

Traffic Characterization

Due Date: Due dates are posted on the web site.

Teams: This lab may be completed in teams of 2 students (Teams of three or more are not permitted. All members receive the same grade).

Purpose of this lab:

In this lab you will study and compare the properties of different types of network traffic. The traffic types considered are:

- Generated Poisson traffic;
- Traffic from compressed video sources; and
- Measured traffic from an Ethernet segment.

The traffic sources from this lab will be used in future labs to generate network traffic.

Software Tools:

- This lab requires the use of MATLAB.
- Note: This lab (and this lab only) requires the use of the “statistics package” of MATLAB. Since the workstations in the lab do not have the statistics package, you need to remotely login to your ECF account.
- Please TAs for further instructions.

What to turn in:

- **Turn in a hard copy of** all your answers to the questions in this lab, including the plots, hard copies of all your MATLAB code, and the anonymous feedback form.

Table of Content

Table of Content	2
Preparing for Lab 1	2
<i>Lab 1</i>	3
Part 1. Poisson Traffic	4
Part 2. Compressed Video Traffic	7
Part 3. Aggregate traffic on an Ethernet network	10
Part 4. (Optional) Metrics for Characterizing Network Traffic	13
<i>Feedback Form for Lab 1</i>	15

Preparing for Lab 1

This lab assumes that you have a basic knowledge of MATLAB. It will be helpful to review your background of MATLAB. Some of the tasks involved in this lab are:

1. Read and process a data file with multiple columns;
2. Use of statistical functions, such as mean value, variance, covariance, etc.

Lab 1

Please read the following comments:

- **Quality of plots in lab report:** This lab asks you to produce plots for a lab report. It is important that the graphs are of high quality. All plots must be properly labeled. This includes that the units on the axes of all graphs are included, and that each plot has a header line that describes the content of the graph.
- **Waiting times with reading large files:** The lab requires reading large files with traffic traces. Since the computers in the UNIX lab are quite old, there may be long delays to load a file. To avoid excessive waiting times, you may truncate the files if loading a file takes longer than 15 minutes.
- **Feedback:** To be able to improve the labs for future years, we collect data on the current lab experience. You must submit an anonymous feedback form for each lab. Please use the feedback form at the end of the lab, and return the form with your lab report.

- **Units:** The following conventions are used:

ms = milliseconds
s = seconds
Mbps = Megabits per second

- **Arrival functions:** In this lab (and in the lecture) we will sometimes describe the arrivals of traffic by a function A .
 - a. We write $A(t)$ to denote the arrivals in the time interval $[0,t]$. We will write $A(t_1, t_2)$ to denote the arrivals in the time interval $[t_1, t_2]$.
 - b. Sometimes arrivals occur only at discrete times $n= 0, 1, 2, 3, \dots$. In this case, we write $A(n)$ to denote the arrivals until n .

- **MATLAB help:**

a. **We have made available:**

- i. sample MATLAB code for Parts 1, 2, and 3.
- ii. Three traffic traces for Poisson traffic used in Part 1


b. **Download all files in:**

<http://www.comm.utoronto.ca/~jorg/teaching/ece466/labs/lab1/>

- **Part 1:** Sample code: [Lab1_Part_1.m](#)
traces: [poisson1.data](#), [poisson2.data](#), [poisson3.data](#)
- **Part 2:** Sample code: [Lab1_Part_2.m](#)
traces: [movietrace.data](#)
- **Part 3:** Sample code: [Lab1_Part_3.m](#)
traces: [Bel.data](#)

Part 1. Poisson Traffic

Suppose we want to predict the delay and throughput at a buffer of a network link that operates at link capacity C Mbps. The situation is illustrated in the figure, where packets are represented as rectangular boxes.



To conduct an analysis of this link, we need to make assumptions on the traffic that arrives to the link. The traffic description must specify the spacing between packets and the size of packets that arrive to the buffer.

A very popular model to represent the number of packet arrivals in a time interval is the Poisson Process. In a **Poisson process** with rate λ , the number of (arrival) events in a time interval $(t, t+\tau]$, denoted by $N(t+\tau) - N(t)$, is given by

$$P[N(t+\tau) - N(t) = k] = \frac{(\lambda\tau)^k}{k!} e^{-\lambda\tau}, \quad k = 0, 1, \dots,$$

A Poisson process has some interesting and useful properties:

- In a Poisson process, the time between two (arrival) events follows an exponential distribution, i.e.,

$$P[\text{Time between two events} \leq X] = 1 - e^{-\lambda X}$$

- If we add the arrivals of two Poisson processes with rate λ and μ , we obtain a Poisson process with rate $\lambda + \mu$.

The following exercises ask you to generate Poisson traffic and study its properties.

Software Tools:

- **This lab requires the use of MATLAB.**
- **Note: This lab (and this lab only) requires the use of the “statistics package” of MATLAB. Since the workstations in the lab do not have the statistics package, you need to remotely login to your ECF account.**
- **Please contact a TA for further instructions or if you have problems with accessing the statistics package.**

Exercise 1.1 Generating Poisson traffic

Consider a packet arrival pattern of a flow, where all packets have a constant size of 100 Bytes and where packet arrivals follow a Poisson process. The average traffic rate of the flow is 1Mbps.

- Determine the rate λ of this Poisson process.
- Write a (Matlab) program that can generate a sample path of the above arrival pattern.
Hint: The MATLAB function `random('poisson',x,1,z)` creates a random vector with z events for a Poisson process with rate x .
- Create a sample path of the above Poisson flow for time intervals of 100 seconds. We will refer to this sample path as the *Poisson trace*.
 - Verify the measured mean bit rate of the flow with the target rate;
 - Calculate the mean and the variance of the times between consecutive arrival events. Compare the values with the theoretically expected values.

Exercise 1.2 Scaling Poisson traffic arrivals

Here, you create sample paths of different lengths, and compare them in scaled plots.

- Consider the Poisson trace from the previous exercise.
- Create three graphs that plot the data generated by the trace, viewed at different time scales. Each graph has 100 data points.
- **Plot 1:** Generate a vector with 100 elements. Each element stores the number of bytes from the Poisson trace that arrive in a time period of 1 second.
 - 1st element: # bytes arriving in time period [0, 1 s];
 - 2nd element: # bytes arriving in time period [1, 2 s];
 - ...
- **Plot 2:** Generate a vector with 100 elements. Each element stores the number of bytes from the Poisson trace that arrive in a time period of 100 milliseconds, beginning at a randomly selected start time.
 - Pick a random starting point, e.g., time 30 s.
 - 1st element: # bytes arriving in time period [30, 30.1 s];
 - 2nd element: # bytes arriving in time period [30.1, 30.2 s];
 - ...
- **Plot 3:** Generate a vector with 100 elements. Each element stores the number of bytes from the Poisson trace that arrive in a time period of 10 milliseconds, beginning at a randomly selected start time.
 - Pick a random starting point, e.g., time 50.2 s.
 - 1st element: # bytes arriving in time period [50.2, 50.21 s];
 - 2nd element: # bytes arriving in time period [50.21, 50.22 s];

-
- Plot the content of the vectors in three separate graphs, with the time intervals on the x-axis, and the number of bytes on the y-axis. The data points should be depicted as vertical bars (e.g., using the MATLAB function `bar()`).
- Describe your observations of the graphs.

Exercise 1.3 Traffic with exponentially sized packets


Consider a packet arrival pattern of a flow, where packets arrive with an equal spacing of 800 μsec (microseconds), and with an exponentially distributed packet size with parameter with average size $1/\mu = 100$ Bytes.

- Repeat the tasks in Exercises 1.1 and 1.2 with the above arrival flows.
- Provide a discussion where you compare the results with those from Exercise 1.1 and 1.2.

Exercise 1.4 Compound Poisson arrival process

Now consider a packet arrival pattern of a flow, which is a combination of the previously discussed flows. Here, the packet arrival events follow a Poisson process with rate $\lambda = 1250$ packets/sec, and the packet size has an exponential distribution with average size $1/\mu = 100$ Bytes.

- Repeat the tasks in Exercises 1.1 and 1.2 with the above arrival flows.
- Provide a discussion where you compare the results, with those of the other exercises.



Lab Report:


- Include your answers to the tasks.
- Include the plots, a description of the plots, and your observations and discussions.

Part 2. Compressed Video Traffic

In this part of the lab, you will explore the properties of a traffic flow that consists of compressed digital video traffic.

Digital video consists of a sequence of video frames that are displayed at a rate of typically 30 frames per seconds. Assuming a frame size of 720 x 480 pixels (a standard format used for digital video) and a depth of 24 bits per pixel, this results in a data rate 240 Mbps per video stream. Using modern compression methods, such as MPEG and H.264, the data rate can be drastically reduced to a few Mbps.

An MPEG encoder (and most other compression algorithms) generates three types of frames: Intra-coded (I), Inter-coded or predictive (P), and bi-directional (B). I frames are coded as still images not unlike JPEG; P frames encode the differences from the most recent I or P frame; and B are interpolations of the next and previous I or P frames. An MPEG encoder generates these frames in a repeating pattern, called the Group-of-Picture (GOP) pattern. With 30 frames per second, the elapsed time between two frames is approximately 33 ms. A typical GOP pattern of IBBPBBPBB, where the dependencies between frames is indicated by arrows, is shown in the following figure:



As indicated in the figure, frame types have different sizes, with I frames being the largest and B frames the smallest. Note that a B frame is constructed from two reference frames of type I or P (e.g., Frame 3 above depends on Frame 1 and Frame 4). Thus, to encode or decode a B frame the transmitter or receiver must have available the next I or P frame. To account for this need, frames are transmitted in a different sequence than they are displayed. In the above example, the display and transmission sequence is as follows:

Display sequence:	1	2	3	4	5	6	7	8	9	10
Transmit sequence:	1	4	2	3	7	5	6	10	8	9

Due to the different sizes of I, P, and B frames, the data rate of an MPEG encoded video sources changes from frame to frame. For this reason, compressed video such as MPEG, is referred to a Variable Bit Rate (VBR) video.

Exercise 2.1 Download a VBR video trace

Your first task is to download a file that contains a trace of a VBR video source. The file contains the frame sizes obtained from a compression of a picture movie. The size of the file is 2.6 MB.

Content: Silence of the Lambs (~ 30 min)

Source Format:	DVD
Frame size and rate:	CIF 352x288 @30 fps (frames per second)
Compression Algorithm:	H.264/AVC (Full)
Quantizer (compression factor):	10
GoP Pattern:	I B B B P B B P P B B B P B B B
Number of Frames:	~54000

Download the trace file from the following URL:

http://www.comm.utoronto.ca/~jorg/teaching/ece466/labs/SilenceOLCIF54k_G16B3F10_verbose.data

- Explore the content of the file. The file contains one line for each frame and has multiple columns. The first four columns are relevant for us:
 - Column 1: Sequence number of the frame (in transmit sequence)
 - Column 2: Display time of the frame
 - Column 3: Frame type
 - Column 4: Frame size in bytes
 - Columns 5 – 7: not used


Exercise 2.2 Determine statistical properties of the video trace

- Compute the following properties of the video trace.
 - Number of frames and total number of bytes;
 - Size of the smallest frame, size of the largest frame, and mean frame size;
 - Size of the smallest, largest and mean I, P, and B frame;
 - Mean bit rate (Computed as the mean frame size divided by the frame duration);
 - Peak bit rate (Computed as the max. frame size divided by the frame duration);
 - Ratio of the peak rate and the average rate. This peak-to-average rate ratio is often used as an indicator how bursty a traffic flow is. A flow with a peak-to-average ratio of 10 or higher is highly variable.
- Generate a set of graphs that show the properties of the video trace:
 1. Generate a graph that shows the frame size as a function of the frame sequence number. (Use the sequence in which the frames are listed in the file, i.e., the transmit sequence.)
 2. Generate a graph that shows the distribution of I frames, P frames, and B frames. (x-axis is the frame size, y-axis is the relative frequency).

Exercise 2.3 Scaling Video traffic

As before with the Poisson traces, you are asked to generate and compare scaled plots of the video trace.

- Create three plots of the video traffic trace viewed at different scales. Each graph has 100 data points.
- **Plot 1:** Generate a vector with 100 elements. Starting with Frame 1, each element stores the amount of data from 500 frames of the video sequence.
 - 1st element: #bytes from the first 500 frames (frames 1, 2, ..., 500).
 - 2nd element: #bytes from the frames 501, 502, ..., 1000.
 - ...
- **Plot 2:** Generate a vector with 100 elements. Starting with a randomly selected frame, each element stores the amount of data from 50 frames of the video sequence.
 - Pick a random starting point, e.g., frame 3000.
 - 1st element: #bytes from frames 3001, 3002, ... 3050.
 - 2nd element: #bytes from frames 3051, 3052, ... 3100.
 - ...
- **Plot 3:** Generate a vector with 100 elements. Starting with a randomly selected frame, each element stores the amount of data from 5 frames of the video sequence.
 - Pick a random starting point, e.g., frame 5010.
 - 1st element: #bytes from frames 5010, 5011, ... 5015.
 - 2nd element: #bytes from frames 5016, 5017, ... 5018.
 -
- Plot the content of the vectors in three separate graphs, with the frame number on the x-axis, and the number of bytes on the y-axis. The data points should be depicted as vertical bars (e.g., using the MATLAB function `bar()`).
- Describe your observations of the graphs, and compare them to the scaled versions of the Poisson plots from Part 1.



Lab Report:

- Include your answers to the exercises. Include the plots, a description of the plots, and your observations and discussions.

Part 3. Aggregate traffic on an Ethernet network

Until the early 1990s, it was common to represent traffic in a packet network by a Poisson process. Measurement studies performed in late 1980s showed that this is not a good assumption. One of the most influential such studies, done by Will Leland and Dan Wilson in 1989, measured the total traffic in an Ethernet network at the Bellcore Morristown Research and Engineering facility. The researchers collected traffic traces that each contains a million packet arrivals seen on an Ethernet network.

These traces, the most extensive measurement study at the time, were studied extensively, and led to the conclusion that aggregate traffic (=traffic from many traffic sources) is not “Poisson-like”. In fact, it was convincingly argued that aggregate packet traffic is “self-similar”, meaning, that it looks roughly the same if viewed at different scales. This discovery of the self-similarity of Ethernet traffic had important implications on the design of networks and packet switches.

Exercise 3.1 Download an Ethernet traffic trace

Download one of the Bellcore traffic traces. The file contains one line for each captured Ethernet packets, with two columns. The first column is a timestamp (in milliseconds), the second column is the size of the captured Ethernet packets (in Bytes).

Content: August 1989 trace from Bellcore

Start date: Aug 29, 1989, 11:25am

Data is recorded for about 3142.82 seconds (until 1,000,000 packets had been captured).

- The total trace can be downloaded as a compressed file (5 MB) is available from <http://www.comm.utoronto.ca/~jorg/teaching/ece466/labs/BC-pAug89.TL.Z>
The file can be uncompressed with the Unix command `uncompress`, yielding a file size of approximately 20 MB.

A smaller version of the file (the first 100,000 lines, corresponding to 25.7 seconds) can be downloaded in uncompressed ASCII format from

<http://www.comm.utoronto.ca/~jorg/teaching/ece466/labs/BC-pAug89-small.TL>

- Determine the file size that your Unix and MATLAB installations can comfortably work with. Reduce the file size to the appropriate length.
 - The following description assumes that you can work with approximately 50,000 entries from the trace file. If you work with a smaller file, you need to adjust the numbers and ranges in the next exercises.

Exercise 3.2 Determine statistical properties of the Ethernet trace

- Extract a portion of 100 seconds worth of collected data from the video trace. Select the starting time at random.

- Analyze the trace of the Ethernet traffic:
 - Compute the number of captured packets and total number of bytes;
 - Compute the mean bit rate of the entire trace;
 - Compute the peak bit rate of the trace (defined as the rate that maximizes the ratio “packet size”/“time since last packet”).
 - Compute the ratio of the peak rate and the average rate. Compare this value to the peak-to-average rate ratio from the video trace in Part 2.


- Generate a set of graphs that show properties of the Ethernet traffic trace:
 1. Generate a graph that depicts the packet size as a function of time.
 2. Generate a graph that shows the distribution of packet sizes (x-axis shows the packet size, y-axis shows the relative frequency).

Exercise 3.2 Scaled depiction of Ethernet traffic

This repeats the exercise from Exercise 2.2 for Ethernet traffic where you generate and compare scaled plots.

- Continue with the 100 second excerpt of Ethernet traffic from the previous exercise.
- Create three plots of the Ethernet traffic trace viewed at different scales. Each graph has 100 data points.
- **Plot 1:** Generate a vector with 100 elements, where each element stores the amount of data captured in a 1 s interval.
 - 1st element: #bytes from the first second of traffic, i.e., interval [0, 1 s].
 - 2nd element: #bytes from the interval [1, 2 s].
 - ...
- **Plot 2:** Generate a vector with 100 elements, where each element stores the traffic from 100 ms worth of captured traffic. Select a random starting time.
 - Pick a random starting point, e.g., $t = 20$ s.
 - 1st element: #bytes from the interval [20, 20.1 s].
 - 2nd element: #bytes from the interval [20.1, 20.2 s].
 - ...
- **Plot 3:** Generate a vector with 100 elements, where each element stores the traffic from 10 ms worth of captured traffic. Select a random starting time.
 - Pick a random starting point, e.g., $t = 90$ s.
 - 1st element: #bytes from the interval [90, 90.01 s].
 - 2nd element: #bytes from the interval [90.01, 90.02 s].
 - ...

- Plot the content of the vectors in three separate graphs, with the packet number on the x-axis, and the number of bytes on the y-axis. The data points should be depicted as vertical bars (e.g., using the MATLAB function `bar()`).
- Describe your observations of the graphs, and compare them to the plots from the Poisson traffic and the video trace.



Lab Report:

- Include your answers to the exercises. Include the plots, a description of the plots, and your observations and discussions.

Part 4. (Optional) Metrics for Characterizing Network Traffic

The comparison of the plots from Poisson, video traffic, and measured Ethernet traffic should have illustrated the vast differences of these sources. There is a whole collection of metrics that help with quickly characterizing the properties of a certain traffic mix on a network. You have already seen the peak-to-average rate ratio, a widely used metrics to describe the burstiness of VBR video traffic. In this part of the lab, you will compute other metrics for the three traffic types.

Exercise 4.1 Autocorrelation function

The **autocorrelation** is a mathematical tool to analyze if a function has a repeating or periodic pattern. The autocorrelation function is the covariance of a random process with a time-shifted version of itself, divided by the variance. The definition is as follows:

$$\begin{aligned} R(k) &= \frac{\text{Cov}[X_n, X_{n+k}]}{\text{Var}[X_n]} \\ &= \frac{E[(X_n - \bar{X})(X_{n+k} - \bar{X})]}{\text{Var}[X_n]} \end{aligned}$$

where, $E[]$, $\text{Var}[]$, and $\text{Cov}[]$ indicate the expected value, the variance and the covariance. We also use $\bar{X} = E[X_n]$.

- Generate graphs that compare the autocorrelation function of the packet sizes from the Poisson traffic in Part 1 (Exercise 1.4), the frame sizes of the video trace in Part 2, and the sizes of Ethernet packets from Part 3, and provide a comparison of the plots.

Exercise 4.2 Burstiness function

The burstiness function plots the worst case average bit rate over a time interval of a given length. The function gives a more detailed view of the burstiness than the peak-to-average rate ratio. The function is defined as follows

$$\gamma(t) = \frac{1}{t} \max_x (A(x+t) - A(x))$$

- Generate graphs that show the index of burstiness function of the arrivals from the Poisson traffic in Part 1 (Exercise 1.2), the frame sizes of the video trace in Part 2, and the sizes of Ethernet packets from Part 3, and provide a comparison of the plots.

Exercise 4.3 Index of dispersion

The index of dispersion can also be used as an indication of the burstiness of traffic. The index of dispersion for an arrival function $A(t)$ (which describes the arrivals in the time interval $[0,t]$) is defined as:

$$\delta(t) = \frac{\text{Var}[A(t)]}{E[A(t)]}$$

- Generate graphs that show the index of dispersion function of the arrivals from the Poisson traffic in Part 1 (Exercise 1.2), the frame sizes of the video trace in Part 2, and the sizes of Ethernet packets from Part 3. The index of dispersion should be shown for times up to 100 seconds.
- Provide a comparison of the plots.

Feedback Form for Lab 1

- Complete this feedback form at the completion of the lab exercises and submit the form when submitting your lab report.
- The feedback is anonymous. **Do not put your name on this form** and keep it separate from your lab report.
- For each exercise, please record the following:

	Difficulty (-2,-1,0,1,2) -2 = too easy 0 = just fine 2 = too hard	Interest Level (-2,-1,0,1,2) -2 = low interest 0 = just fine 2 = high interest	Time to complete (minutes)
Part 1. Poisson Traffic			
Part 2. Compressed Video Traffic			
Part 3. Aggregate traffic on an Ethernet network			
Part 4. Characterizing Traffic			

Please answer the following questions:

- What did you like about this lab?
- What did you dislike about this lab?
- Make a suggestion to improve the lab.