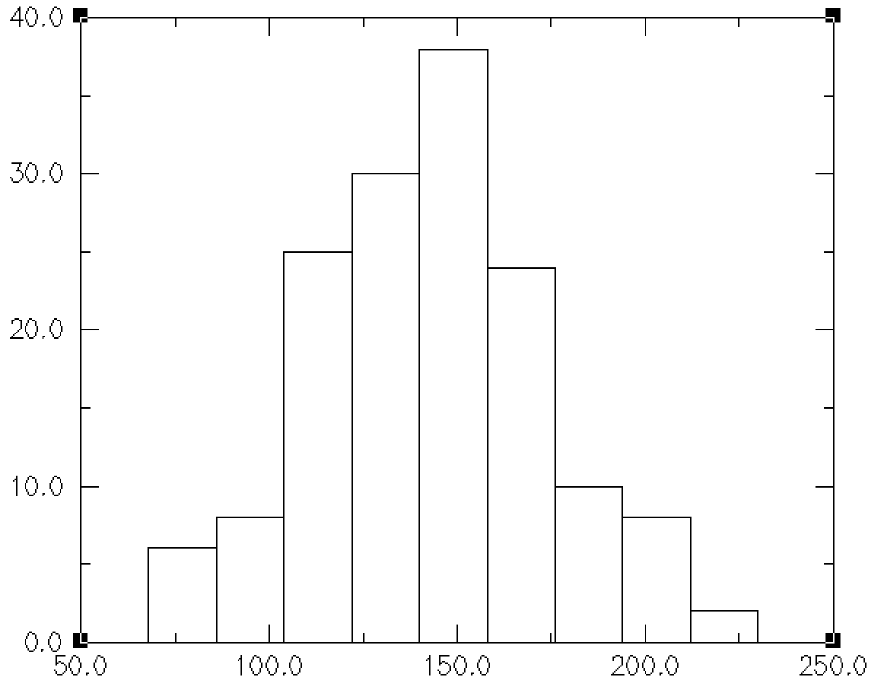

Figure 2.13: The histogram of the busy period

quantile plot which will be detailed in Section 2.7.4. On Figure 2.14 there is the Q-Q plot of a busy period. Here the sample average and the empirical square of variance were used as estimators of the mean and variance of the normal distribution.

2.7.3 Estimating parameters of cut-normal distribution

Let the quantile be α such that

$$\alpha = \Phi(\frac{c-\mu}{\sigma}), \tag{2.21}$$

where c equals the censoring level.

Here we assume that the originating distribution is Normal(μ, σ^2).

$$\mathrm{E}(X) = \int_c^{+\infty} t \frac{1}{\sqrt{2\pi}\sigma(1-\alpha)} e^{-\frac{(t-\mu)^2}{2\sigma^2}}\, dt =$$

$$\frac{1}{1-\alpha} \int_{\frac{c-\mu}{\sigma}}^{+\infty} (\sigma u + \mu) \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}\, du =$$

$$= \frac{\sigma}{1-\alpha} \int_{\frac{c-\mu}{\sigma}}^{+\infty} u \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}\, du \;+\; \frac{\mu}{1-\alpha} \int_{\frac{c-\mu}{\sigma}}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}\, du =$$

$$= \frac{\sigma}{1-\alpha} \frac{1}{\sqrt{2\pi}} e^{-\frac{(c-\mu)^2}{2\sigma^2}} \;+\; \frac{\mu}{1-\alpha} \left(1 - \Phi(\frac{c-\mu}{\sigma})\right). \tag{2.22}$$

If we definine

$$a = \frac{1}{\sqrt{2\pi}(1-\alpha)} e^{-\frac{(c-\mu)^2}{2\sigma^2}} = \frac{1}{\sqrt{2\pi}(1-\alpha)} e^{-\frac{\left(\Phi^{-1}(\alpha)\right)^2}{2}},$$

where we used α for a.

By this definition and from Equation 2.22 we get 2.23:

$$\mathrm{E}(X) = \sigma a + \mu, \tag{2.23}$$

$$\mu = c - \sigma \Phi^{-1}(\alpha). \tag{2.24}$$

From Equation 2.23 and 2.24 we obtain the following estimation for σ

$$\hat{\sigma} = \frac{\mathrm{E}(X) - c}{a - \Phi^{-1}(\alpha)}, \tag{2.25}$$

and substituting σ in Equation 2.24 we get $\hat{\mu}$.

Finally we estimate α and $\mathrm{E}(X)$ such that

$$\hat{\alpha} = \sum_{i=1}^{N} \chi(X_i < c), \tag{2.26}$$

$$\overline{X} = \frac{1}{N} \sum_{i=1}^{N} X_i. \tag{2.27}$$

Then we are ready to estimate $\hat{\sigma}$, and $\hat{\mu}$ by Equation 2.23, and 2.24.

2.7.4 Validating

Q-Q plot is a clear, and graphic, but not so exact way of matching the empirical distributions of two samples. We have a function $q_{(X_1,\ldots),(Y_1,\ldots)}$ such that

$$q_{(X_1,\ldots),(Y_1,\ldots)} : \mathbf{R} \to \mathbf{R}$$

and if the graph of this function fits the graph of the $u(x) = x$ function well, then the two distributions are quite similar. However if there is curvature in the $q(\alpha)$ function, we are able to determine where the difference can be. Let us define the $q(\cdot)$ function as follows.

Consider the estimated qunatile function $Q_X(u)$ from the sample

$$X_1,\, X_2,\, \ldots\, X_m$$

and $Q_Y(u)$ from the sample

$$Y_1,\, Y_2,\, \ldots\, Y_n.$$

Now let $q(\cdot)$ be the following

$$q(Q_X(u)) = Q_y(u),$$

that is

$$q(x) = Q_Y Q_X^-(x), \tag{2.28}$$

where

$$Q_X^-(x) = \frac{1}{m}\sum_{i=1}^{m} \chi(X_i \le x).$$

There are some advised restrictions on plotting $q(\alpha)$. For the sake of good visibility, one should evaluate fewer or an equal number of α's than the minimum of m and n, and use rather identical increases in α.

In order to show that the assumtions that were made when we saw the qunatile functions - and the histograms - of the periods hold, we plot the quantile-quantile plot of the measured data, and the fitted distributions.

2.7.5 Consequences

Although I shall go into more detail in the next chapter, I must make several remarks in advance. TCP traffic strongly differs from those traffic types -

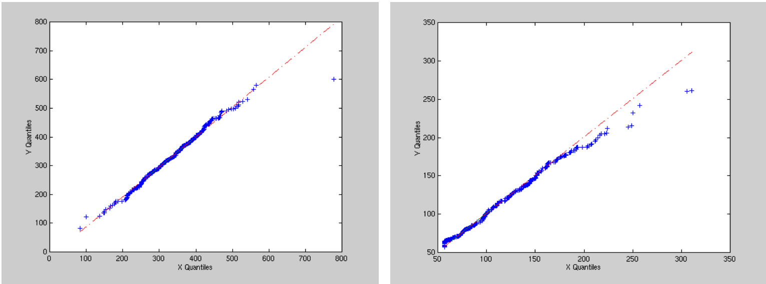

Figure 2.14: Q-Q plot with an estimated normal, and cut-normal distribution

in a wider sense - that we could observe so far. (For example road traffic, arrivals to a shop or to a telephone exchange, customers' demands.)

Here - though we have a daily profile - the governing distribution changes suddenly, beacuse of the properties of the TCP connection. When a new connection is established, it explores the available capacities. This causes considerably increasing traffic - in slow start state the connection increases its traffic exponentially. This is the reason why change-point detection can be used to examine the series and search for the begining of connections.

Another important attribute is the normality of the queue length. The conventional queue length distributions are rather more exponential, or in a discrete case, more negative binomial, than normal. The reason for this difference is that in our case the packet arrivals are not independent of the past, furthermore they may have very strong correlations. To obtain suggestions about how we can get normal-like queue length, I try to use a bit different approach than simply examining the packet arrivals.

As I mentioned above, the size of the congestion window is the number of packets in the network that are unacknowledged. This implies that the number of packets in the bottleneck queue plus the number of packets in the "wire" elsewhere gives us the congestion window.

In fact the number of sequential packets and the acknowledgements togather equals the size of the congestion window. However, the proportion of the sequential packets and acknowledgements is more or less constant, because it depends on the route to and back from the receiving computer.

Assuming that the "wire" contains a fixed number of packets, we get that the addition of a single connection to the queue length in the bottleneck is related to the congestion window, which is controlled by the sender TCP, and therefore somehow determinative. Assuming either that the congestion windows of different TCP connections are independent, we can use the central limit theorem of the sum of independent, identically distributed random variables. Thus the number of packets of the different conncetions in the queue and the wire together may be iid random variables, and their sum is normally distributed. In this case, queueing occurs when the wire is filled with packets, and so we get the appropriate cut normal distribution.

Chapter 3

Arrival Series

In the previous chapter I examined the packet delay which in our case relates to the length of the bottleneck buffer. There I assumed that I was sampling the buffer seldom enough to obtain independence. However, it may be interesting to see what happens at packet level time, that is, between two samplings.

I measured the packet arrival times of a TCP connection, where we had infinitely much data to send. At the same time, I collected some parts of the sequence packets, and the acknowledgements, so that I was able to calculate the size of the congestion window. In this way I had two time series. One is the arrival times of the packets, and the other the size of the congestion window at these times.

I counted the interarrival times of the acknowledgements, because the the receiver acknowlegdes two packets in the same acknowledgement packet, therefore the sender sends two more packet when receiving that packet. Sometimes, however, the receiver does not wait for the second packet to be sent an acknowlegdement, and in this case the sender sends only one packet as well. Conseqently, they are the interarrival times of the acknowledgements that are reliable, not the interarrival times of the sequence packets.

3.1 General description

At the first glance on the interarrival times (Figure 3.1) one can see that they differ from the delay series very much. Here it is typical that the series starts from a large value, then decreases till it reaches a pretty small value.

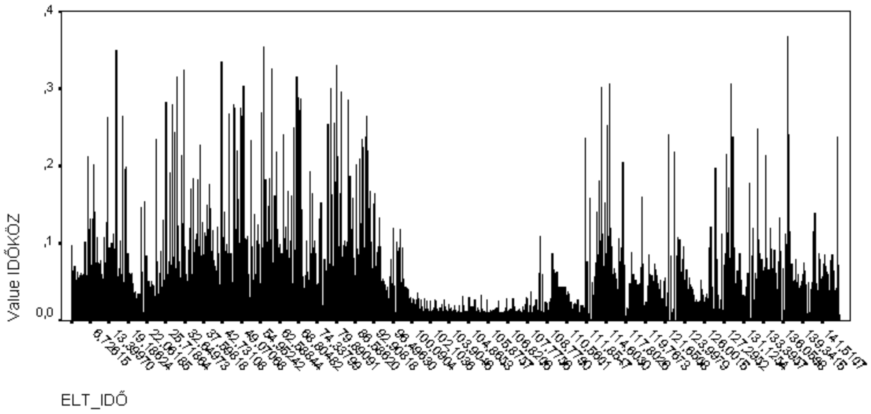

Figure 3.1: The interarrival series

After that it goes on with a much larger value, as previously (see Figure 3.7). Sometimes this periodical behaviour changes, and the interarrival time remains constant for a certain period. These are more visible on the smoothed series (Figure 3.8).

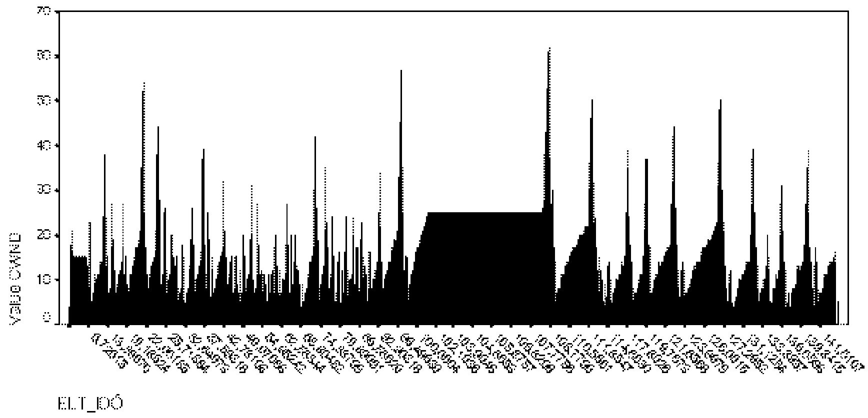

Figure 3.2: The congestion window series

The congestion window series (Figure 3.2 or Figure 3.9) is a bit different. Here we can also see the sesonality in the previous sense, but on a reversed way. The series starts from a small value, and gradually increases when the (moving average of the) other series decreases. This implies that the two series may be negatively correlated, and indeed this is the case.

3.2 The stabilized state

I wrote in the previous section that the series shows two different modes of behaviour. One is the oscillation, the other the approximately fixed average and lack of oscillaction in the previous sense. The reason of the second state is the following.

The congestion window oscillates because of the specification of TCP. TCP constantly tries to increase it until packet loss occurs. Then it reduces the congestion window and starts to increase it again. The size of the congestion window, however, is upper bounded by the advertised window of the receiver. This number expresses the space that the receiver allocated to store the received packets before the application takes it away from the TCP. So if there is no packet loss in the network, the sender stops increasing its congestion window when it reaches the size of the advertised window.

When the congestion window remains fixed, TCP sends a packet when it receives an acknowledgement, and doesn't send any packet because of something else. This means that the sending rate remains constant, and doesn't increase either. This is the stabilized state.

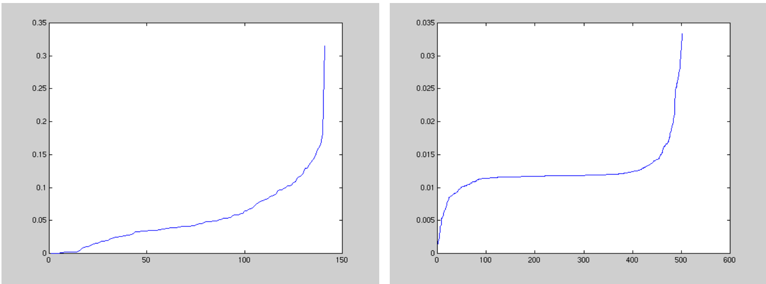

Figure 3.3: The qunatile function of a stable state

On Figure 3.3 and 3.4 I plotted the quantile functions and histograms of two different series, when they are in stable state. The difference between them is the size of the sequential packets and the time when the data were collected. The first series comes from a measurement at around midnight,

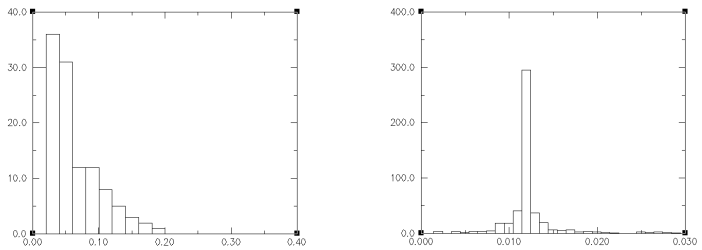

Figure 3.4: The histogram of a stable state

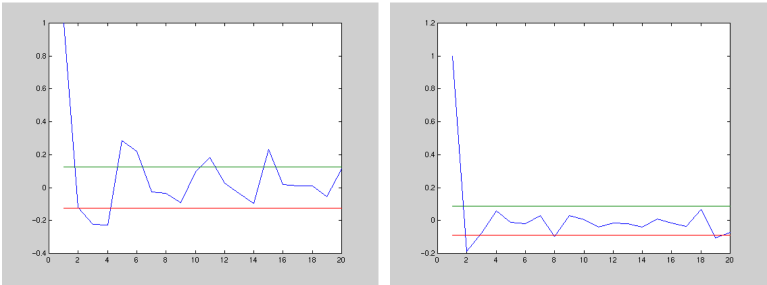

Figure 3.5: The autocorrelation function of a stable state

when the traffic was the only measurement-generated traffic, while the second series was measured at 10 in the morning under busier traffic circumstances.

The reason I changed the size of the packets was that the size of the advertised window was fixed - actually it was 8096 bytes - and I wanted to observe different numbers of packets that a congestion window contains.

In the first case, the packet size was 1024 bytes, which means that 8 packet yield a congestion window. However I measured the interarrival times of the acknowledgements, and because in the stable state each acknowledgement was delayed each allows two more packets to be sent. Therefore in the first case 4 acknowledgement yields one congestion window.

In the second case, the packet size was 512 bytes, where 16 packet yield a congestion window, and the number of acknowledgements is 8. There are the corresponding autocorrelation functions on Figure 3.5. I shall go into more details in the next section.

These functions show several interesting effects of the packet switched networks. One of these is the short-range negative correlation, another is the repeated positive correlation on longer lags. The reason is explained in Figure 3.6.

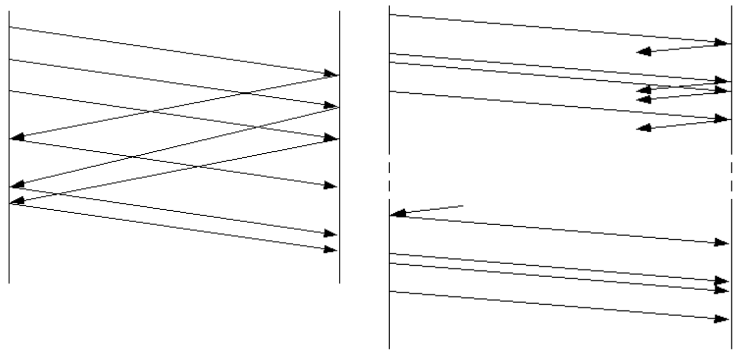

Figure 3.6: The explanation for the correlation pattern

Assume a sequence of three packets that follow each other at equal distances. Then assume that the second acknowledgement is delayed, but the third is not. The enlarged interval between the first and second acknowledgement is followed by an interval, which should be smaller than the average. If later there is no delay, this pattern remains.

The explanation for the repeated positive correlation is the following. If the number of acknowledgements of a congestion window is small enough and there is no buffering, then the interarrival times of the packets remain more or less strictly the same the next time they cause additional packets to be sent.

In the first case there was no traffic, and 4 packets completed the congestion window. We can see the periods on the corresponding autocorrelation function figure. In the second case the similar number was 8, moreover, there was traffic, and buffering occured, and so the estimated autocorrelation function does not show this periodicity.

3.3 The connection between the two series

Recall that the congestion window is the number of bytes that were sent out to the network, and unacknowledged. Consider that the interarrival time of two packets, where there is a congestion window - many packets between them, relates approximately to the packet delay in the network. The explanation is the following. When we send a packet, it flows through the network, becomes accepted, and the acknowledgement to this packet comes back to the sender. The time of this sequence of events is the round-trip time, or the packet delay. On the other hand, after sending our packet and before receiving its acknowledgement we get acknowledgements for the whole current congestion window. In acknowledgements this means that we get as many packet, as the size of the congestion window allows.

Unfortunately, this is not very exact statement, because the congestion window is expressed in bytes, not in packets. Other problem is that the receiver answers every second sequence packet, except some, so the expression "congestion window-many" in sequence packets means the number of packets N that were sent and unacknowledged so far. Let the number M be that amount of packets (acknowledgements) that we shall receive until we get the answer to the appropriate sequenece packet. That is

$$cwnd = \overline{seq}N \approx 2\overline{seq}M,$$

where $\overline{seq}$ is the (average) size of a sequence packet.

Let $X_i, \quad i = 1, 2, \ldots$ be the interarrival series of acknowledgements, and $rtt_i, \quad i = 1, 2, \ldots$ be the round trip times of the triggered packets. Then

$$rtt_i = \sum_{j=1}^{M_i} X_{i+j}.$$

On the other hand, rtt_i can be decomposed into several factors:

$$rtt_i = prop_i^{\mathrm{SEQ}} + blen_i + prop_i^{\mathrm{ACK}},$$

here:

$prop_i^{\mathrm{SEQ}}$ is the propagation delay of the i-th packet before it arrives to the receiver

$blen_i$ is the length of the buffer in service time, when the i-th packet arrives

$prop_i^{\mathrm{ACK}}$ is the propagation delay of the i-th acknowledgement back from the receiver

It is not too strong assumption that the propagation delays are fixed. I shall assume this for simplicity and if there is no other traffic, then this holds. Therefore

$$rtt_i = blen_i + c.$$

We need some more definitions

seq_i the number of the sequence packets in the buffer when the i-th packet arrives

$cseq$ the number of the sequence packets in the "wire", which is constant

ack_i the (average) number of the acknowledgemnts when the i-th packet flows through the network

$cwnd_i$ the size of the congestion window when the i-th packet was sent

$\overline{X}_i$ is the average value of i-th and the several next interarrival times

Then

$$cwnd_i \approx \overline{seq}(seq_i + cseq + ack_i),$$

and

$$ack_i \approx (seq_{i-K} + cseq)l,$$

that is the number of the acknowledgements depends linearly from a previous number of sequence packets in the network. Then we get

$$cwnd_i \approx \overline{seq}(seq_i + cseq + (seq_{i-K} + cseq)l).$$

Now

$$blen_i = seq_i s$$

where s is the service time in the buffer.

$$M_i = \frac{cwnd_i}{2\overline{seq}} \approx \frac{seq_i + cseq + (seq_{i-K} + cseq)l}{2},$$

$$blen_i = s(2M_i - cseq - (seq_{i-K} + cseq)l),$$

$$rtt_i \approx M_i \overline{X}_i,$$

$$rtt_i \approx s(2M_i - cseq - (seq_{i-K} + cseq)l) + c,$$

$$\overline{X_i} \approx 2s + \frac{c - s(cseq + seq_{i-K} l)}{M_i}.$$

If we assume that

$$seq_i = l^* seq_{i-K},$$

then

$$M_i = \frac{seq_i + \frac{l}{l^*} seq_i + cseq}{2} = \frac{\left(1 + \frac{l}{l^*}\right) seq_i + cseq}{2},$$

$$M_i \overline{X}_i \approx \frac{2M_i - cseq}{1 + \frac{l}{l^*}} s + c = \frac{2s}{1 + \frac{l}{l^*}} M_i - \frac{s\ cseq}{1 + \frac{l}{l^*}} + c,$$

finally

$$\overline{X}_i \approx \frac{2s}{1 + \frac{l}{l^*}} + \frac{c - \frac{s\ cseq}{1 + \frac{l}{l^*}}}{M_i} = \mathrm{const}_1 + \frac{\mathrm{const}_2}{M_i}. \tag{3.1}$$

Here is the explanation for the negative correlation result. Equation 3.1 implies that the larger congestion window we have, the smaller the interarrival times they are. Actually Equation 3.1 suggests linear dependence, which I try to check now. First recall the "detailed" interval of the series, and the result of smoothing on the same period.

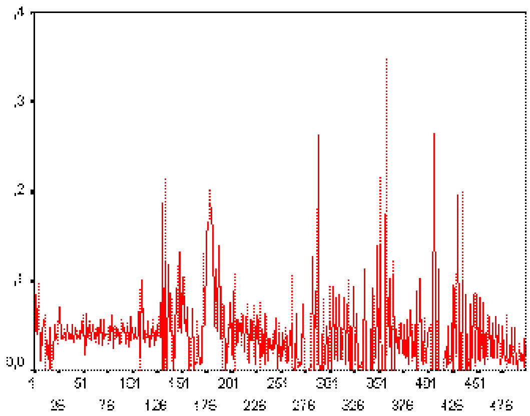

Figure 3.7: The first 500 member of an interarrival series

Figure 3.9 shows the congestion window series of this interval.

I used the program SPSS 7.5 to estimate the parameters of the linear regression by the model

$$\overline{X}_i = \mathrm{const}_1 + \frac{\mathrm{const}_2}{M_i} + \mathrm{error}_i,$$

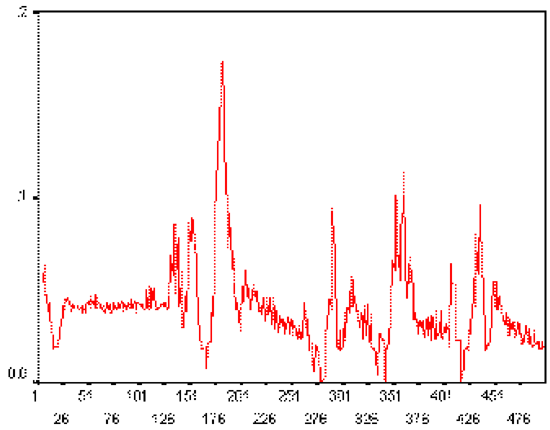

Figure 3.8: The smoothed series - prior moving average at lag 5

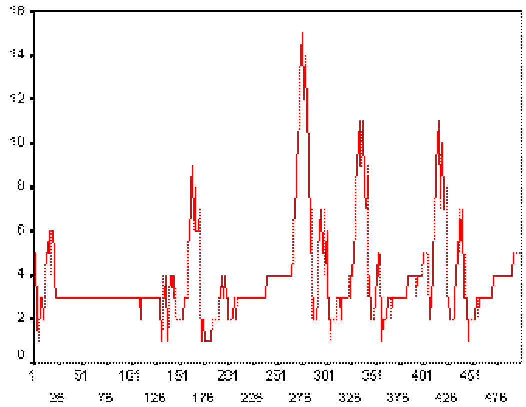

Figure 3.9: The first 500 member of the corresponding congestion window series

where in our case

$$\overline{X}_i = \frac{1}{5}\sum_{j=0}^{4} X_{i+j}$$

The result is the following:

$$\text{const}_1 = 8.15 \cdot 10^{-4}, \quad \text{const}_2 = 0.1159.$$

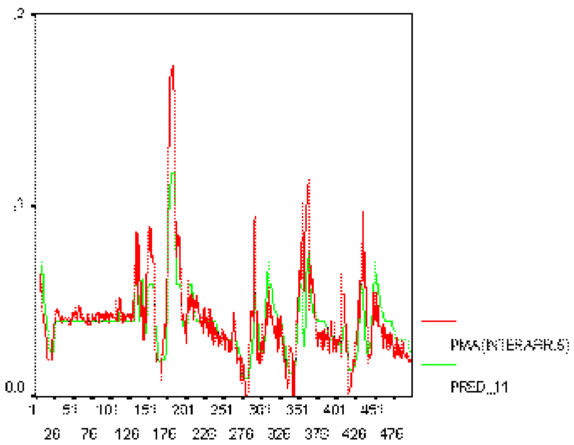

Figure 3.10: The interarrival series, and the result of the regression together

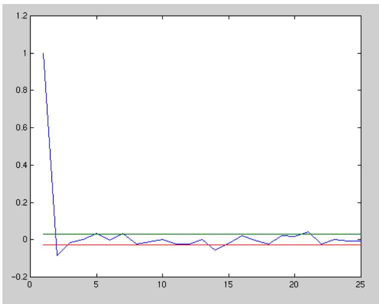

Figure 3.11: The autocorrelation function of the error series

Figure 3.10 shows the series

$$\text{const}_1 + \frac{\text{const}_2}{M_i}, \quad \text{and} \quad X_i$$

and Figure 3.11 shows the estimated autocorrelation function of the error series.

Bibliography

[1] D. C. Lynch: Historical Perspective (Internet System Handbook 1993, 3-14)

[2] W. Richard Stevens: TCP/IP Illustrated, Volume 1 (Addison-Wesley professonal computing series)

[3] Miklós Csörgő and Lajos Horváth: Invariance Principles for Changepoint Problems (J. of Multivariate Analysis 27, 151-168)

[4] L. Dümbgen: The Asymptotic Behavior of Some Nonparametric Change-point Estimators (The Annals of Stat. 19, 1471-1495)

[5] D. Ferger: Nonparametric Tests fir Bibstabdart Change-point Problems (The Annals of Stat. 23, 1848-1861)

[6] Emanuel Parzen: Nonparametric Statistical Data Modeling (Journal of the American Statistical Association, Marc 1979, 105-126)