

Recitation 12

Internet Technology (Section 01)

Assignment 4: Packet Trace Analysis

- Pcapng - a file format for storing packet traces

The screenshot displays the Wireshark network protocol analyzer interface. At the top, the 'Menu' bar includes File, Edit, View, Go, Capture, Analyze, Statistics, Telephony, Tools, Internals, and Help. Below the menu is a toolbar with various icons for file operations and analysis. The main window is divided into three panes:

- Packet List:** A table showing a list of captured packets. The columns are No., Time, Source, Destination, Protocol, Length, and Info. The table contains 12 rows of data, including ARP, NBNS, ICMP, IGMP, DNS, SSDP, and TCP packets.
- Packet Details:** A pane showing the hierarchical structure of the selected packet (Frame 11). It includes Ethernet II, Internet Protocol, and Transmission Control Protocol (TCP) fields. The TCP field shows flags: 0x02 (SYN) and window size value: 64240.
- Packet Bytes:** A pane showing the raw hex and ASCII data of the selected packet. The hex data is displayed in columns, and the ASCII data is shown in a separate column.

Red arrows point to the Menu bar, the Packet List table, the Packet Details pane, and the Packet Bytes pane, with labels 'Menu', 'Packet List', 'Packet Details', and 'Packet Bytes' respectively.

Reading pcap files

- Pcapng - a file format for storing packet traces
 - Python Scapy lib can be used to read pcap files

```
from scapy.all import *  
pcap = rdpcap("pcap1.pcap")  
print(pcap)
```

Reading pcap files

- Pcapng - a file format for storing packet traces
 - Python Scapy lib can be used to read pcap files
 - Obtain a dictionary of sessions

```
from scapy.all import *
pcap = rdpcap("pcap1.pcap")
sessions = pcap.sessions()
print(sessions)
print(sessions["Other"])
for p in sessions["Other"]:
    print(p)
```

A complete example


```
#!/usr/bin/python3
# Example code using Scapy Python library
# counts packets, TCP packets, UDP packets, and shows the time-of-arrival of HTTP requests
# (c) 2023 R. P. Martin, GPL version 2
from scapy.all import *
import sys
import time
import math

# make sure to load the HTTP layer or your code will silently fail
load_layer("http")

# name of the pcap file to load
pcap_filename = "pcap1.pcap"
# example counters
number_of_packets_total = 0
number_of_tcp_packets = 0
number_of_udp_packets = 0
processed_file = rdpcap(pcap_filename) # read in the pcap file
sessions = processed_file.sessions() # get the list of sessions
for session in sessions:
    for packet in sessions[session]: # for each packet in each session
        number_of_packets_total = number_of_packets_total + 1 # increment total packet count
        if packet.haslayer(TCP): # check if the packet is a TCP packet
            number_of_tcp_packets = number_of_tcp_packets + 1 # count TCP packets
            source_ip = packet[IP].src # note that a packet is represented as a python hash table with keys corresponding to
            dest_ip = packet[IP].dst # layer field names and the values of the hash table as the packet field values
            if packet.haslayer(HTTP):
                if HTTPRequest in packet:
                    arrival_time = packet.time
                    print("Got a TCP packet part of an HTTP request at time%0.4f for server IP%s" % (arrival_time, dest_ip))
                    packet.show()
        else:
            if packet.haslayer(UDP):
                number_of_udp_packets = number_of_udp_packets + 1
print("Got %d packets total, %d TCP packets and %d UDP packets" % (number_of_packets_total, number_of_tcp_packets, number_of_udp_packets))
```

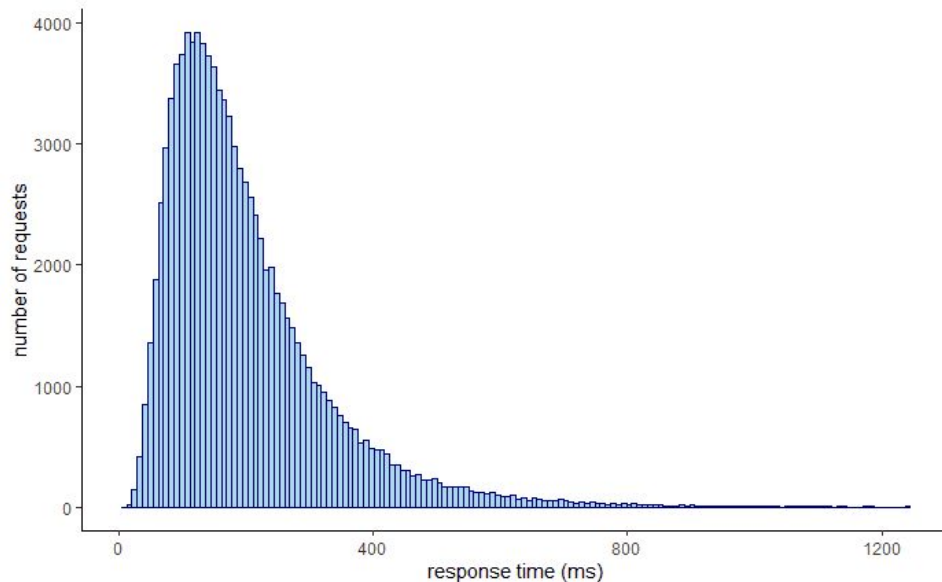
Tail Latency

- High percentile latencies of a distribution of latencies




Tail Latency

- High percentile latencies of a distribution of latencies
 - Why is this important?



Tail Latency

- High percentile latencies from a distribution of latencies
 - Why is this important?
- Maybe: The most number of requests come from a very important client, they expect least issues

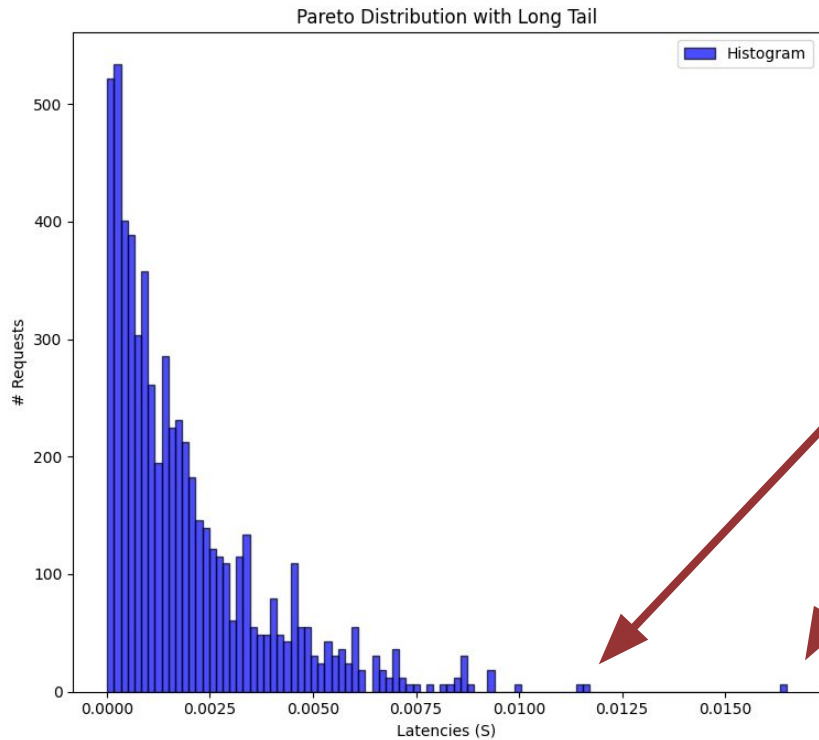


Visualizing Tail Latency

- We can use matplotlib

```
import numpy as np
import matplotlib.pyplot as plt
# Parameters for the Pareto distribution
alpha = 500 # shape parameter (controls tail length)
x_min = 0.0 # minimum value
# Generate random samples from Pareto distribution
np.random.seed(42) # for reproducibility
pareto_samples = np.random.pareto(alpha, 1000) + x_min
# Plot the histogram of the samples
plt.hist(pareto_samples, bins=100, density=True, alpha=0.7, color='blue', edgecolor='black')
# Set plot labels and title
plt.title('Pareto Distribution with Long Tail')
plt.xlabel('Latencies (S)')
plt.ylabel('# Requests')
plt.legend(['Histogram'])
# Show the plot
plt.show()
```

Visualizing Tail Latency



Tail Latencies!

Always Visualize Data!

- Anscombe's quartet
 - 4 datasets with identical summary stats

Always Visualize Data!





- Anscombe's quartet
 - 4 datasets with identical summary stats

Property	Value	Accuracy
Mean of x	9	exact
Sample variance of x : s_x^2	11	exact
Mean of y	7.50	to 2 decimal places
Sample variance of y : s_y^2	4.125	± 0.003
Correlation between x and y	0.816	to 3 decimal places
Linear regression line	$y = 3.00 + 0.500x$	to 2 and 3 decimal places, respectively
Coefficient of determination of the linear regression: R^2	0.67	to 2 decimal places

Always Visualize Data!

- Anscombe's quartet
 - 4 datasets with identical summary stats

Property	Value	Accuracy
Mean of x	9	exact
Sample variance of x : s_x^2	11	exact
Mean of y	7.50	to 2 decimal places
Sample variance of y : s_y^2	4.125	± 0.003
Correlation between x and y	0.816	to 3 decimal places
Linear regression line	$y = 3.00 + 0.500x$	to 2 and 3 decimal places, respectively
Coefficient of determination of the linear regression: R^2	0.67	to 2 decimal places



- Yet when plotted, the data looks very different!