



DEPARTMENT OF INFORMATICS

Master Thesis

Exploring MultiPath TCP Through Discrete Event Simulation

Henrik Libeck Hexeberg
Supervisor: Øyvind Ytrehus

June 1, 2021

Abstract

Global internet usage is rapidly becoming more mobile. Devices are, as a consequence, equipped with multiple network interfaces to meet the demand for mobility. With the emergence of 5G-technology, this trend will most likely continue. However, most internet traffic makes use of standard TCP. TCP does not allow a connection to span several interfaces, even though multiple network interfaces are available. MPTCP is a proposed protocol that enables the use of several interfaces. By distributing data through several paths, MPTCP aims at making internet usage more robust while providing higher throughput and multihoming features. In this thesis, we develop a discrete event simulator to simulate a network. Using this simulator, we analyze and evaluate the performance of MPTCP in various situations. Our findings show that MPTCP has several pitfalls and highlight possible mitigations.

Acknowledgements

I want to express my special thanks of gratitude to my supervisor Øyvind Ytrehus and everyone at Simula@UiB. Thank you for all the mentoring, supervision, and patience.

Besides my supervisor, I would like to thank my family and Mariah, my friends Christopher, Jonas, Håkon, Anders, and Sondre, in addition to all the people I have gotten to know during my time at the University in Bergen. Your help, support, and encouragement have truly been appreciated.

Contents

1	Introduction	7
1.1	Motivation	7
1.2	Goal	8
1.3	Overview	8
2	Background	9
2.1	The Internet Protocol Suite	9
2.1.1	Internet Layer	10
2.1.2	Transport Layer	10
2.2	TCP	11
2.2.1	Flow Control	12
2.2.2	Congestion Control	13
2.3	MultiPath TCP	15
2.4	Queuing Theory	16
2.4.1	Kendall Notation	16
2.4.2	Notations	18
2.4.3	Arrival	19
2.4.4	Service	21
2.4.5	General Results	23
2.4.6	Event-Oriented Bookkeeping	26
2.4.7	Markov Chains	26
2.5	Discrete Event Simulator	30
2.5.1	Advantages	30
2.5.2	Disadvantages	30
3	Related Work	31
3.1	MPTCP Research	31
3.2	Contribution	32
4	Method	33
4.1	The Discrete Event Simulator	33
4.1.1	Events	34
4.1.2	Event Handler	34

4.1.3	Simulated Time	35
4.1.4	Random Generator	36
4.2	Router Modeling	36
4.2.1	Queue	36
4.2.2	Transmission and Queue Delay	37
4.3	Channel Modeling	37
4.3.1	Propagation Delay	37
4.3.2	Packet Loss	38
4.4	TCP Modeling	39
4.4.1	Receiving Window	39
4.4.2	Retransmission	40
4.4.3	Sending Window	40
4.4.4	Transmission Delay	41
4.5	MPTCP Modeling	41
4.5.1	Sending Window	41
4.5.2	Receiver Window	42
4.5.3	Delay	42
4.6	Collecting Results	42
5	Results	44
5.1	Setup	44
5.1.1	Path Variations	44
5.1.2	Simulation Results	46
5.2	TCP	46
5.2.1	Short Path Low Loss	46
5.2.2	Short Path High Loss	50
5.2.3	Long Path Low Loss	53
5.2.4	Long Path High Loss	56
5.3	MPTCP	59
5.3.1	Homogeneous Paths Low Loss	60
5.3.2	Homogeneous Paths High Loss	67
5.3.3	Heterogeneous Paths Low Loss	74
5.3.4	Heterogeneous Paths High Loss	81
5.4	Model Validity	88
5.4.1	Modeling Loss	89
5.4.2	Modeling Router Queue Size	89
5.4.3	Packet Size Distribution	89
5.4.4	Processing Delay	90
5.4.5	TCP as Subflow	90
5.4.6	Discrete Event Simulation	91
5.5	Evaluation of Data	91
6	Conclusion	93
6.1	Research Conclusion	93

6.2	Load Imbalance Problem	93
6.3	Out-of-Order Problem	94
6.4	Redundant ACK	94
6.5	Possible Solutions	95
6.6	Future Work	95

Acronyms		97
-----------------	--	-----------

List of Tables

- 2.1 Interarrival Time Distributions 17
- 2.2 Service Time Distributions 17
- 2.3 Queuing Disciplines 18
- 2.4 Key Queueing Theory Notations 19
- 2.5 Basic Bookkeeping Relations 26

- 4.1 Collected and Calculated Results 43

- 5.1 TCP Short Path Low Loss Results 47
- 5.2 TCP Short Path High Loss Results 51
- 5.3 TCP Long Path Low Loss Results 54
- 5.4 TCP Long Path High Loss Results 57
- 5.5 MPTCP Homogeneous Paths Low Loss Results 61
- 5.6 MPTCP Homogeneous Paths High Loss Results 68
- 5.7 MPTCP Heterogeneous Paths Low Loss Results 75
- 5.8 MPTCP Heterogeneous Paths High Loss Results 82

List of Figures

2.1	TCP/IP Data Flow	10
2.2	TCP Three-Way Handshake	11
2.3	TCP Reno and Tahoe Congestion Control [6]	14
2.4	MPTCP [10]	16
2.5	Little's Law	23
2.6	Little's Law Applied on the Queue	24
2.7	Lindley's Equation Visualized [12]	25
2.8	Markov Chain With Four States and Transition Probabilities	27
4.1	Simulator Flow Chart	35
4.2	Gilbert-Elliot Burst Model	38
5.1	TCP Short Path Low Loss Packet Arrivals	47
5.2	TCP Short Path Low Loss Packet Departures	48
5.3	TCP Short Path Low Loss Packet Interarrival Time	48
5.4	TCP Short Path Low Loss Packets in System	49
5.5	TCP Short Path Low Loss Packet Time in System	49
5.6	TCP Short Path Low Loss Congestion Window	50
5.7	TCP Short Path High Loss Packet Arrivals	51
5.8	TCP Short Path High Loss Packet Departures	51
5.9	TCP Short Path High Loss Packet Interarrival Time	52
5.10	TCP Short Path High Loss Packets in System	52
5.11	TCP Short Path High Loss Packet Time in System	53
5.12	TCP Short Path High Loss Congestion Window	53
5.13	TCP Long Path Low Loss Packet Arrivals	54
5.14	TCP Long Path Low Loss Packet Departures	54
5.15	TCP Long Path Low Loss Packet Interarrival Time	55
5.16	TCP Long Path Low Loss Packets in System	55
5.17	TCP Long Path Low Loss Packet Time in System	56
5.18	TCP Long Path Low Loss Congestion Window	56
5.19	TCP Long Path High Loss Packet Arrivals	57
5.20	TCP Long Path High Loss Packet Departures	57
5.21	TCP Long Path High Loss Packet Interarrival Time	58
5.22	TCP Long Path High Loss Packets in System	58

5.23	TCP Long Path High Loss Packet Time in System	59
5.24	TCP Long Path High Loss Congestion Window	59
5.25	MPTCP Homogeneous Paths Low Loss Packet Arrivals	62
5.26	MPTCP Homogeneous Paths Low Loss Packet Departures	63
5.27	MPTCP Homogeneous Paths Low Loss Packet Interarrival Time	64
5.28	MPTCP Homogeneous Paths Low Loss Packets in System	65
5.29	MPTCP Homogeneous Paths Low Loss Packet Time in System	66
5.30	MPTCP Homogeneous Paths Low Loss Congestion Window	67
5.31	MPTCP Homogeneous Paths High Loss Packet Arrivals	69
5.32	MPTCP Homogeneous Paths High Loss Packet Departures	70
5.33	MPTCP Homogeneous Paths High Loss Packet Interarrival Time	71
5.34	MPTCP Homogeneous Paths High Loss Packets in System	72
5.35	MPTCP Homogeneous Paths High Loss Packet Time in System	73
5.36	MPTCP Homogeneous Paths High Loss Congestion Window	74
5.37	MPTCP Heterogeneous Paths Low Loss Packet Arrivals	76
5.38	MPTCP Heterogeneous Paths Low Loss Packet Departures	77
5.39	MPTCP Heterogeneous Paths Low Loss Packet Interarrival Time	78
5.40	MPTCP Heterogeneous Paths Low Loss Packets in System	79
5.41	MPTCP Heterogeneous Paths Low Loss Packet Time in System	80
5.42	MPTCP Heterogeneous Paths Low Loss Congestion Window	81
5.43	MPTCP Heterogeneous Paths High Loss Packet Arrivals	83
5.44	MPTCP Heterogeneous Paths High Loss Packet Departures	84
5.45	MPTCP Heterogeneous Paths High Loss Packet Interarrival Time	85
5.46	MPTCP Heterogeneous Paths High Loss Packets in System	86
5.47	MPTCP Heterogeneous Paths High Loss Packet Time in System	87
5.48	MPTCP Heterogeneous Paths High Loss Congestion Window	88

Chapter 1

Introduction

This chapter will give an introduction to the thesis, including motivation, goal, and overview.

1.1 Motivation

The internet was created in a time with no mobile internet devices. Today, mobile devices are widespread, but the fundamental internet protocols are essentially the same. Most devices have the capability to use multiple network interfaces at once but rarely do. However, utilizing multiple interfaces to communicate could effectively eliminate many of the challenges wireless, and mobile devices have.

Firstly, mobile devices move around and will often change IP addresses, which identifies where the data should be sent. The change of IP address can result in a lost connection and the need to establish a new connection with the new IP. The process is relatively time-consuming and is not acceptable in latency-sensitive applications such as Voice over Internet Protocol (VoIP) or real-time video applications.

Secondly, wireless devices are prone to more losses than wired devices due to signal degradation. Using multiple channels combined makes it possible to offload a faulty channel when a failure is detected, making the device more reliable.

Thirdly, applications are becoming more data demanding. Aggregating the potential bandwidth of two or more interfaces could significantly increase throughput while also improving resource usage within the network.

MPTCP is a protocol that enables the use of multiple paths during communication, with the potential to solve the problems mentioned above. In practice, the protocol allows for simultaneously using 4G and WiFi or other combinations. With the emergence of 5G, MPTCP becomes increasingly relevant.

1.2 Goal

This thesis explores the possible benefits and obstacles of using MPTCP. By using a discrete event simulator, we hope to gain the insight needed to test and suggest improvements to MPTCP.

1.3 Overview

Chapter 1 Introduction - Outline of this thesis.

Chapter 2 Background - Introducing concepts, notations, and terminology.

Chapter 3 Related Work - Summary of related research in the field.

Chapter 4 Method - Description of the developed simulator and how it is used.

Chapter 5 Results - Presentation of the obtained results.

Chapter 6 Conclusion - Conclusion and future work.

Chapter 2

Background

This chapter will introduce concepts, notations, and terminology needed to understand the thesis. The internet protocol suite and TCP are sections that describe how the internet works. Then, an introduction to MPTCP is given. At last queueing theory and discrete event, simulation is presented.

2.1 The Internet Protocol Suite

The Internet protocol suite is a number of protocols used to communicate over computer networks. It is often referred to as TCP/IP because of the two fundamental protocols Transmission Control Protocol (TCP) and Internet Protocol (IP). TCP/IP's key characteristic is the use of encapsulation through the concept of network layers. Several different models describe the internet in a layered fashion. In the context of TCP/IP, the model specified by RFC1122 [1] is often used:

1. Application layer
2. Transport layer
3. Internet layer
4. Link layer

All layers depend on the layers beneath, and should in theory, not interfere with each other. Information is only passed between adjacent layers. Each layer has layer-specific protocols that handle the responsibilities assigned to the layer. For example, the transport layer should not handle the routing of packets; this is the internet layer's responsibility. Equally, the application layer should not affect the rate at which packets are sent, as it is the transport layer's responsibility. [2]




Figure 2.1: TCP/IP Data Flow

2.1.1 Internet Layer

The internet layer is responsible for communication between hosts through the use of IP addresses. IP addresses allow for communication between hosts on the same network as well as communication between networks through connected gateways. The Internet Protocol (IP) is the protocol used to find the next host on the path to the destination of the packet. When the next host is found, the packet is passed to the link layer. Here the packet is sent to the given host. The next host will, in turn, do the same until the destination is reached. The Internet layer provides best-effort delivery and therefore does not guarantee the eventual delivery of the sent packets. [2]

IP Fragmentation

IP fragmentation is the capability of splitting and merging packets in order not to exceed the links Maximum Transmission Unit (MTU), i.e., the maximum size of data that can be transmitted at once. Different links on the Internet have different MTU, which affects the number and size of the packets to be sent. All protocols that directly or indirectly use IP must consider this as IP fragmentation can cause higher propagation and processing delay. Fragments of packets can be lost, leading to retransmission of the whole packet. [2]

2.1.2 Transport Layer

The Transport layer is situated between the Application layer and the Internet Layer in the TCP/IP model. In contrast to the Internet layer, the Transport layer is only operating on endpoints in the communication. The layer provides host-to-host communication for the application layer to use. This layer contains the two protocols TCP, which provides connected and reliable communication, and User Datagram Protocol (UDP) that provides connectionless and best-effort communication. The layer can facilitate the following

features:

1. Connection based communication
2. Reliable communication
3. Fairness to other users of the Internet

[2]

2.2 TCP

The Transmission Control Protocol (TCP) is a connection-based and reliable protocol situated in the Transport layer. It features flow control and congestion control mechanisms that ideally should be fair to other internet users. However, the fairness of the mechanisms are disputed; connections with longer paths can, for instance, experience lower throughput than competing connections with shorter path [3]. Also, there are multiple versions of TCP that mainly differ in the way the congestion control mechanism is implemented. This thesis will only concentrate on the TCP Tahoe and TCP Reno implementations.

TCP rely on a virtual connection to communicate with the other part. The virtual connection is only maintained by the endpoints and is needed in the process of making TCP reliable. TCP connections are established with a three-way handshake as shown in figure 2.2.

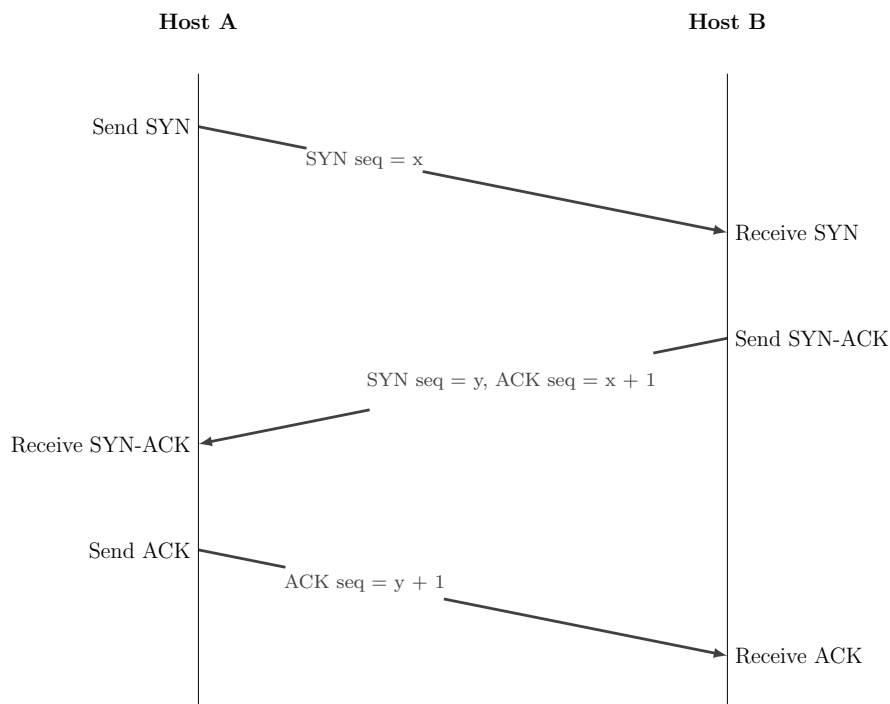


Figure 2.2: TCP Three-Way Handshake

Host A sends a Synchronize (SYN) packet to host B. The SYN packet is what initiates a connection. Upon receiving the SYN, host B will acknowledge (ACK) the SYN while opening a connection with Host A. this packet is called SYN-ACK. Host A will, in turn, acknowledge the incoming connection, and the two parts have established a connection. [2]

2.2.1 Flow Control

Flow Control is a mechanism in TCP for ordered and reliable data transfer. Flow control established the maximum number of packets that can be sent and controls what sending discipline is used to make communication reliable. The mechanism often uses a form of Automatic Repeat Request (ARQ), but Forward Error Correction (FEC) can also be used. ARQ's and FEC's are error control mechanisms used to either notify the sender that an error has occurred or assume an error rate and add redundancy. In ARQ's this means acknowledging received packets (ACK). The absence of an ACK tells the sender that retransmission is needed. For FEC's, the redundant packets are added in the hopes of recreating missing packets. However, this scheme must know the packet loss rate in forehand to know the redundancy needed [4]. Hybrids of the two error control mechanisms have also been implemented, with the benefit being that the loss rate can be unknown and adjusted [5]. Further, the TCP's flow control will notify the sender how many packets can be sent at any given time. This limit is called the Receiver Window (RWND). It prohibits packet loss caused by congestion at the receiver's end. Since sending one packet at a time is slow, different methods of sending multiple packets are often used. The technique must, however, still respect the restriction given by the RWND. [2]

Retransmission

Retransmissions occur when packets are lost, and an ARQ or ARQ-FEC hybrid mechanism is used. Packet loss can happen due to network congestion or link degradation caused by interference or other conditions. Regardless, the occurrence of a lost packet is detected by either not receiving an ACK or by getting multiple ACKs called DupACKs on a specific packet (caused by the receiver getting multiple packets out of order). A DupACK signals the sender to retransmit a specific packet. Not receiving an ACK will, after a timeout, trigger retransmission. [2]

ARQ

Automatic Repeat Request (ARQ) can be implemented in several different ways. The most common are Stop-and-wait, Go-Back-N, and Selective Repeat. The Stop-and-wait ARQ is the simplest. It sends only one packet at a time and waits for the acknowledgment before sending a new packet. A lost packet will trigger a retransmission of the packet sent.

The Go-Back-N ARQ is reasonably easy to implement while significantly improving the sending rate compared to Stop-and-wait. It features a sliding window of size N on the

sending side of the communication. The sender sends N packets before requiring an ACK, acknowledging all the N packets. If one of the N packets or the ACK is lost, all N packets will be retransmitted. While this ARQ is more efficient compared to Stop-and-wait, it can cause excessive retransmissions.

The Selective Repeat ARQ can reduce the number of retransmissions needed but at the cost of higher complexity and higher memory requirements. It uses a sliding window with size N at both sides of the communication. N packets are sent, and all packets are ACKed individually. When a packet is lost, the packets out of order, i.e., the packets that in sequence are after the lost packet, are stored in the receiver's buffer. Either through a Duplicate Acknowledgement (DupACK) or through a Retransmission Timeout (RTO). The lost packet is retransmitted, and the stored packets can be sent to the application layer in the correct order. [2]

2.2.2 Congestion Control

Congestion control consists of four mechanisms that limit the number of packets that can be sent simultaneously through a scheme called Additive-Increase/Multiplicative-Decrease (AIMD) and a Congestion Window (CWND). The AIMD scheme increases CWND by adding a fixed integer a and decreases CWND by multiplying a fixed number b . This implies that a and b have the following properties. $a \in \mathbb{N}$, and $0 < b < 1$. Often, a is set to 1, and b is set to 0.5 but can have different values depending on implementation.

The main four congestion control mechanisms are slow start, congestion avoidance, fast retransmit, and fast recovery. They are aimed at limiting congestion while achieving high throughput and being fair to other users of the internet. By using the Congestion Window (CWND), the mechanisms can dynamically adjust the sending rate. It is important to note that while congestion control limits the number of packets that can be sent at a time, the size of the RWND will always work as an upper bound to the congestion control algorithms. $CWND \leq RWND$. [2]




Figure 2.3: TCP Reno and Tahoe Congestion Control [6]

Slow Start

The slow start mechanism is initiated when a new connection is established and when retransmission caused by a Retransmission Timeout (RTO) occurs. It starts with a small initial congestion window and stops its operation when the slow start threshold ($ssthresh$) is hit or congestion is detected. Since the initial $ssthresh$ value is set to an arbitrarily high value, i.e., as high as possible, the initial slow start operation is operating until congestion is detected. This can be thought of as testing the waters before diving in. Despite its name, slow start is increasing the sending rate aggressively. The congestion window size is increasing by one for every acknowledged packet. In practice the congestion window will double in size every Round Trip Time (RTT), making it an exponential growth. [7] [2]

Congestion Avoidance

Congestion avoidance is initiated when slow start hits $ssthresh$. The mechanism will continue to adjust the sending rate until congestion is detected, in which case the CWND is multiplicatively decreased. When operating without loss, the algorithm increases the CWND by adding $a/CWND$ for each acknowledged packet, i.e., adding a to CWND each RTT. Consequently, congestion avoidance operates with linear growth. When congestion occurs, the sender is notified in two ways, either explicitly through a DupACK or a RTO. This is where the two main implementations of TCP differ. TCP Tahoe will treat three

successive DupACK the same way as a normal RTO, effectively reducing the CWND to the initial value and starting the slow start mechanism. TCP Reno will, however, treat three successive DupACK's in another way. CWND will be multiplicatively decreased with the factor b , then a fast retransmit of the lost packet is triggered, and the fast recovery mechanism is initiated. If a RTO occurs, then the slow start mechanism starts with the initial CWND, just like in TCP Tahoe. [7] [8] [2]

Fast Retransmit

Fast retransmit is a mechanism that improves the throughput by fast-retransmitting packets that most likely are lost due to congestion. When a DupACK is received, it can mean that the packet is lost, but it can also mean that packets must be reordered. If reordering is needed, the retransmission of the packet is not necessary. However, if multiple DupACK's occur, the packet is more likely to have been lost in transmission. The fast retransmit mechanism is therefore initiated if three or more DupACK's are received. After the presumed lost packet is retransmitted, TCP initiates either slow start or fast recovery depending on implementation. [7] [2]

Fast Recovery

Fast recovery is a mechanism that, much like fast retransmit, exists to improve the throughput. It works as an alternative to slow start by going directly to the congestion avoidance mechanism and is deployed in TCP Reno. The reasoning behind this mechanism is that the occurrences of DupACKs are proof that the lost packet is not part of a major congestion problem in the network, but rather a rare event. DupACKs are only sent if packets are received out of order, essentially proving that no serious network congestion has occurred. [7] [2]

2.3 MultiPath TCP

MultiPath TCP (MPTCP) is an emerging protocol situated in the transport layer that aims to utilize multiple paths during communication. It is implemented using TCP extensions defined in RFC8684 [9]; however, it has been divided opinions on whether the implementation should be embedded in the payload instead. The goals of MPTCP are many, with the primary goal being to make the internet more robust with better performance and load-balancing. This should, however, be implemented in a way that presents itself as TCP to the layer above, namely the application layer. Other goals of MPTCP are to perform at least as good as standard TCP in terms of resource usage and reliability.

One of the main benefits of using multiple paths aside from higher throughput is the redundancy of multiple paths. If one path no longer exists, all data can be sent through the remaining paths. Mobile devices often have WiFi and 4G(LTE) available at the same

time. Using MPTCP it is possible to use both interfaces at the same time as shown in Figure 2.4.

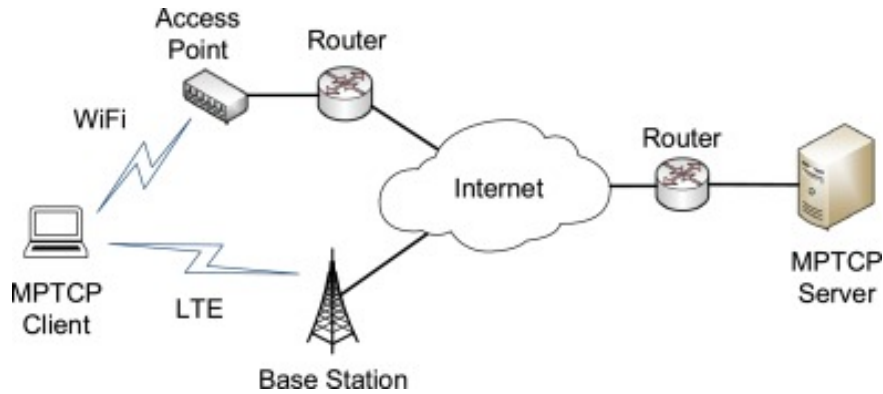


Figure 2.4: MPTCP [10]

Due to the many middleboxes on the internet, such as NAT and load-balancers, the implementation must touch all TCP functions like flow control, congestion control, and connection establishment and tear-down. This increases the complexity needed to deploy a working MPTCP implementation successfully. [11]

2.4 Queuing Theory

This section will go through the fundamentals of queueing theory and is influenced by the book Fundamentals of Queueing Theory [12]. This source is used throughout this section and is the used source if not stated otherwise.

2.4.1 Kendall Notation

In queueing theory, a node's queueing process can be described using the Kendall notation. This notation has six parameters and is divided by slashes like this: $A/S/c/K/N/D$. The position and value describe how the queueing process handles and experiences arrivals (A), Service (S), number of servers (c), capacity (K), source size (N), and queueing discipline (D). However, the three last parameters are often left unspecified. In which case the the values are assumed to be $K = \infty$, $N = \infty$ and $D = \text{FIFO}$. Each parameter can take multiple different values.

Interarrival Time Distribution (A)

Interarrival time distribution (A) is a symbol that denotes the interarrival time distribution used in the queueing system and is the first parameter in the Kendall notation. See table 2.1

Symbol	Description	Explanation
M	Exponential	Poisson arrival process. Inherits the Markovian or memoryless property and is discrete.
M^X	Batch or Bulk Markovian	Poisson arrival process with a random X arrivals between 1 and ∞ . Inherits the Markovian or memoryless property and is discrete.
D	Deterministic	Deterministic arrival process. Arrivals occur at a fixed rate.
G	General	Independent or unknown arrival process.
E_k	Erlang type-k	An Erlang distribution with k as shape parameter. i.e the average time it takes for k Poisson arrivals to occur ($E_1 = M$)

Table 2.1: Interarrival Time Distributions

Service Time Distribution (S)

Service time distribution (S) is a symbol that denotes the service time distribution used in the queuing system and is the second parameter in the Kendall notation. See table 2.2

Symbol	Description	Explanation
M	Exponential	Exponential service time. Has the Markovian or memoryless property.
M^X	Batch or Bulk Markovian	Exponential service time with a random X between 1 and ∞ customers being served at once. Has the Markovian or memoryless property.
D	Deterministic	Deterministic service time. The time it takes to serve a customer is fixed
G	General	Independent or unknown service time.
E_k	Erlang type-k	An Erlang distribution with k as shape parameter. i.e the average time it takes to serve k customers ($E_1 = M$)

Table 2.2: Service Time Distributions

Parallel Servers (c)

Parallel servers (c) is an integer value that denotes the number of parallel servers in the system and is the third parameter in the Kendall notation. Each parallel server has an independent but identical service distribution. For instance, an M/M/2 queue has a Poisson arrival process (section 2.4.3) and an exponential service rate (section 2.4.4). Both the arrival and service distributions are independent of the number of parallel servers in

the system. The arrival rate represents the arrivals to the system and not a server. The service rate describes the time a server uses to serve a customer and not the time the system uses.

System Capacity (K)

System capacity (K) is an integer value that denotes the capacity of the system and is the fourth parameter in the Kendall notation. It describes the maximum number of customers in the system at any time. An M/M/2/3 queue can have no more than one customer in the queue since this implies that two customers are being served.

Source Population (N)

Source population (N) is an integer value that denotes the source population and is the fifth parameter in the Kendall notation. Other values than ∞ mean that there is a finite source population which in turn will affect the arrival rate to the queuing system.

Queuing Discipline (D)

An acronym that denotes the queuing discipline used in the queuing system and is the sixth parameter in the Kendall notation. See table 2.3

Acronym	Description
FIFO/FCFS	Fist In Fist Out/First Come First Serve
LIFO/LCFS	Last In First Out/Last Come First Serve
SIRO/RSS	Service In Random Order/Random Selection for Service
PR	Priority
GD	General Discipline

Table 2.3: Queuing Disciplines

2.4.2 Notations

Table 2.4 present the key notations and results for G/G/c queues.

Notation	Description
A	Random arrival time
S	Random service time
$E[x]$	Expected value or mean of x
$\lambda \equiv 1/E[A]$	Average arrival rate
$\mu \equiv 1/E[S]$	Average service rate
c	Number of servers
$r \equiv \lambda/\mu$	Offered load
$\rho \equiv \lambda/c\mu$	Traffic intensity or utilization
T	Random time a customer spends in the system
T_q	Random time a customer spends in the queue
W	Average time a customer spends in the system
W_q	Average time a customer spends in the queue
N	Random number of customers in the system
N_q	Random number of customers in the queue
L	Average number of customers in the system
L_q	Average number of customers in the queue

Table 2.4: Key Queueing Theory Notations

2.4.3 Arrival

An arrival process defines the rate at which customers arrive at a queue. λ is used to denote the average arrival rate of the arrival process, which in turn is used in the calculation of several important metrics. Note that the term customer is used as a general term to describe an entity that enters a queueing system. Looking at a queue in a simplistic manner, a customer that arrives at a queue can either be served at once or must wait in queue for service. However, this assumption does not take into consideration that the queue might, for example, have a finite capacity or that customers can choose not to enqueue for some reason. Either way, if the arrivals to the system happen at a higher rate than customers are being served, the queue can grow indefinitely.

Knowing the exact arrival process to a system is often impossible. For this reason, it is common to use a probability distribution to model the way customers arrive at the system. The model will never be perfect, but many natural processes have shown to follow certain distributions closely. Counting processes are used to model the arrival processes since arrivals are discrete events, i.e., events that occur in time.

The Poisson Process

The Poisson process is a counting process that is commonly used to model the arrivals to a queue. A counting process $N(t)$ is defined as a stochastic process in which $N(t)$ takes on non-negative integers and is non-decreasing in time. Counting processes are used to count the number of events that occur in a given time interval.

Definition 2.4.1. A Poisson process with rate $\lambda > 0$ is a counting process $N(t)$ with the following properties:

1. $N(0) = 0$.
2. $P\{1 \text{ event between } t \text{ and } t + \Delta t\} = \lambda\Delta t + o(\Delta t)$.
3. $P\{2 \text{ or more events between } t \text{ and } t + \Delta t\} = o(\Delta t)$
4. The number of events in non-overlapping intervals are statically independent; that is, the process has independent increments.

The quantity $o(\Delta t)$ can intuitively be seen as essentially zero when $\Delta t \rightarrow 0$. Formally the relation between Δt and $o(\Delta t)$ is:

$$\lim_{\Delta t \rightarrow 0} \frac{o(\Delta t)}{\Delta t} = 0$$

While definition 2.4.1 introduce the properties of a Poisson process, it does not give the probability distribution for the counting process $N(t)$. First, the Poisson random variable needs to be defined.

Definition 2.4.2. A Poisson random variable is a discrete random variable with a probability mass function.

$$p_n = e^{-A} \frac{A^n}{n!} \quad (n = 0, 1, 2, \dots)$$

where $A > 0$ is a constant.

Using the defined properties of a Poisson process 2.4.2 and the definition of a Poisson random variable 2.4.1, it is possible to obtain the probability distributing of the counting process $N(t)$

Theorem 2.4.1. Let $N(t)$ be a Poisson process with rate $\lambda > 0$. The number of events occurring by the time t is a Poisson random variable with mean λt . That is,

$$p_n = e^{-\lambda t} \frac{(\lambda t)^n}{n!} \quad (n = 0, 1, 2, \dots)$$

where $p_n(t) \equiv P\{N(t) = n\}$.

The Poisson properties have, as a bonus, some additional characteristics that make them unique and useful when modeling arrivals.

Theorem 2.4.2. Let $N(t)$ be a Poisson process with rate λ . Then the times between successive events are independent and exponentially distributed with rate λ

Theorem 2.4.2 states that the Poisson process has a relation to the exponential distribution. This results in the Poisson process inherently having the memoryless property (Theorem 2.4.7). Another interesting characteristic is that a Poisson process has stationary increments, which in essence means that the probability of an event's occurrence in

time is dependent on the length of the interval, and non-dependent on the exact moment the event occurs. This characteristic is defined as follows:

Theorem 2.4.3. *A Poisson process has stationary increments. That is, for $t > s$, $N(t) - N(s)$ is identically distributed as $N(t+h) - N(s+h)$ for any $h > 0$.*

The Poisson process can be thought of as a random process, due to its exponentially distributed interarrival times. While the concept of randomness is hard to define, the experienced randomness of a Poisson process is a result of the following characteristic.

Theorem 2.4.4. *Let $N(t)$ be a Poisson process. Given that k events have occurred in a time interval $[0, T]$, the times $t_1 < t_2 < \dots < t_k$ at which the events occurred are distributed as the order statistics of k independent uniform random variables on $[0, T]$*

A significant side effect of Theorem 2.4.4 is the Poisson Arrivals See Time Averages (PASTA) property [13]. This property states that the observations taken according to a stochastic process have the same probability of observing a system in a given state as observing the system at Poisson selected points. More intuitively, the PASTA property states that customers arriving in a queue will have the same probability of finding the queue at a certain length as an outsider looking at the queue at a random point in time.

The last characteristic of the process is that splitting a Poisson process into sub-processes will result in several Poisson processes. Equivalently, the outcome of combining several Poisson processes will yield a Poisson process. This characteristic is formally defined in the following theorems.

Theorem 2.4.5 (Splitting). *Let $N(t)$ be a Poisson process with rate λ . Suppose that each event is labeled a type- i event with probability p_i , independent of all else. Let $N_i(t)$ be the number of type- i events by time t . Then $N_i(t)$ is a Poisson process with rate λp_i , $i = 1, \dots, n$. Furthermore, $N_i(t)$ and $N_j(t)$ are independent, for all $i \neq j$*

Theorem 2.4.6 (Combining). *Let $N_1(t), \dots, N_n(t)$ be independent Poisson processes with rates $\lambda_1, \dots, \lambda_n$, respectively. Then $N(t) \equiv N_1(t) + \dots + N_n(t)$ is a Poisson process with rate $\lambda \equiv \lambda_1 + \dots + \lambda_n$.*

2.4.4 Service

The service time of a queue is the time it takes for a customer to complete service. In other terms, the time a customer occupies a server. Once a customer has completed service, the next customer can be served. When dealing with queuing systems with multiple servers, the service rate is provided per server and not the service rate as a whole. For this reason, the traffic intensity (ρ) measure is a product of the number of servers while the offered load (r) is not Table 2.4. In many cases, the service time distributions can seem random and are often modeled as a stochastic process, much like the arrival process. However, unlike the arrival processes, the service patterns should model time and not events in time. The probability distributions must consequently be continuous.

The Exponential Distribution

The exponential distribution is continuous and is commonly used to model the service rate of a queueing system. Additionally, it is used as the continuous-time between two events in the Poisson process (Theorem 2.4.2) and is fundamental in continuous-time Markov chains (subsection 2.4.7). The probability distribution has the memoryless property, which states that the current and future states are independent of the previous. This property is formally defined as follows:

Definition 2.4.3. *A random variable T has the memoryless property if*

$$P\{T > t + s | T > s\} = P\{T > t\} \quad (s, t \geq 0).$$

The exponential distribution is often used to model time. This could, for example, be the time between two successive events, i.e., the Poisson process, or the time an arbitrary process needs to complete, as is the case for service patterns. This pseudo-random value is defined as an exponential random variable.

Definition 2.4.4. *An exponential random variable is a continuous random variable with:*

$$\text{PDF :} \quad f(t) = \lambda e^{-\lambda t} \quad (2.1)$$

$$\text{CDF :} \quad f(x) = 1 - e^{-\lambda x} \quad (2.2)$$

$$\text{CCDF :} \quad f(x) = e^{-\lambda x} \quad (2.3)$$

where $\lambda > 0$ is a constant and $t \geq 0$

Given the memoryless property (Definition 2.4.3) and Complementary Cumulative Distribution Function (CCDF) (Definition 2.4.4), it is possible to show that the exponential distribution is in fact memoryless.

$$\begin{aligned} P\{T > t + s | T > s\} &= \frac{P\{T > t + s, T > s\}}{P\{T > s\}} \\ &= \frac{P\{T > t + s\}}{P\{T > s\}} \\ &= \frac{e^{-\lambda(t+s)}}{e^{-\lambda s}} \\ &= e^{-\lambda t} \end{aligned} \quad (2.4)$$

Theorem 2.4.7. *An exponential random variable has the memoryless property*

2.4.5 General Results

Little's Law

Little's Law is a fundamental queueing theory result that provides a relationship between the average arrival rate λ , the average time a customer spends in the queueing system W , and the average number of customers in the system L . The result is general, meaning that it applies to all possible queues given some constraints. The relationship is based on the long-run averages; hence the three values must be finite when the time and the cumulative number of customers approach infinity. Little's Law is therefore not applicable with systems where the queue size increase without bound.

Theorem 2.4.8 (Little's Law). *In a queueing process, let $1/\lambda$ be the mean time between the arrivals of two consecutive customers, L be the mean number of customers in the system, and W be the mean time spent by a customer in the system. It is shown that if the three means are finite and the corresponding stochastic processes strictly stationary, and if the arrival process is metrically transitive with nonzero mean [14].*

$$L = \lambda W$$

The quantities L , λ and W are the long-run averages and are defined as follows:

$$\lambda \equiv \lim_{t \rightarrow \infty} \frac{A(t)}{t}, \quad W \equiv \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k W^k, \quad L \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T N(t) dt. \quad (2.5)$$

$A(t)$ is the cumulative number of arrivals at time t , W^k is the time customer number k spent in the system, and $N(t)$ is the number of customers in the system at the time t . Note that Little's Law is an intuitive result. For example, suppose customers arrive at a system exactly every second and customers finish service precisely every second, given that the system starts with one customer in the system. In that case, there will be precisely one customer in the system at all times. However, the law is dependent on several assumptions about the system. In the example given, it is not specified that a customer must be served upon entering the system. Likewise, it is not specified that service must be completed.




Figure 2.5: Little's Law

Figure 2.5 illustrates the law's usability, as well as its limitations. As shown, customers arrive at the system and depart from the system after a given time. No customer can leave in any other way. Thinking of the system as a black box can help in understanding its general nature. Consequently, Little's Law can be used in several different situations. Figure 2.6 and Equation 2.6 shows how the law can be used on the queue only, which will yield the time or number of customers in the queue, thereby excluding the customers being served. In fact, since Little's law is a relation among the variables L , W , L_q and W_q additional equations can be derived shown in Equation 2.7 and Equation 2.8

$$L_q = \lambda W_q \quad (2.6)$$

$$L = L_q + r \quad (2.7)$$

$$W = W_q + 1/\mu \quad (2.8)$$

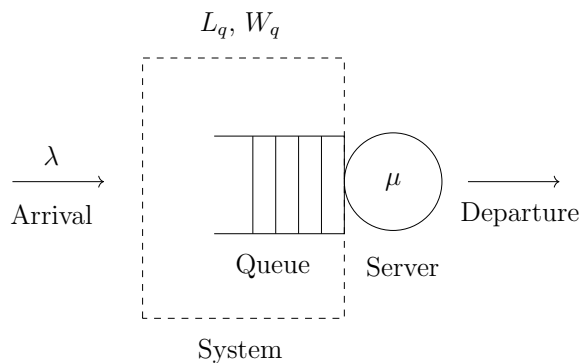


Figure 2.6: Little's Law Applied on the Queue

Lindley's equation

Lindley's equation is an essential general result in Queueing Theory. The equation provides a relation between the time spent in the queue of customer n ($W_q^{(n)}$), the time used to serve this customer ($S^{(n)}$), the time spent in the queue of the next customer $n + 1$ ($W_q^{(n+1)}$), and the interarrival time between the two customers ($T^{(n)}$). However, the equation is not general when it comes to the number of servers and queueing discipline since it assumes one server and a FIFO discipline — a $G/G/1/FIFO$ queue. Note that the subtracting of the interarrival time can, if large enough, result in a negative waiting time. Since customers cannot wait a negative amount of time, a maximum of the equation and zero is performed [15]. Figure 2.7 illustrates the relation given by the equation. Formally Lindley's equation can be described as follows:

Theorem 2.4.9. *A customer arriving at a $G/G/1/FIFO$ queue will experience a waiting time equal to the max of the previous arrived customer plus the service time of this*

customer minus the interarrival time between the two customers, and zero. [15]

$$W_q^{(n+1)} = \max\{W_q^n + S^n - T^n, 0\}$$

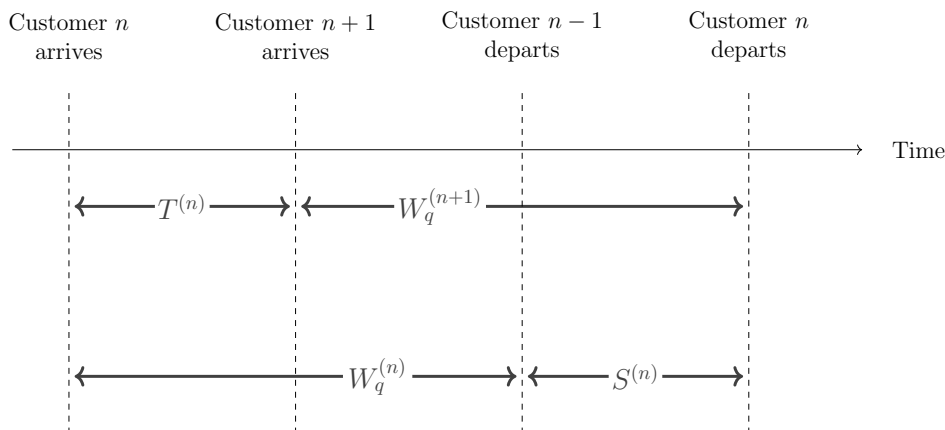


Figure 2.7: Lindley's Equation Visualized [12]

Pollaczek–Khinchine Formula

Until now, the general result has been general in both arrival and service patterns, with Little's Law being completely general and the fact that Lindley's equation is working under assumptions of a single server FIFO queue. An M/G/1 queue has general service distributing and a Poisson arrival process. This motivates the following set of formulas that can be said to extend Little's Law in an M/G/1 queue.

$$L_q = \frac{\lambda^2 E[S^2]}{2(1 - \rho)} \quad (2.9)$$

$$W_q = \frac{\lambda E[S^2]}{2(1 - \rho)} \quad (2.10)$$

$$W = \frac{\lambda E[S^2]}{2(1 - \rho)} + \frac{1}{\mu} \quad (2.11)$$

$$L = \frac{\lambda^2 E[S^2]}{2(1 - \rho)} + \rho \quad (2.12)$$

Consequently, it is possible to obtain results by only knowing the arrival process and the service process of the system. Recalling Table 2.4, the variables λ , μ , ρ , and S are all conditioned on either the service or arrival distribution. Simple justifications of the given formulas can be derived by altering Little's Law and using the PASTA property. Recall that in an arrival context, the PASTA property states that new arrivals will, on average, find the queue in the same state as the time-average state of the queue.

2.4.6 Event-Oriented Bookkeeping

Event-oriented bookkeeping is a method to secure that state change only happens when an event occurs. How queues behave in time is often the objective when analyzing a queue. Observing the arrival of customers and the time it takes to serve this customer is enough to calculate the interarrival time between customers, the time a customer starts service, the departure time of a customer, and the total time a customer used in the system (Table 2.5). However, calculating when the next customer starts service assumes a $G/G/1/FIFO$ queue since this equation uses the previous customer in service. Consequently, another observation is needed to calculate all variables in a $G/G/c$ queue.

Variable	Definition	Sample Relationship	Applicable Queue
$A^{(n)}$	Arrival time of customers n		$G/G/c$
$S^{(n)}$	Service time of customer n		$G/G/c$
$T^{(n)}$	Interarrival time between customer n and customer $n + 1$	$T^{(n)} = A^{(n+1)} - A^{(n)}$	$G/G/c$
$U^{(n)}$	Time customer n starts service	$U^{(n+1)} = \max\{D^{(n)}, A^{(n+1)}\}$	$G/G/1/FIFO$
$D^{(n)}$	Departure time of customer n	$D^{(n)} = U^{(n)} + S^{(n)}$	$G/G/c$
$W_q^{(n)}$	Time customer n is in the queue	$W_q^{(n)} = U^{(n)} - A^{(n)}$	$G/G/c$
$W^{(n)}$	Time customer n is in the system	$W^{(n)} = W_q^{(n)} + S^{(n)}$	$G/G/c$

Table 2.5: Basic Bookkeeping Relations

2.4.7 Markov Chains

Markov chains are stochastic processes that describe the possible change of state in a model. They can be both discrete and continuous with respect to time. Discrete-time Markov chains can only change state at discrete points in time, unlike the continuous-time Markov chains that will remain in a state in an arbitrary amount of time before transitioning to another state. The fundamental basis of all Markov chains is the Markovian property.

Theorem 2.4.10 (Markovian property). *A series of random variables is said to be a Markov chain if it has the following property:*

$$P\{X_{n+1} = j | X_0 = i_0, X_1 = i_1, \dots, X_n = i_n\} = P\{X_{n+1} = j | X_n = i_n\}$$

This property states that the current state is independent of previous states. Knowing the probability of transitioning from the current state to any other state is the same

as knowing the whole history of prior states. There is a relation to the memoryless property 2.4.3 where the next variable is entirely independent of the past. However, Markov chains have a state, unlike the exponential distribution, which makes them more versatile. Further, the arbitrary time spent in a continuous-time Markov chain follows the exponential distribution.

Transition probabilities

The transition probabilities of a Markov chain is the probability of transitioning to a state once a state change happens. Notably, for every event in the case of discrete-time Markov chains, or after an arbitrary amount of time, as is the case for continuous-time Markov chains. Figure 2.8 is an example of a four-state Markov chain with transition probabilities. Note that the sum of all probabilities going out from a state must be equal to 1 as shown in Equation 2.13

$$\sum_{j=0}^m X_{ij} = 1 \quad (2.13)$$




Figure 2.8: Markov Chain With Four States and Transition Probabilities

The single-step probabilities $P\{X_{n+1} = j|N_n = i\}$ is often assumed to be independent of n , which is the time of the system. In which case the Markov chain is homogeneous, and the transition probabilities can be written as $p_{ij} \equiv P\{X_{n+1} = j|N_n = i\}$. Figure 2.8 transition probabilities can be written in matrix form as shown in Equation 2.14 and is called the transition matrix.

$$\mathbf{P} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0.4 & 0.6 & 0 \\ 0.5 & 0.2 & 0 & 0.3 \\ 0.2 & 0.8 & 0 & 0 \end{bmatrix} \quad (2.14)$$

Properties and Terminology

Although the defining property for all Markov chains is the Markovian property Theorem 2.4.10, Markov chains have several other properties or characteristics associated with them. The following definitions will be helpful in the next sections in deciding how to work with particular chains.

Definition 2.4.5 (Accessibility). *A state j is accessible from a state i if there exist a valid path between j and i ($i \rightarrow j$).*

Definition 2.4.6 (Communication). *The two states i and j communicate if i is accessible from j and j is accessible from i ($i \leftrightarrow j$). A Markov chain can have multiple communication classes which means that it is impossible to go back and forth between a set of states.*

Definition 2.4.7 (Reducible). *A Markov chain is irreducible if all states communicate with each other, i.e., all states are within the same communication class. A Markov chain is said to be reducible if the opposite is true.*

Definition 2.4.8 (Recurrent). *A state j is said to be recurrent if the probability of returning to state j , when starting at j , is 1. Formally this is defined as follows:*

$$f_{jj} = \sum_{n=1}^{\infty} f_{jj}^{(n)}$$

If $f_{jj} = 1$ the state is recurrent. The state is transient if $f_{jj} < 1$

Definition 2.4.9 (Positive Recurrent). *A state is positive recurrent if the mean recurrence time m_{jj} is less than ∞ .*

$$m_{jj} = \sum_{n=1}^{\infty} n f_{jj}^{(n)} < \infty$$

If $m_{jj} = \infty$ the state is null recurrent.

Definition 2.4.10 (Periodic). *A state j is periodic if the chain is guaranteed to transition back to state j after a fixed number of steps d , where d is the greatest common divisor of all possible m steps so that the probability of being in state j is $p_{jj}^{(m)} > 0$. A state with period $d = 1$ is aperiodic*

Long-Run Behavior

When working with Markov chains, it is often interesting to know how the chain behaves when the time approaches infinity ($n \rightarrow \infty$). The n -step transition matrix is the matrix that is formed when raising the transition matrix to the power of n , essentially finding the probability of being in certain states after time n . When n approaches infinity ($n \rightarrow \infty$), this matrix is called the limiting matrix. The limiting matrix of the transition matrix \mathbf{P}

Equation 2.14 is:

$$\lim_{n \rightarrow \infty} \mathbf{P}^n = \lim_{n \rightarrow \infty} \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0.4 & 0.6 & 0 \\ 0.5 & 0.2 & 0 & 0.3 \\ 0.2 & 0.8 & 0 & 0 \end{bmatrix}^n \approx \begin{bmatrix} 0.15 & 0.47 & 0.28 & 0.08 \\ 0.15 & 0.47 & 0.28 & 0.08 \\ 0.15 & 0.47 & 0.28 & 0.08 \\ 0.15 & 0.47 & 0.28 & 0.08 \end{bmatrix} \quad (2.15)$$

In Equation 2.15 all rows converge, meaning that in the long run, the probability of being in a state is independent of the starting state. The values can be seen as the fraction of time spent in the system states. Further, if the limiting matrix is converging with equal rows, and each row sum to one $\sum_j \pi_j = 1$, then π_j is a limiting distribution. It can be shown that instead of raising the transition matrix to a high degree, π_j can be found by solving one of the following equation:

$$\pi_j = \sum_i \pi_i P_{ij} \quad (2.16)$$

$$\boldsymbol{\pi} = \boldsymbol{\pi} \mathbf{P} \quad (2.17)$$

Where Equation 2.17 is the vector for of Equation 2.16

Theorem 2.4.11 (Stationary equation). *An irreducible and positive recurrent discrete-time Markov chain has a unique solution to the stationary equations*

$$\boldsymbol{\pi} = \boldsymbol{\pi} \mathbf{P} \quad \text{and} \quad \sum_j \pi_j = 1 \quad (2.18)$$

Furthermore, if the chain is aperiodic, the limiting probability distribution exists and is equal to the stationary distribution.

Another long-run behavior concept is the idea of ergodicity. A process is ergodic if the ensemble average of a process is equal to the average of a process. Intuitively, if measuring an infinite number of processes at a fixed point is equal to the average of a process, then the process is said to be ergodic. Often, processes are time-dependent, and so the approximation is a time average. In which case, ergodicity can be defined as follows:

Definition 2.4.11 (Ergodicity). *A time-dependent process $X(t)$ is ergodic if time averages equal ensemble averages.*

The properties defined in section 2.4.7 decides some resulting properties and how to use the stationary distribution. If the Markov chain is irreducible and positive recurrent, there exists a unique stationary distribution as defined in Theorem 2.4.11. Further, if $\boldsymbol{\pi}^0 = \boldsymbol{\pi}$, i.e., the starting probability vector is equal to the stationary vector, then the process has a unique stationary distribution, and the process is stationary and ergodic. If, on the other hand, the chain is irreducible, positive recurrent, and aperiodic, there exist a limiting distribution equal to the unique stationary distribution. The process is also ergodic.

2.5 Discrete Event Simulator

Discrete event simulation is described in the book *Discrete-Event System Simulation* [16] to be a way of simulation a system in which the state of the system only changes at discrete points in time. These state changes are determined by the so-called events running in the simulator, hence the name discrete event simulator. Events are also tasked with creating and scheduling new events to run in the simulator at a later time. Since an event can create multiple events, a call queue is needed to store and order future events. Simulated time is often used to order the event; however, real-time is also used. [17]

The simulator output is composed of measurements gathered during the simulation, acting as an artificial history of the system generated by model assumptions. These measurements are used to analyze the performance of the modeled system, which in turn is used to estimate the true performance of the system of interest. Generally, simulation is used when it's impractical to experiment on the real system. There may be too high a risk associated with the experiments, or it is impossible to acquire the needed results from the experiments. Regardless of the reason, simulation can be an efficient and practical way of experimenting since all unknown factors can be eliminated, giving precise and reproducible results. Simplified assumptions can, however, lead to misleading results. A good model is therefore key to obtain usable results. [16]

2.5.1 Advantages

The advantages of using discrete event simulation, or just simulation, to analyze a system are many. Two of them are, as mentioned, the potential risk associated with performing tests on the real system and the feasibility of acquiring the needed results from the experiments. Simulation allows for risk-free testing of a system with control over all parameters. Further, insight can be obtained about the interaction and importance of variables in the system and how it affects performance. In addition, questions concerning potential changes to the system can be answered without the need to implement them. In summary, a simulation is a versatile tool that allows the user insight into every aspect of the simulated process. [16]

2.5.2 Disadvantages

Simulating has some clear disadvantages that affect the way one should interpret the obtained results. Simulators are used when complicated systems are analyzed. As a consequence, many of the processes happening within a system are too complicated to model correctly. This motivates the use of random components that can model the seemingly random processes that cannot be modeled effectively. Therefore it can be hard to know if the results are caused by inherent randomness or by interrelationships. Simulation cannot identically replicate the real-world environment. Consequently, the results must be analyzed with this in mind [16]

Chapter 3

Related Work

This chapter will introduce some of the related work previously done in the field of MPTCP and present the contribution of this thesis.

3.1 MPTCP Research

The effort towards developing MultiPath TCP is an active research topic with many issues that need solving to ensure good throughput and efficient use of the added available bandwidth. Ideally, the throughput of MPTCP should be the added potential throughput of all available TCP flows. In reality, this is not achieved and can, in some cases, underperform the throughput of a normal TCP connection. This section will shed light on some of the work done in this field.

The Linux kernel implementation created by C. Paasch, S. Barre, et al. [18] is an effort to enable the simultaneous use of several TCP interfaces. This implementation is aiming at providing MPTCP in a way that resembles TCP to the application layer, while in fact, sending data across multiple paths. The project is a product of the research being done in the field, and the implantation is actively worked on.

Fountain Code-Based MultiPath Transmission Control Protocol (FMTCP) is proposed [19] to tackle the fact that subflows experiencing high loss and delay are becoming bottlenecks of the MPTCP connection. The researchers conducted an extensive simulation-based study, which found that the heterogeneity of paths can cause a significant reduction in the overall performance. The proposed implementation takes advantage of fountain codes to mitigate the negative impact of the heterogeneity of the different paths.

A MPTCP implementation using network coding for mobile devices in heterogeneous networks is presented in this paper [20]. Empirical data was collected to help understand the mobile environment when three heterogeneous networks are available to the mobile device. WiFi, WiMax, and an Iridium satellite network are used to gather the data. Further, a reliable multi-path protocol called MPTCP with Network Coding (MPTCP/NC)

is proposed. The results compared to MPTCP without network coding show that the proposed implementation gives users a higher quality of service since it can overcome packet losses due to lossy wireless network connections.

A research group has presented the experience in implementing MPTCP [11]. The paper focuses on receive buffer tuning and segment reordering to efficiently send data along different paths with different characteristics. The results are then compared to regular TCP. Based on the results and the improvements made through new algorithms that solve important MPTCP problems, it is concluded that MPTCP is ready for adaptation.

The performance of MPTCP over heterogeneous subpaths is the topic of another research paper [21]. The researchers are exploring the performance of MPTCP over real-world networks. The results mainly point to obstacles that need to be addressed if MPTCP is to be widespread. In addition, the results show how heterogeneous paths can reduce the overall performance of MPTCP.

The way packets are scheduled across the available subpaths is important for the performance of a MPTCP implementation. This paper [22] propose two schedulers that aim to optimize the latency to match an applications desired latency.

The asymmetric nature of heterogeneous paths is causing problems such as the out-of-order problem and the load imbalance problem. This paper [23] is proposing a Receive Buffer Pre-division based flow control mechanism (RBP) for MPTCP as an alternative to scheduling algorithms. The control mechanism divides the receive buffer according to the subflows characteristics and controls the data transmission in each subflow. The NS-3 simulator is used to verify the scheme, which produces results that show a significant increase in throughput.

3.2 Contribution

The focus of this thesis is to analyze a simulated implementation of MPTCP in various situations. The results obtained from the simulator can hopefully give valuable insight into problems and possible solutions. As part of this thesis, a discrete event simulator was developed to simulate a network in an event-based fashion. The TCP and MPTCP implementations were tested using the simulated network, and the results are analyzed and compared using results from queueing theory.

Chapter 4

Method

This chapter will describe how the discrete event simulator, developed as part of the work behind this thesis, is implemented and explain the simplifications and assumptions needed to understand the obtained results. The simulator is used to analyze simulated implementations of both TCP and MPTCP in various situations.

4.1 The Discrete Event Simulator

The simulator created for this thesis is written in java using an object-oriented programming paradigm to structure the code as objects in a similar manner to how the internet actually looks like. For this reason, logic is implemented in the various classes:

- Router
- ClassicTCP
- MPTCP
- Channel
- Packet

In addition, the Router, ClassicTCP, and MPTCP classes are implementing the NetworkNode interface that defines the common methods needed to communicate with each other through Channel objects. Furthermore, the ClassicTCP and MPTCP classes are implementing the TCP interface as well to define the common methods needed in TCP communication. The Packet and Channel classes are not implementing any interface but are used throughout the simulator.

The use of a discrete event simulator allows for reproducible results as it is deterministic if the random sequence is known. This is achieved by using a single random generator with a pre-defined seed. The reproducibility of the results is key during analyses and for the reader to verify the obtained results. Further, real-time is not a factor in the simulator,

nor are hardware concerns which could be a decisive factor in other simulator techniques. To run the simulator, the `EventHandler` class employs a `PriorityQueue` that sorts `Event` objects in simulated time. Together these classes form the core of the simulation loop and will be described further in the following sections.

4.1.1 Events

Events are created to run the objects in the simulator. This ensures a distinction between the simulator and the simulated network. The objects public methods are used to run the objects as well as generating new events that are added to the call queue. This is structured around two main methods that all events are implementing, namely the `run` method and the `generateEvent` method (see Algorithm 1). The objects are responsible for the actual state changes that happen in the `run` method. In other words, the state of the system is stored in the networking objects, while the events only call the objects and enqueue new events based on the feedback from the objects themselves.

4.1.2 Event Handler

An event handler is the main controller of the simulation. It employs a call queue that keeps track of the next events to be called as well as storing future events. The queue is implemented using the standard `PriorityQueue` provided by java, with events sorted in time. Events are given timestamps (t) upon creation which is based on the objects delay-method (d), as well as the simulation time (T).

$$t = T + d$$

This operation ensures that events are scheduled in the correct order based on the type of event added and the time of creation. Further, if the next event is scheduled with a long delay, i.e., a long simulated time to the event should occur, the simulator can jump directly to this event without any waiting needed (see Algorithm 1 and Figure 4.1). This is doable due to the use of simulated time in place of actual time.

Algorithm 1: Main simulation loop

```

Function run ()
  while True do
    var event = eventQueue.poll();
    if event == null then
      return;
    Util.tickTime(event);
    event.run();
    event.generateNextEvent(eventQueue);

```

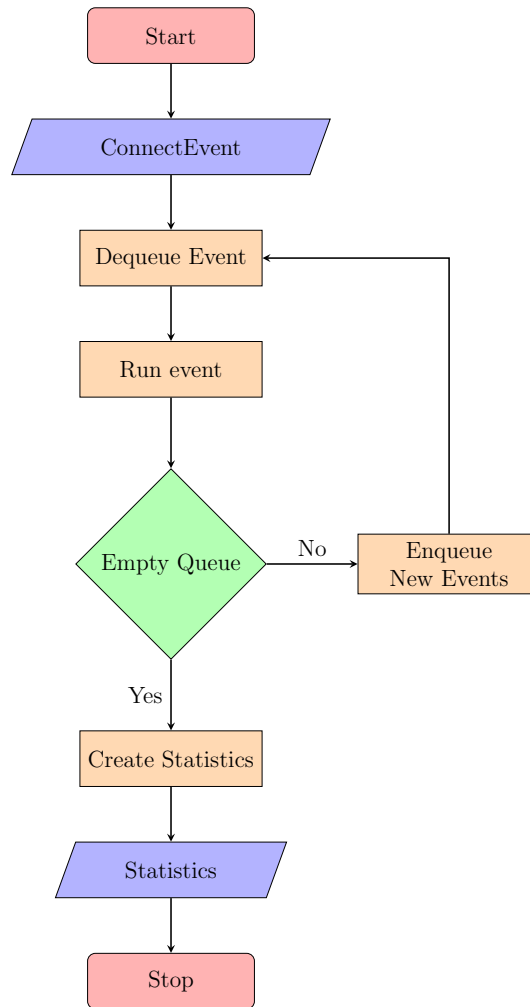


Figure 4.1: Simulator Flow Chart

4.1.3 Simulated Time

The simulator employs a logical clock in place of a physical clock. This logical clock ticks based on events and has no relation to real-world time. A logical clock is useful in simulation as hardware speeds compared to real-world time are eliminated, making results reproducible using different hardware. However, the delays described in the literature are based on real-world time. The equivalent to a real-world microsecond is set to 1 simulated time, meaning that the time unit used is microseconds. This is done in the effort to scale the delays according to each other. Consequently, all delays associated with events can be set to approximations of values in microseconds found in the literature.

Each time the event handler runs an event, the event timestamp is set to the new simulated time as shown in the main simulation loop algorithm 1. This ensures that events happen at the right time, as well as eliminating the possibility of scheduling events in the past. This also makes it easy to introduce randomness in the model without dramatically increasing

the complexity of the model [24].

4.1.4 Random Generator

A random generator is used throughout the model and is the only source of randomness. Randomly generated input is needed in modeling both the queue sizes and the loss in the model. It is important to generate the same random sequence in order to create reproducible results. This is achieved by using seeds given to the random generator and, in that way, reproducing the same pseudo-random generated sequence.

4.2 Router Modeling

The Router objects are tasked with the routing of packets, i.e., finding the right channel to send a Packet object through, as well as modeling the queuing and transmission delay. The actual routing operation is not important for the tests conducted in this thesis, as the networks created are linear paths. However, larger simulated experiments that needed this functionality were planned but are left for future work. Nevertheless, the queuing and transmission delays are what affect the output from the simulation. The following sections will describe how the two delays are modeled.

4.2.1 Queue

The presented model is using a simple Poisson process to model the router queue sizes to allow for comprehensible results. However, it is widely accepted that a Poisson process cannot accurately model the arrival rate of internet routers. The traffic seen in routers has a Long-Range Dependence (LRD) which means that the traffic has some sort of memory [25]. Considering that much of the traffic comes from TCP connections, it is not surprising that the perceived traffic distribution has some kind of memory. The TCP congestion control mechanisms are actively adjusting the sending rate to maximize throughput while not contributing to congestion. However, modeling the queues of routers has been shown to be difficult. Models that promise LRD properties and a closer fitting to real-life router queue sizes are generally complicated. They may have a lot of parameters and may only fit certain scenarios well [25]. Too realistic assumptions and complex models can make the analysis difficult.

Each time a new packet is arriving the modeled routers, the experienced queue size is equal to the Poisson random variable given by the Poisson process. Recall the PASTA property discussed in section 2.4.3. Since the Poisson process possesses this property, it is suitable to only generate the queue size when new arrivals from the modeled TCP endpoint arrive. That is, there is no need to keep track of the queue size in times outside of interesting arrivals because of the PASTA property.

4.2.2 Transmission and Queue Delay

There are three sources of delay that take place in the router of a network: transmission delay, queueing delay, and processing delay.

The transmission delay is the time it takes to push the bits in one packet onto the channel and is typically in the order of microseconds to milliseconds. The delay is proportional to the size of the packet L and the transmission rate R of the link, that is L/R [2]. The model does not differentiate the size of packets, nor does it have transmission rates. For this reason, the transmission delay is set to be a constant of 10.

The queueing delay is where the noticeable variation in packet delay originates from. This delay is the duration of time a packet waits in the queue while other packets are being transmitted. Arriving packets can experience queue sizes from zero to full; in the latter case, packets are dropped or lost. The varying queue size causes varying delays from zero to several milliseconds. For this reason, it is an interesting and important part of the simulator. Given the queue size experienced by the arriving packet $L_q^{(p)}$ and the previously mentioned transmission delay S the delay is set to be:

$$Delay = S + L_q^{(p)} \cdot S$$

where $L_q^{(p)}$ is modulated as a Poisson process with $\lambda = 100$.

The processing delay is commonly neglected because of the small, approximately constant delay in the order of microseconds or fewer [2]. However, in routers that perform complex payload modifications it can be of significance [26] [27]. This model will nonetheless assume a simple packet forwarding network, so the processing delay is omitted in the model.

4.3 Channel Modeling

The responsibility of the Channel objects is to deliver a Packet object to and from the linked NetworkNodes, in addition to model loss and propagation delay of the sent Packet object. NetworkNodes are objects that implement the NetworkNode interface, i.e., Router, ClassicTCP, or MPTCP objects. Loss and propagation delay manifest themselves in the output of the simulation and will be described in the following sections.

4.3.1 Propagation Delay

Propagation delay is the delay associated with a channel. This delay is closely approximated to be constant and is directly dependent on the length of the channel and the propagation speed of the link. The medium in which signals are sent is what decides the propagation speed, where fiber optics are the fastest. Regardless, the speed is a little less than the speed of light. In wide-area networks, propagation delay is said to be in the order of milliseconds. [2]. The model has a propagation delay set to be between 10

and 1000, depending on the cost associated with the channel. The cost is a parameter of the channel and can be thought of as both length and medium combined, where a high cost leads to a longer propagation delay. Although it is hard to exactly carry over the propagation delay to a simulation with a simulated time, the delay should be appropriate compared with the other delays in the network.

4.3.2 Packet Loss

Packet loss occurs when one or more packets of data fail to reach the destination. Loss can be caused by routing changes, link degradation, or congestion in the network [28]. The particular types of loss are often experienced as burst loss, i.e., loss follows loss with a high probability.

A 2-state hidden Markov model is used to model the loss that occurs in a channel, namely the Gilbert-Elliot Burst Model [29] [30]. This model is widely used for describing error patterns in transmission channels due to its flexibility and simplicity. The Gilbert-Elliot burst model achieves its burstiness by semantically having a "Good" and a "Bad" state (see Figure 4.2). Being in the good state means that loss rarely happens while being in the bad state means that loss occurs with a higher probability. Also, the probability of transitioning from the good state to the bad state is relatively low compared to the probability of going from the bad state to the good state, as shown in the figure. The properties as mentioned above, along with the Markovian property Theorem 2.4.10 gives the model its characteristics which have shown to be appropriate in emulating packet loss patterns [31] [32].




Figure 4.2: Gilbert-Elliot Burst Model

Implementation

Considering the simplicity of the Gilbert-Elliot burst model, the model is implemented with a simple check of the current state of the channel (see Algorithm 2). If the channel is in the Good-state, the channel loss parameter is used for both the transition probability of the Markov chain as well as the probability for loss. If the channel is in the bad state, then there is a 20% chance to stay in the Bad-state as well as packets being lost. This motivates the following transition matrix:

$$\begin{bmatrix} 1 - p & p \\ 0.8 & 0.2 \end{bmatrix} \quad (4.1)$$

Where p is the *loss* parameter in each channel in the network, hence, it is possible to calculate the fraction of time spent in the different states by using the known *loss* and the stationary equations Theorem 2.4.11. This result can again be used to calculate the expected loss per channel while having a burst loss model.

Algorithm 2: Gilbert-Elliot burst model implementation

```

Function loss()
  if goodState then
    goodState = Random.double() ≥ loss;
    return Random.double() < loss;
  goodState = Random.double() ≥ 0.5;
  return Random.double() < 0.5;

```

4.4 TCP Modeling

TCP modeling is achieved by implementing a class called ClassicTCP that employs the needed functions to simulate a TCP communication. Some of the logic concerning sending and receiving are separated out as own classes to reduce the complexity of ClassicTCP class. The following sections will describe how the ReceivingWindow and the SendingWindow are implemented and how this affects the results, as well as describe the delays defined in the Implementation of ClassicTCP.

4.4.1 Receiving Window

The receiving window implements the Flow Control mechanism, Selective Repeat ARQ, in the TCP model. The TCP objects in the model hold a receiving window that should be interchangeable with various ARQ schemes; however, only the Selective Repeat ARQ is implemented. The receiving window's responsibility is to check if incoming packets are ordered and if the packet should be stored in the buffer. Also, it is deciding which packet to acknowledge. If a packet is out of order but inside the defined receiving window, the packet is stored in the buffer. This can happen in the model due to packet loss. When the packet or packets blocking the stored packet arrives, the stored packet is moved from the buffer to the list of received packets. A single ACK is sent, acknowledging the last packet added to the received packet list, i.e., acknowledging the stored packet if it was the only stored packet.

The TCP model also holds a sending window that also should be interchangeable with other ARQ schemes. This object is created when a connection is made and controls how fast TCP should send packets. Since the receiving window capacity is the upper bound of packets possible to send without receiving an acknowledgment, the receivers receiving window capacity is communicated during the connection phase. This ensures that the receiving window capacity is not exceeded.

4.4.2 Retransmission

Retransmission of packets is performed if the transmitter receives three Duplicate Acknowledgements or through a Retransmission Timeout (RTO). Usually, the TCP endpoints are run as events when packets arrive; however, retransmissions are events that should occur without a packet triggering the event. Consequently, a retransmit event is scheduled with a delay equal to the RTO, which checks if the packet should be retransmitted or not. The event then triggers a retransmit if necessary. In contrast, retransmissions caused by duplicate ACKs can fast-retransmit the packet in the event triggered by the third DupACK.

RTO

RFC6298 [33] describe the initial calculation of the RTO value as follows:

$$SRTT = R \tag{4.2}$$

$$RTTVAR = \frac{R}{2} \tag{4.3}$$

$$RTO = SRTT + \max\{G, K \cdot RTTVAR\} \tag{4.4}$$

Where R is the initial RTT measurement, $SRTT$ is a smoothed round-trip time, $RTTVAR$ is round-trip time variation, G is a clock granularity, and K is set to 4. The initial value of the RTT and RTO is calculated during connection, but in contrast to the real-world implementations, the RTT and RTO values are not updated. The $SRTT$ and $RTTVAR$ are consequently not needed. The clock granularity G is neglected, so the RTO is:

$$RTO = R + 4 \cdot \frac{R}{2} \tag{4.5}$$

$$RTO = 3 \cdot R \tag{4.6}$$

The RTO of the model is, for this reason, set to $3 \cdot RTT$.

4.4.3 Sending Window

A sending window is created when a connection is made to match the ARQ scheme of the receiver and to operate the congestion control mechanisms. The TCP model simply calls the sending window to send the appropriate number of packets according to the congestion control mechanism and the ARQ scheme which is operated within the sending window. Consequently, the sending window operates the sender-side TCP flow control mechanism, as well as the congestion control mechanism, which is implemented only on the sender's side.

Congestion Control

The Congestion Control mechanism adjusts the sending window size or Congestion Window (CWND), effectively adjusting the sending rate. The adjustments are made according

to the ACKs received or by the absents of receivedACK and the eventual retransmit caused by the RTO. Therefore the model allows for communication between the sending window and the receiving window objects. Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery are all implemented, along with the possibility of using the Tahoe or Reno implementation

4.4.4 Transmission Delay

The transmission delay of the modeled TCP is set to be the same as for the routers. Since there is no concept of transmission rates, the transmission delay should be set equal to how to router's transmission delay is set. Further, since all packets have the same size, the transmission delay is constant at both routers and the TCP endpoints. The transmission delay is set to a constant of 10. 4.2.2

4.5 MPTCP Modeling

The MPTCP model must be able to send TCP-like traffic along different paths. For this reason, it was decided to implement the modeled MPTCP using The ClassicTCP class as the subflows, i.e., separated TCP flows per path. This approach is somewhat similar to using an existing TCP implementation on multiple interfaces. However, some modifications to the ClassicTCP implementation were needed. All the ClassicTCP subflows share a receiving window, while the sending window is isolated in the ClassicTCP subflows to separate the congestion control mechanisms as much as possible. The following sections will describe how the MPTCP implementation differs from the ClassicTCP implementation.

4.5.1 Sending Window

Since each subflow in MPTCP is operating on different paths, Each TCP subflow in the model has its own congestion control mechanisms. If congestion is detected in one path, it should not directly affect the sending rate of the other active subflows. The congestion control mechanisms are implemented in the sending window, so it follows that each TCP subflow in the MPTCP model has its own sending window. The congestion control mechanisms are therefore responding to ACK and RTO in the same way as the normal TCP implementation. In fact, using real-life existing TCP implementations as subflows limit the way data can be balanced over the available subflows. Packets are retained in the sending buffer until ACK is received and cannot directly be moved from one subflow to another to offload a subflow if serious congestion is detected [11].

Packet scheduler

The model does not implement a scheduler as is common in many MPTCP implementations [11] [21] [22]; however, the TCP subflows are all sending packets from a common list of packets to send. Subflows with shorter RTT and better throughput will consequently

send more packets, and the other subflows are offloaded. Then again, since all packets should be received in order at the shared receiver window, the lack of scheduler can result in sub-flows being bottle-necked by other slower subflows, i.e., forced to wait on a packet from another subflow.

4.5.2 Receiver Window

The receiver window is shared between all the subflows to ensure that packets are received in order. The Linux kernel implementation does the same, where all subflows compete against each other. This can cause unreasonable distribution of the receive buffer, which is undesirable. For this reason, a pre-division of the receive buffer is performed in the model [23]. However, unlike other pre-division methods that use the RTT to divide the receiving window, this model assigns equal receiving window size for all subflows.

4.5.3 Delay

The transmission delay is set to be the same as for the standard TCP and Router implementations, since the subflows of the MPTCP model are TCP subflows. which is a delay of 10. The processing delay is throughout this model neglected. This may be a simplification which is not suitable when analyzing MPTCP as the receiver window can, with benefit, be implemented with more complexity. Packets in normal TCP are assumed to be received in order [11], this assumption is not valid in the case of MPTCP where packets often arrive out of order. This can cause higher processing delay, and more complex data structures can help this problem. The processing delay is nevertheless neglected.

4.6 Collecting Results

To collect results, a statistics class is used to do event-based bookkeeping as described in subsection 2.4.6 and Table 2.5 To retrieve results from the simulation runs. The arrival and departure times are collected and used to calculate the interarrival times of packets, the time spent in the system per packet, and the number of packets in the system. Since the service time, i.e., the transmission delay is a constant 10, this is not collected. Also, the time in the system and time in the queue of a packet is only differing with a constant 10, so the time in the queue is not needed. Additionally, the time service starts (U) is reformulated as $U^{(n)} = D^{(n)} - S^{(n)}$ and is used to calculate the time in system using just departure and arrival time Table 4.1. Further, graphs are made to visualize arrival and departure times, the number of packets in the system, the time of packets in the system, and the interarrival times. In addition, TCP statistics such as goodput, which is the number of useful packets delivered in per time unit, and loss rate are collected, along with a graph of the sender's congestion window.

Variable	Definition	Calculation
$A^{(n)}$	Arrival time of packet n	Collected
$D^{(n)}$	Departure time of packet n	Collected
$L^{(t)}$	Number of packets in the system at time t	Collected
$S^{(n)}$	Service time of packet n	Constant 10
$T^{(n)}$	Interarrival time between packet n and packet $n + 1$	$T^{(n)} = A^{(n+1)} - A^{(n)}$
$U^{(n)}$	Time packet n starts service	$U^{(n)} = D^{(n)} - S^{(n)}$
$W^{(n)}$	Time packet n is in the system	$W^{(n)} = D^{(n)} - A^{(n)}$
$W_q^{(n)}$	Time packet n is in the queue	$W_q^{(n)} = W^{(n)} - S^{(n)}$
T	Number of packets sent (Transmitted)	Collected
RT	Number of packets retransmitted	Collected
FRT	Number of packets fast retransmitted	Collected
Goodput	The information throughput	$Goodput = T/t$
Loss rate	The perceived loss	$Loss\ rate = (RT + FRT)/T$

Table 4.1: Collected and Calculated Results

Chapter 5

Results

In this chapter, the results gathered from the simulation are presented. The setup of the various tests is described in detail before the results from the simulated TCP and MPTCP implementations are presented.

5.1 Setup

The tests which led to the results are performed using equal values of the parameters throughout the different experiments with the exception of the paths being different as well as the use of both TCP and MPTCP. In this way, the importance of the parameters that changed from test to test could be observed. All routers are using a Poisson queue length process with different means to simulate the queues and consequently the queueing delay per packet. During development, it was found that 2000 packets is a good amount of packets to send since it is enough that loss can occur, even in environments that have a low probability of loss, and allows for clearer graphs as more packets will compact the timeline and consequently obscure the results. Further, it simulates downloading a file of size roughly equal to 2-3GB, if the average packet size is assumed to be 1300 byte [34]. Even though the measure of packet size does not affect the simulator in any way, it can give the reader and the results something tangible to relate to. For this reason, the average packet size of 1300 byte is used whenever packet size is referenced in the results. The simulated time is divided by 1000, making the results a measure in "milliseconds" instead of "microseconds" to make the results more readable. Additionally, only the TCP Reno implementation is presented since the Tahoe and Reno results did not differ in a substantial way.

5.1.1 Path Variations

Four variations of paths is the main difference between all the test done on MPTCP and TCP. Each path has its own properties that will show how the implementations work in different environments. The properties that are different in the four path variation are the

length and loss parameters, giving four different variations. When testing the MPTCP implementation, the heterogeneous path tests are using combinations of the declared tests, while the homogeneous tests are using two of the same paths. The results can, in turn, be compared to the TCP implementation used on the same path. The properties of the variation of the paths will be presented in the following sections.

Short Path Low Loss

This path is using four routers, with router 1 and 2 having a 50% queue utilization while router 2 and 3 have a 98% queue utilization. The loss parameter is set in only one channel to 0.001, which should give a loss of a little more 0.12% when calculating the fraction of time being at the two-state the Gilbert-Elliot channel Equation 4.1 and multiplying the probability of loss in each state.

$$0.999 \cdot 0.001 + 0.001 \cdot 0.2 = 0.001199 \quad (5.1)$$

$$\approx 0.12\% \quad (5.2)$$

However, in the results that will be presented, the loss will be higher since this is the implementation's perceived loss. This is the number of retransmissions divided by the number of packets sent as shown in Table 4.1. The cost of the channels is set to 1 in the first and last channel, while the inner channels have a cost of 10.

Short Path High Loss

This path is equal to the short path with low loss (section 5.1.1) except that the loss parameter is set to 0.01 giving the channel a loss of approximately 1.2% when taking the Gilbert-Elliot burst channel into consideration.

$$0.988 \cdot 0.01 + 0.012 \cdot 0.2 = 0.01228 \quad (5.3)$$

$$\approx 1.2\% \quad (5.4)$$

Long Path Low Loss

This path is set up using ten routers, with the outermost routers at both ends being at 50% queue utilization while the rest is at 98%. This makes it similar to the short path but with extra routers in between. Additionally, channels 1 and 11 have a cost of 1, channels 2 and 10 have a cost of 10, while the rest have a cost of 50, making the path substantially longer. Further, 3 of the channels connecting the inner routers are lossy, in contrast to the single lossy channel in the short paths. This loss is set to the same as in the short path low loss test, but this time in 3 channels making it more probable that loss will occur. Using the Gilbert-Elliot burst channel transition matrix from Equation 4.1 and the stationary equations Theorem 2.4.11, the fraction of time in each state is found. Further, since loss can occur three places on the path, the overall probability of loss is the reverse of not having loss at all:

$$1 - (0.999 \cdot 0.999 + 0.001 \cdot 0.8)^3 = 0.00359 \quad (5.5)$$

$$\approx 0.36\% \quad (5.6)$$

The higher loss compared with the short path with low loss section 5.1.1 is expected in real life since longer paths have more possibilities of experiencing interference, congestion, and other situations leading to loss.

Long Path High Loss

This path is equal to the long path with low loss section 5.1.1 with the exception of the loss parameters. Instead of having three channels with the same loss of 0.001, two of the channels now have a loss of 0.001 and one with a loss of 0.01, which is the same loss as in the short path high loss section 5.1.1. The loss for this channel is:

$$1 - ((0.999 \cdot 0.999 + 0.001 \cdot 0.8)^2(0.988 \cdot 0.99 + 0.01 \cdot 0.8)) = 0.01465 \quad (5.7)$$

$$\approx 1.5\% \quad (5.8)$$

5.1.2 Simulation Results

General results are presented in separate tables per test. The arrival plot, departure plot, interarrival plot, number of packets in system plot, time in system plot, and the congestion window size plot are all graphs that visualize how the implementation behaves in time in the given environment. The arrival plot and departure plot are describing the arrival and departure time of the n-th packet. This is interesting because it highlights the fact that arrivals and departures are essentially the same in TCP. The interarrival plot is plotting the interarrival time between packets received. This is interesting because it should be easy to spot when loss occurs as well as seeing the effect of the sending window. The number of packets in the system and the amount of packets in the system is of interest because it shows how full the receiving window is, as well as highlighting the effect lost packets has on the time out of order packets are held in the system. Lastly, the congestion window plot is showing the congestion control mechanisms. In light of the other graphs, the congestion window is especially interesting as it is a key part of how TCP works. The MPTCP results are presented as separate subflows since the performance difference or lack of difference is interesting to analyze. Consequently, there are two tables and double the number of graphs when MPTCP results are presented.

5.2 TCP

Tests were performed on simulated environments to validate the TCP implementation used in the MPTCP model. The tests included environments with high and low loss, as well as long and short paths, resulting in four tests described in the following sections.

5.2.1 Short Path Low Loss

This test is using the declared short path with low loss section 5.1.1. From the graphs, it can be seen that loss was detected two times. The congestion window is perhaps the best

indicator Figure 5.6. The first experienced loss is serious enough to trigger the RTO seen by the appearance of the slow start mechanism, i.e., an exponential growth until ssthresh is reached. The second experienced loss is not as serious as seen by the lack of slow start, but rather the fast recovery mechanism being initiated due to 3 or more DupACKs. All the other graphs are observations made at the receivers end and are especially interesting to look at in the light of the sender's congestion window.

From the graphs listed below, it is clear that loss is a vital part of the performance of TCP. In Figure 5.1, Figure 5.2, Figure 5.3, Figure 5.4, Figure 5.5, and Figure 5.6, a packet loss is evident at the same simulated time represented by the x-axis. Given the loss of 0.12%, the expected amount of packets to be lost when 2000 packets are sent is $2000 \cdot 0.0012 = 2.4$. Looking at the graphs, at least two loss bursts can be seen.

Variable	Result
λ	3.093 <i>t</i>
W	0.053 <i>t</i>
L	0.164 <i>Packets</i>
<i>Packets Sent</i>	2000 <i>Packets</i>
<i>ACKs Received</i>	1804 <i>Packets</i>
<i>RT</i>	32 <i>Packets</i>
<i>FRT</i>	9 <i>Packets</i>
<i>Goodput</i>	3.641 <i>Packets/t</i>
<i>Loss rate</i>	2%

Table 5.1: TCP Short Path Low Loss Results

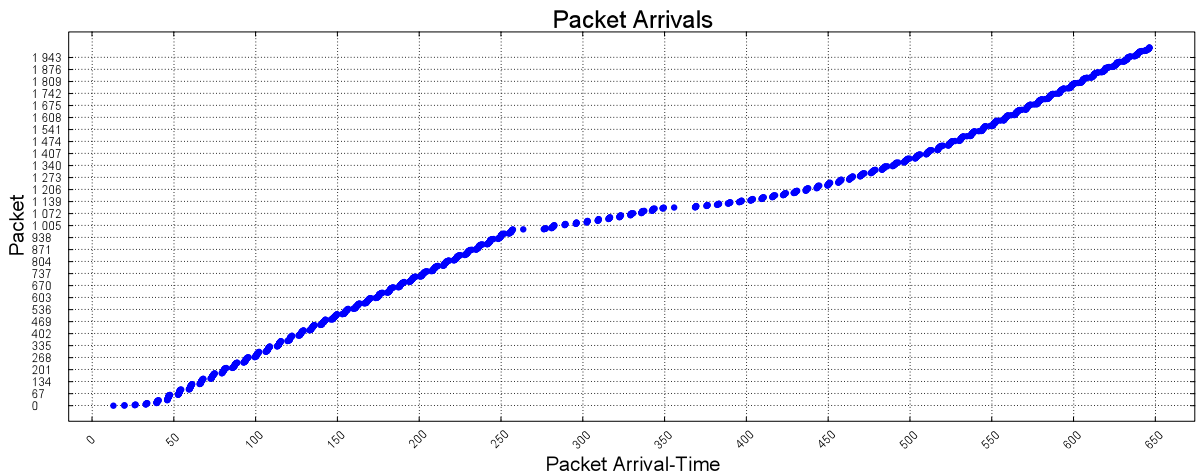


Figure 5.1: TCP Short Path Low Loss Packet Arrivals




Figure 5.2: TCP Short Path Low Loss Packet Departures

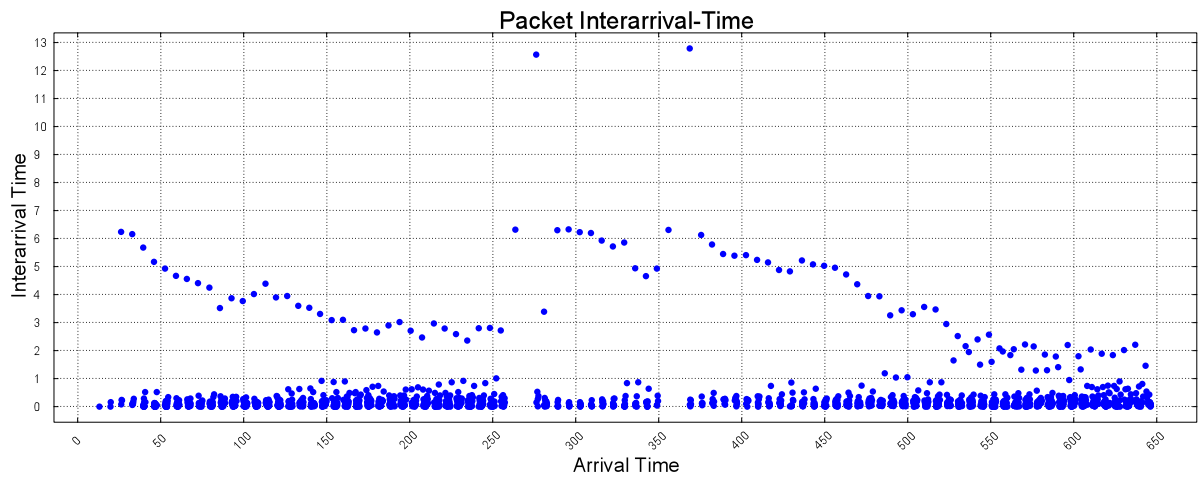


Figure 5.3: TCP Short Path Low Loss Packet Interarrival Time




Figure 5.4: TCP Short Path Low Loss Packets in System

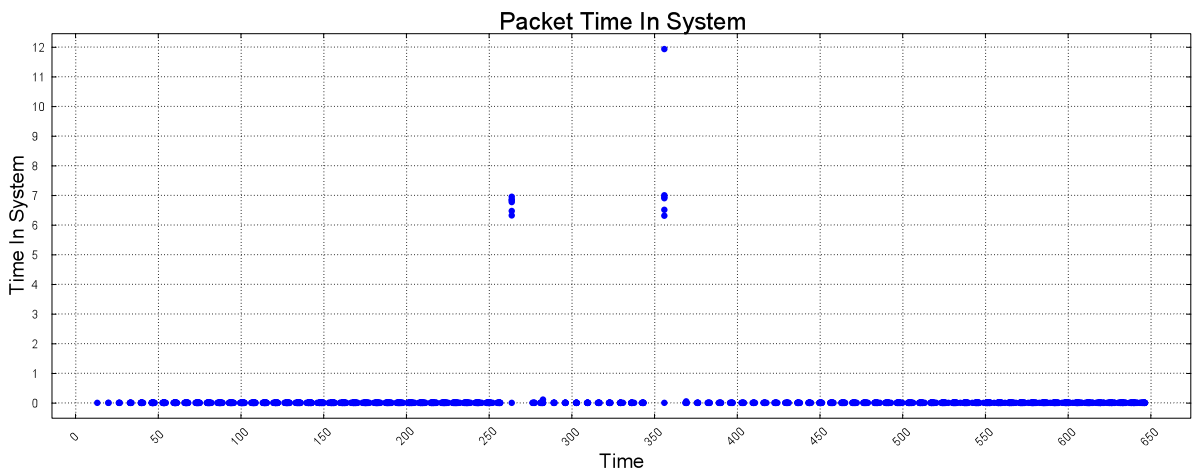


Figure 5.5: TCP Short Path Low Loss Packet Time in System




Figure 5.6: TCP Short Path Low Loss Congestion Window

5.2.2 Short Path High Loss

This test is using the short path with high loss section 5.1.1. Comparing this test with the test with low loss, the congestion window in this test does not stabilize at the RWND upper bound, apart from the small time frame during slow start. This is because loss often occurs, resulting in the reduction of the congestion window. The fast recovery mechanism can be observed in the congestion window plot Figure 5.12. Other points of interest are the increase in both retransmissions and fast retransmissions, resulting in a higher loss as well as more ACKs received compared with the previous test. Additionally, the average time in the system W and the average number of packets in system L are also increased. This is due to more packets being lost, and consequently, the packets need to be retransmitted, causing a block in the receivers end, resulting in a higher average time and number of packets in the system. This is can be seen by looking at the Packets in system graph Figure 5.10 and the time in system graph Figure 5.11. The interarrival times shown in Figure 5.9 are as expected with less predictability caused by the increase in loss.

Variable	Result
λ	1.11 <i>t</i>
W	0.59 <i>t</i>
L	0.658 <i>Packets</i>
<i>Packets Sent</i>	2000 <i>Packets</i>
<i>ACKs Received</i>	1995 <i>Packets</i>
<i>RT</i>	133 <i>Packets</i>
<i>FRT</i>	55 <i>Packets</i>
<i>Goodput</i>	1.109 <i>Packets/t</i>
<i>Loss rate</i>	9.4%

Table 5.2: TCP Short Path High Loss Results

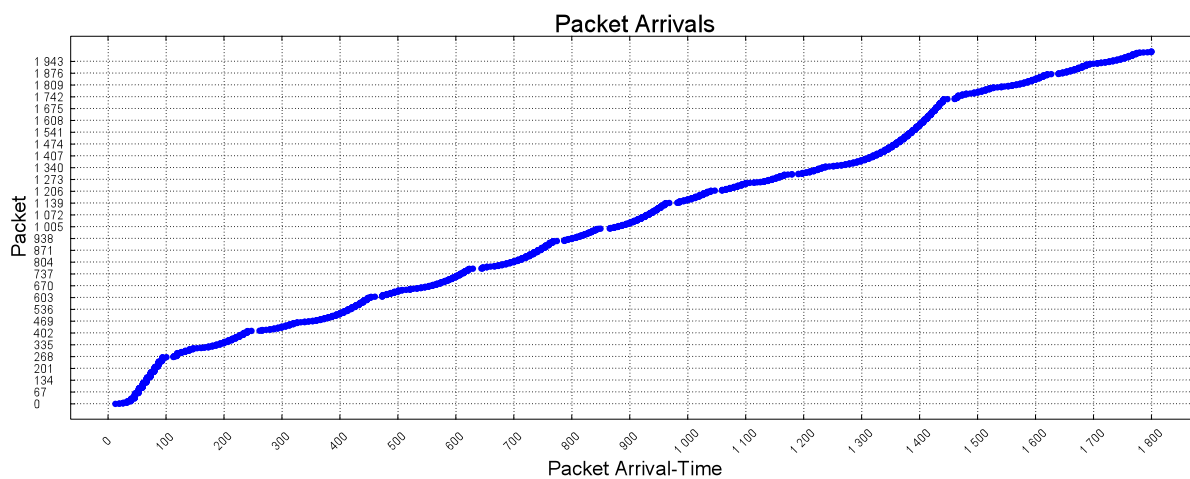


Figure 5.7: TCP Short Path High Loss Packet Arrivals




Figure 5.8: TCP Short Path High Loss Packet Departures

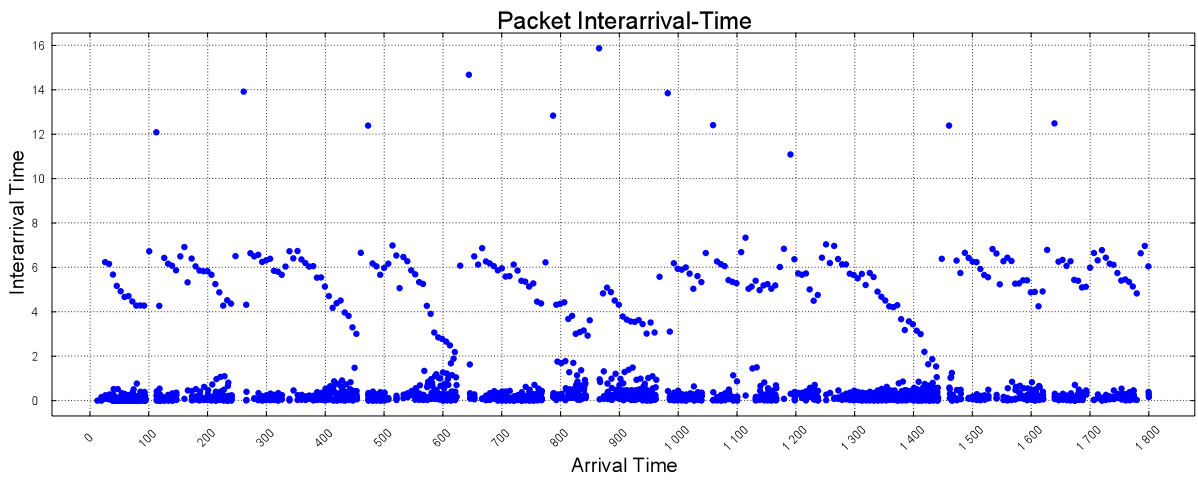


Figure 5.9: TCP Short Path High Loss Packet Interarrival Time




Figure 5.10: TCP Short Path High Loss Packets in System




Figure 5.11: TCP Short Path High Loss Packet Time in System

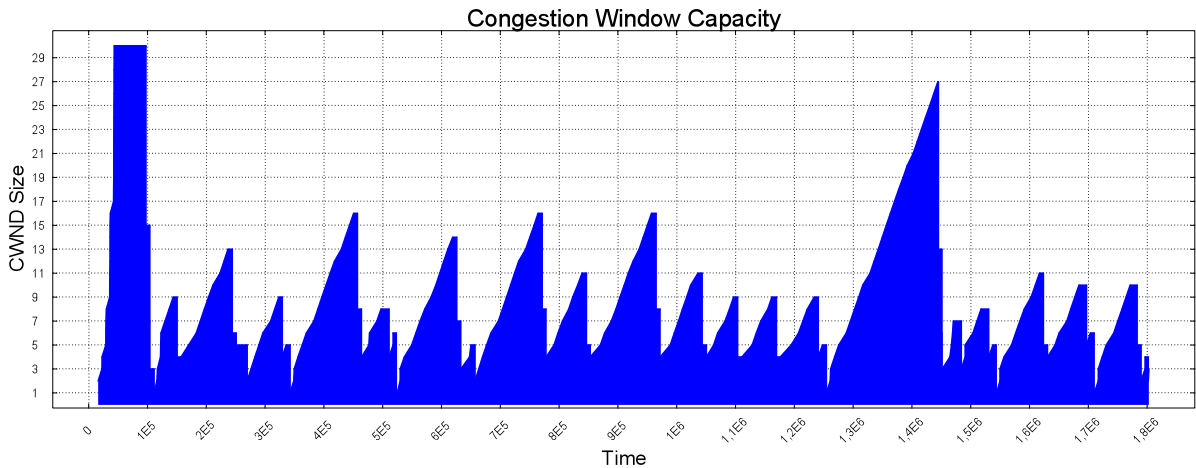


Figure 5.12: TCP Short Path High Loss Congestion Window

5.2.3 Long Path Low Loss

This test is using the long path with low loss section 5.1.1. Again looking at just the congestion window plot Figure 5.18 it is clear that there are four bursts of loss, where four have been serious enough to trigger RTO, while fast retransmission has solved one loss burst. Numerical results are presented in Table 5.3. One important point of interest is the dramatic decrease in goodput. This is caused by the longer path inherently reducing the number of packets being sent per time unit.

Variable	Result
λ	0.623 <i>t</i>
W	0.317 <i>t</i>
L	0.197 <i>Packets</i>
<i>Packets Sent</i>	2000 <i>Packets</i>
<i>ACKs Received</i>	1985 <i>Packets</i>
<i>RT</i>	67 <i>Packets</i>
<i>FRT</i>	24 <i>Packets</i>
<i>Goodput</i>	0.632 <i>Packets/t</i>
<i>Loss rate</i>	4.6%

Table 5.3: TCP Long Path Low Loss Results




Figure 5.13: TCP Long Path Low Loss Packet Arrivals




Figure 5.14: TCP Long Path Low Loss Packet Departures

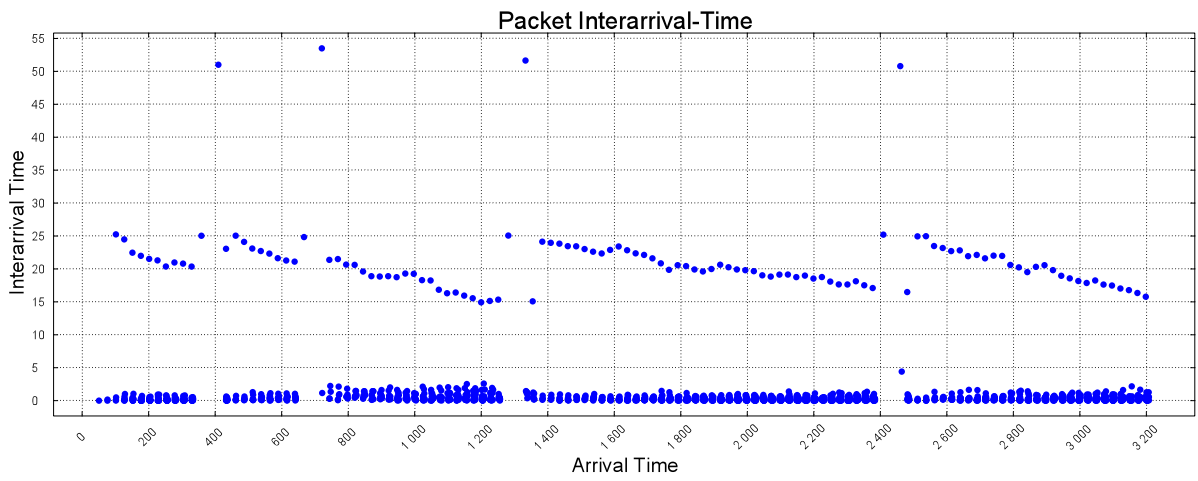


Figure 5.15: TCP Long Path Low Loss Packet Interarrival Time

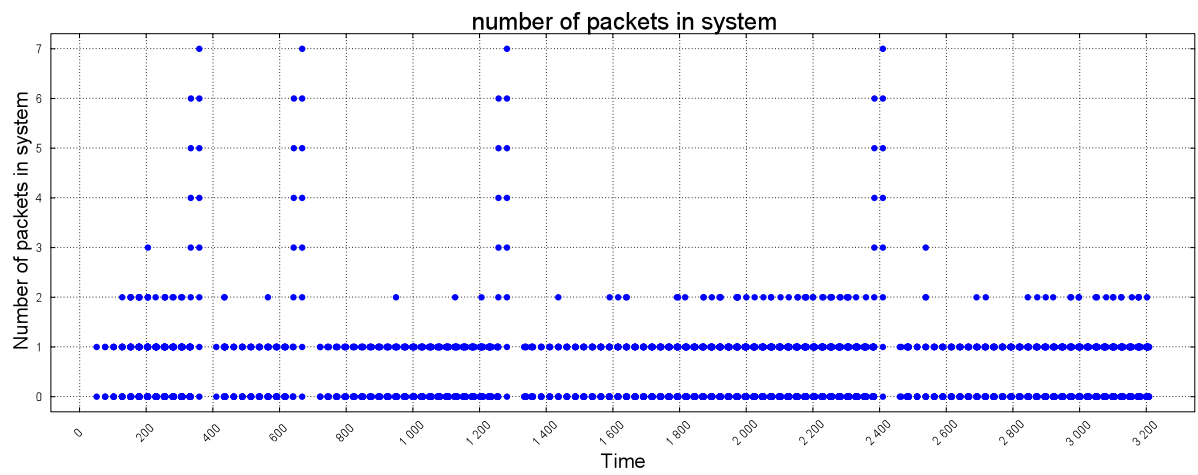


Figure 5.16: TCP Long Path Low Loss Packets in System

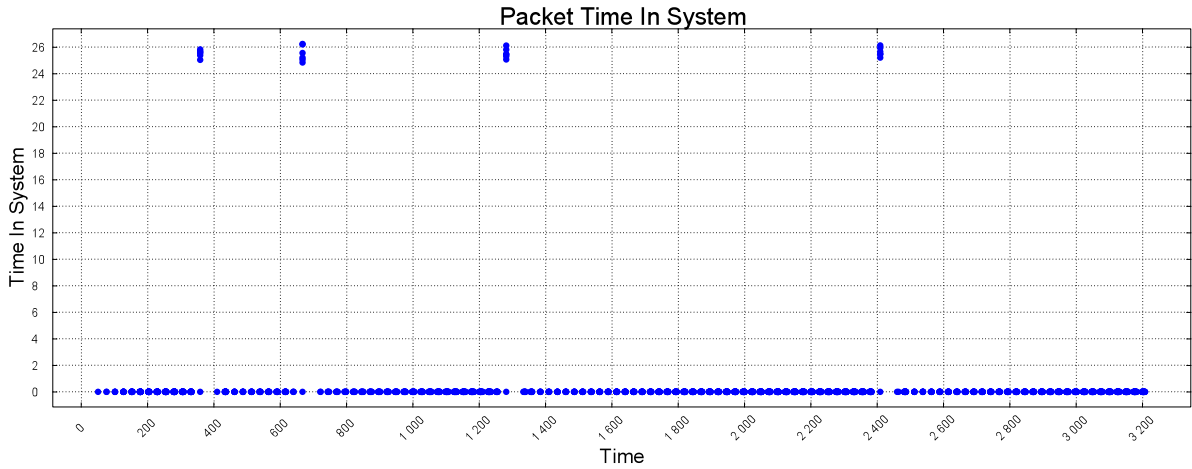


Figure 5.17: TCP Long Path Low Loss Packet Time in System




Figure 5.18: TCP Long Path Low Loss Congestion Window

5.2.4 Long Path High Loss

This test is using the long path with high loss section 5.1.1. Comparing this test with the results from the short path with high loss and the long path with low loss, the results are not that surprising. This is the most unforgiving path, with the highest loss and longest path leading to bad goodput. This is also the first test where the number of ACKs is higher than the number of packets sent.

Variable	Result
λ	0.223 t
W	3.523 t
L	0.787 Packets
Packets Sent	2000 Packets
ACKs Received	2074 Packets
RT	132 Packets
FRT	64 Packets
Goodput	0.224 Packets/t
Loss rate	9.8%

Table 5.4: TCP Long Path High Loss Results

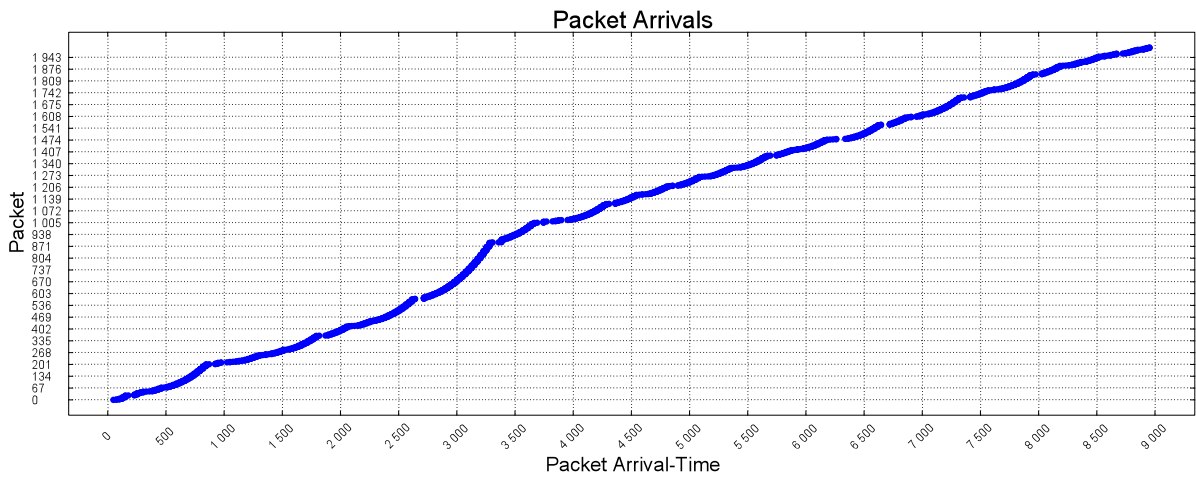


Figure 5.19: TCP Long Path High Loss Packet Arrivals




Figure 5.20: TCP Long Path High Loss Packet Departures




Figure 5.21: TCP Long Path High Loss Packet Interarrival Time

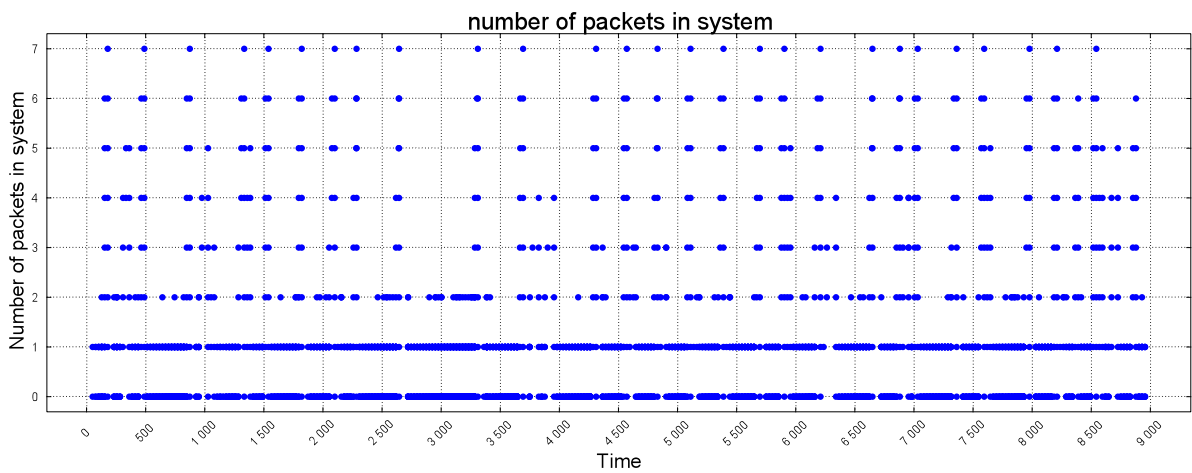


Figure 5.22: TCP Long Path High Loss Packets in System




Figure 5.23: TCP Long Path High Loss Packet Time in System




Figure 5.24: TCP Long Path High Loss Congestion Window

5.3 MPTCP

MPTCP is different from TCP in that it is using at least two paths to communicate. This raises issues concerning the behavior differences between homogeneous and heterogeneous paths. Using MPTCP over homogeneous paths, or paths with equal properties, will, in theory, distribute the traffic equally, effectively doubling the throughput compared to using just one of the paths, as is the case for TCP. However, real-life MPTCP situations will more often than not use heterogeneous paths, paths with different properties. All the presented tests have been completed using two subflows as this is the most likely use case in practice. However, the implementation has been tested using several subflows, but this did not dramatically change the overall results. The following sections will test the

implemented MPTCP simulator on multiple environments with high and low loss, long and short path, and homogeneous and heterogeneous paths.

5.3.1 Homogeneous Paths Low Loss

The results achieved for homogeneous paths in the long and short path environments introduced in section 5.2 are essentially the same when adjusting for the added delay that comes from the longer path. For this reason, only the results from the short path are shown. Firstly the two subflows are not distributing the load equally, with subflow 0 sending 1217 out of 2000 packets while subflow 1 is sending 783 packets out of 2000 seen in Table 5.5. This is a problem since the two subflows have equal prerequisites. However, testing with 20000 yields roughly the same packet offset. In addition, the combined goodput of the two subflows is, in fact, lower than the TCP baseline.

The loss of the two subflows is at 2.8% and 4.9%, respectfully. This is not as expected since the standard TCP implementation experienced a 2% loss in the same environment. Further, there are no retransmissions caused by the RTO, but rather an excessive amount of fast retransmissions. This raises an issue that will become clearer as more loss and heterogeneous paths are introduced. When looking at the congestion windows Figure 5.30 it is clear that the subflows are getting unnecessary response and are therefore reducing the congestion window. This response comes from packets received from the other subflow, effectively triggering DupACKs to be sent and consequently interfering with the other subflow. The slow start mechanism at the start of subflow 1's congestion window graph stops much earlier than subflow 0, giving more evidence of interference in the congestion control mechanisms.

The subflows of MPTCP must in some way interfere with each other since packets must arrive in order. Looking at the arrival graph Figure 5.25, departure graph Figure 5.26, interarrival graph Figure 5.27, and especially the number of packets in system graph Figure 5.28 and the packet time in the system graph Figure 5.29 it is clear that the subflows are not operating independently of each other.


(a) Subflow 0

Variable	Result
λ	1.478 <i>t</i>
W	0.193 <i>t</i>
L	0.285 <i>Packets</i>
<i>Packets Sent</i>	1217 <i>Packets</i>
<i>ACKs Received</i>	1287 <i>Packets</i>
RT	0 <i>Packets</i>
FRT	34 <i>Packets</i>
<i>Goodput</i>	1.473 <i>Packets/t</i>
<i>Loss rate</i>	2.8%


(b) Subflow 1

Variable	Result
λ	0.955 <i>t</i>
W	0.180 <i>t</i>
L	0.172 <i>Packets</i>
<i>Packets Sent</i>	783 <i>Packets</i>
<i>ACKs Received</i>	958 <i>Packets</i>
RT	0 <i>Packets</i>
FRT	38 <i>Packets</i>
<i>Goodput</i>	0.953 <i>Packets/t</i>
<i>Loss rate</i>	4.9%

Table 5.5: MPTCP Homogeneous Paths Low Loss Results




(a) Subflow 0




(b) Subflow 1

Figure 5.25: MPTCP Homogeneous Paths Low Loss Packet Arrivals




(a) Subflow 0




(b) Subflow 1

Figure 5.26: MPTCP Homogeneous Paths Low Loss Packet Departures




(a) Subflow 0




(b) Subflow 1

Figure 5.27: MPTCP Homogeneous Paths Low Loss Packet Interarrival Time




(a) Subflow 0




(b) Subflow 1

Figure 5.28: MPTCP Homogeneous Paths Low Loss Packets in System




(a) Subflow 0




(b) Subflow 1

Figure 5.29: MPTCP Homogeneous Paths Low Loss Packet Time in System



(a) Subflow 0



(b) Subflow 1

Figure 5.30: MPTCP Homogeneous Paths Low Loss Congestion Window

5.3.2 Homogeneous Paths High Loss

For the same reason as discussed in subsection 5.3.1 the results from the short path will be presented rather than give the results for both the long and short path. This test is using the short path high loss path section 5.1.1 for both subflows. The results seen in Table 5.6 are repeating the key take away from the previous test (subsection 5.3.1). Subflow 1 has over double the amount of perceived loss as subflow 0, which increases the difference compared to the previous test with low loss. This can again be connected to the fact that excessive DupACKs are sent, which is evident in the congestion graph and especially during slow start. This tells that increased loss will contribute to more redundant retransmissions.


However, looking at the number of packets sent in Table 5.6, higher loss contributes to

a better load balancing between the two subflows. Perhaps related is the fact that the combined goodput is substantially better compared with the results from normal TCP on the same path (subsection 5.2.2). Individually though, normal TCP outperforms the MPTCP subflows. Nevertheless, this test is promising compared to the test using homogeneous paths with low loss (subsection 5.3.1).


One possible reason that this test has better goodput and balancing compared with the previous test might be that the increased loss limits the dominant subflow and consequently allows the other subflow to send uninterrupted for a while. Another possible explanation might be the effect the increased loss has on the random generation. The additional loss will contribute to even more calls to the random generator, making it unclear if the results are due to lucky randomness or implementation related.

(a) Subflow 0		(b) Subflow 1	
Variable	Result	Variable	Result
λ	0.92 <i>t</i>	λ	0.622 <i>t</i>
W	0.84 <i>t</i>	W	0.990 <i>t</i>
L	0.773 <i>Packets</i>	L	0.615 <i>Packets</i>
<i>Packets Sent</i>	1191 <i>Packets</i>	<i>Packets Sent</i>	809 <i>Packets</i>
<i>ACKs Received</i>	1418 <i>Packets</i>	<i>ACKs Received</i>	1391 <i>Packets</i>
<i>RT</i>	25 <i>Packets</i>	<i>RT</i>	26 <i>Packets</i>
<i>FRT</i>	73 <i>Packets</i>	<i>FRT</i>	135 <i>Packets</i>
<i>Goodput</i>	0.918 <i>Packets/t</i>	<i>Goodput</i>	0.623 <i>Packets/t</i>
<i>Loss rate</i>	8.2%	<i>Loss rate</i>	19.9%

Table 5.6: MPTCP Homogeneous Paths High Loss Results




(a) Subflow 0




(b) Subflow 1

Figure 5.31: MPTCP Homogeneous Paths High Loss Packet Arrivals




(a) Subflow 0




(b) Subflow 1

Figure 5.32: MPTCP Homogeneous Paths High Loss Packet Departures




(a) Subflow 0




(b) Subflow 1

Figure 5.33: MPTCP Homogeneous Paths High Loss Packet Interarrival Time




(a) Subflow 0




(b) Subflow 1

Figure 5.34: MPTCP Homogeneous Paths High Loss Packets in System




(a) Subflow 0




(b) Subflow 1

Figure 5.35: MPTCP Homogeneous Paths High Loss Packet Time in System



(a) Subflow 0



(b) Subflow 1

Figure 5.36: MPTCP Homogeneous Paths High Loss Congestion Window

5.3.3 Heterogeneous Paths Low Loss

This test is using both the short path with low loss (section 5.1.1) and the long path with low loss (section 5.1.1), making it a test of MPTCP on heterogeneous paths with low loss. The general results presented in Table 5.7 show that this simulated implementation of MPTCP performs poorly when the paths are of unequal length. While load balancing between the two paths is somewhat equal with subflow 0 at 1198 packets sent and subflow 1 at 802 packets sent, subflow 1 should have sent most of the packet as this flow sends over the short path. With this in mind, it is clear that subflow 0 is the dominant subflow that severely damages the potential of subflow 1.

The results generated from subflow 0 are somewhat equal to the standard TCP results. Almost all metrics are worse but not unexpected or abnormal, except for no RTO retrans-

missions and a congestion window that seems to stutter. Subflow 1, on the other hand, is producing completely unexpected results. None of the graphs resemble the standard TCP tests, and the performance metrics are flawed. 595 fast retransmissions, a 76.7% loss rate, and 3348 AKCs received (Table 5.7) are results that does not make sense


(a) Subflow 0

Variable	Result
λ	0.365 <i>t</i>
<i>W</i>	0.957 <i>t</i>
<i>L</i>	0.35 <i>Packets</i>
<i>Packets Sent</i>	1198 <i>Packets</i>
<i>ACKs Received</i>	1386 <i>Packets</i>
<i>RT</i>	0 <i>Packets</i>
<i>FRT</i>	68 <i>Packets</i>
<i>Goodput</i>	0.367 <i>Packets/t</i>
<i>Loss rate</i>	5.7%


(b) Subflow 1

Variable	Result
λ	0.245 <i>t</i>
<i>W</i>	3.422 <i>t</i>
<i>L</i>	0.84 <i>Packets</i>
<i>Packets Sent</i>	802 <i>Packets</i>
<i>ACKs Received</i>	3348 <i>Packets</i>
<i>RT</i>	20 <i>Packets</i>
<i>FRT</i>	595 <i>Packets</i>
<i>Goodput</i>	0.245 <i>Packets/t</i>
<i>Loss rate</i>	76.7%

Table 5.7: MPTCP Heterogeneous Paths Low Loss Results




(a) Subflow 0




(b) Subflow 1

Figure 5.37: MPTCP Heterogeneous Paths Low Loss Packet Arrivals




(a) Subflow 0




(b) Subflow 1

Figure 5.38: MPTCP Heterogeneous Paths Low Loss Packet Departures




(a) Subflow 0




(b) Subflow 1

Figure 5.39: MPTCP Heterogeneous Paths Low Loss Packet Interarrival Time




(a) Subflow 0




(b) Subflow 1

Figure 5.40: MPTCP Heterogeneous Paths Low Loss Packets in System




(a) Subflow 0




(b) Subflow 1

Figure 5.41: MPTCP Heterogeneous Paths Low Loss Packet Time in System



(a) Subflow 0



(b) Subflow 1

Figure 5.42: MPTCP Heterogeneous Paths Low Loss Congestion Window

5.3.4 Heterogeneous Paths High Loss

This test is using both the short path with high loss and the long path with high loss, making it a test of MPTCP on heterogeneous paths with high loss. It is natural to compare this test to the previous test on heterogeneous paths with low loss. Surprisingly, this test performs better when it comes to balancing the packets equally, with the short path (subflow 1) sending more packets this time. Subflow 1 also has a surprisingly lower loss rate, while subflow 0 has a much higher loss rate, as is expected. All results point to a better performing subflow 1 and a worse performing subflow 0 compared to the previous test. The results are for this reason pointing to better cooperation between the subflows when the loss of the channels are higher.

As discussed earlier in subsection 5.3.1 and subsection 5.3.2 the excessive amount of

retransmissions can cause problems and limit the other subflow. If the dominant subflow limits the other subflow, then loss occurring in the dominant subflow may allow the other subflow to send without as much interference. This may be the reason why subflow 1 is performing better in this test compared to the previous test subsection 5.3.3.


(a) Subflow 0

Variable	Result
λ	0.236 <i>t</i>
W	2.554 <i>t</i>
L	0.603 <i>Packets</i>
<i>Packets Sent</i>	948 <i>Packets</i>
<i>ACKs Received</i>	1249 <i>Packets</i>
RT	21 <i>Packets</i>
FRT	101 <i>Packets</i>
<i>Goodput</i>	0.235 <i>Packets/t</i>
<i>Loss rate</i>	12.9%


(b) Subflow 1

Variable	Result
λ	0.263 <i>t</i>
W	3.749 <i>t</i>
L	0.985 <i>Packets</i>
<i>Packets Sent</i>	1052 <i>Packets</i>
<i>ACKs Received</i>	3439 <i>Packets</i>
RT	68 <i>Packets</i>
FRT	559 <i>Packets</i>
<i>Goodput</i>	0.263 <i>Packets/t</i>
<i>Loss rate</i>	59.6%

Table 5.8: MPTCP Heterogeneous Paths High Loss Results




(a) Subflow 0




(b) Subflow 1

Figure 5.43: MPTCP Heterogeneous Paths High Loss Packet Arrivals




(a) Subflow 0




(b) Subflow 1

Figure 5.44: MPTCP Heterogeneous Paths High Loss Packet Departures




(a) Subflow 0




(b) Subflow 1

Figure 5.45: MPTCP Heterogeneous Paths High Loss Packet Interarrival Time




(a) Subflow 0




(b) Subflow 1

Figure 5.46: MPTCP Heterogeneous Paths High Loss Packets in System



(a) Subflow 0



(b) Subflow 1

Figure 5.47: MPTCP Heterogeneous Paths High Loss Packet Time in System

5.4.1 Modeling Loss

Loss is modeled using Gilbert-Elliot burst channels implemented in the channels as explained in section 4.3. In the short paths (section 5.1.1 and section 5.1.1), the loss parameter is set in one channel, giving the path one model for loss in contrast to the long paths (section 5.1.1 and section 5.1.1) that have three channels with loss. The literature [31] [31] [32] only describes the use of Gilbert-Elliot channels as modeling the overall loss, however my implementation is using multiple burst models when using the long paths. Since communication over longer paths is prone to more error, more variation in RTT, and less predictability, it can be justified that multiple burst models make sense. Then again, this usage is not described in any literature found.

Further, the loss that occurs in real life is often due to congestion. However, when a wireless link is part of the path, noise and interference are often the cause of loss. The model does not differentiate the cause of loss. As a result, the modeled loss has no relation to the router's congestion at all. This is not realistic since many of the losses that occur in wired connections are from dropped packets. For this reason, it can be argued that there should be a relation between queuing delay and congestion in the network; however, this is not implemented.

Lastly, the Gilbert Elliot burst model is a simplification of how loss occurs. Other loss models could have been used, such as a realistic fading model [32]. That being said, the Gilbert-Elliot channel is simplistic, which in itself is beneficial.

5.4.2 Modeling Router Queue Size

The number of packets situated in the routers is modeled using a Poisson process section 4.2. The Poisson process is known not accurately to model queues of routers [25]. However, as discussed earlier, it is chosen for its simplicity and because it is widely used as an arrival process in simulation. The queuing delay is described as the factor that has the most variance, thus contributing the most to the unpredictability of internet communication. For this reason, it is a problem that the accuracy of the Poisson process is inadequate. Also, as discussed in subsection 5.4.1 the size of the queues, and consequently the queuing delay, is not contributing to the loss modeling. As a result, the modeling of queue size and the modeling of queuing delay are not good enough to correctly model internet traffic. Nevertheless, for this simulation, it is sufficient as simplifications are made in every part of the simulator. As seen in section 5.2 the result for TCP resembles something similar to what can be thought of as communication over the internet with random components adding to the results. However, given more time, a better model with a relation between queue size and loss could be implemented.

5.4.3 Packet Size Distribution

It was decided not to model the size of packets to keep the simulator as simple as possible. That is why the packets in the model do not have a packet size associated with them.

However, in actual TCP traffic, the packet sizes are varying a lot causing varying transmission delays and processing delays [34]. In addition, small packets like ACKs should have less probability of error as they can more easily fit the buffers of congested routers and be less exposed to interference. In the simulator, ACKs are handled the same way as packets that contain larger payloads. For the reasons above, the modeled packets size, or rather the lack of modeled packets size, is a simplification of the model and needs to be taken into consideration when evaluating the results

5.4.4 Processing Delay

For the purpose of this simulator, it was decided to neglect the processing delay as it is described to be in the order of microseconds or less [2]. Other sources are, however, describing the processing delay to reach the magnitude of long-distance propagation delay in some cases where complex payload operations, content adaptation, or encryption is needed [26]. Regardless, the fact that the packet processing delay is neglected is a shortcoming of the model and must be considered when evaluating the presented results.

5.4.5 TCP as Subflow

Using TCP directly in a naive way presented some expected issues and unexpected issues that notably reduced the potential performance of the MPTCP implementation.

Scheduling the TCP subflows according to RTT or other performance metrics to receive packets in order is an active research topic and was not implemented [22] [21]. Another solution to this is a pre-division of the receiving buffer based on RTT [23]. The presented model uses the latter but divides the buffer in half regardless of RTT or other metrics. This simplifies the model, limiting the overall performance of MPTCP, especially in heterogeneous paths where delay differences become essential. Consequently, the time in system metric W increases in the heterogeneous test compared to homogeneous tests, being an indicator of packets needing to wait for other packets, i.e., packets out of order.

An unexpected effect caused by using TCP directly in an MPTCP implementation was ARQ combined with congestion control. TCP is implemented to use the acknowledgments or an absence of acknowledgment as indicators of loss, giving ACKs two responsibilities, namely, (1) avoiding congestion and being fair to other users, and (2) providing flow control that guarantees that packets are delivered in the correct order. In an MPTCP use case, this is not working. To illustrate the problem, suppose we have subflow A. A has sent all its packets and is waiting for an ACK. This ACK will signal the A to send more packets. The receiver has received all packets from A, but the packets from another subflow B are missing. Suppose the packets from A are to be received before packets from B. Then only a DupACK can be sent to the sender, effectively telling the sender that loss has occurred even though it's an out-of-order problem. When B finally can send all packets successfully, A must receive an ACK to, in turn, continue to send as normal. This leads to a precarious situation where the subflows must trigger ACKs going to other subflows.

At the same time, ACKs are flow specific and will trigger congestion control mechanisms. One of the goals of MPTCP is to effectively use two or more paths to send data faster and more reliable. This implementation does not achieve this. The main reason is the ambiguity of an ACK and how MPTCP is implemented in the simulator.

5.4.6 Discrete Event Simulation

Discrete event simulators have some properties that make them useful when simulating events that occur in time. This is presented in detail in section 4.1. However, simulation and modeling is, by nature, a simplification of actual processes. For this reason, it should be analyzed with this in mind.

Time is modeled using a logical clock with an imaginary unit simulated as microseconds to scale the delays appropriate to each other. The time unit is not essential, but the proportional delays are. The delays, or time until an event happens, are modeled using approximations from the literature [26] [2], and will inherently have smaller inconsistencies. The high dynamic variance in queuing delays and the variance between channel propagation delays are hard to model. Also, the transmission and processing delays are small but should vary based on packet size and other factors. Together the mentioned inconsistencies can add up to significant disproportions that can affect the validity of the results. Since the delays are what drives time in the model, and the fact that the modeled delays are inconsistent with real life, the simulated time can not be compared with real-life time. Consequently, results that depend on time, such as goodput, can not be compared with results from outside the simulated environment.

5.5 Evaluation of Data

The data gathered from the test runs are based on the implementation of the discrete event simulator and the way MPTCP and TCP were implemented. As discussed in section 5.4 there are several simplifications made to the model as well as the fact that the simulator is only replicating the behavior of real-life communication over the internet and consequently how TCP and MPTCP will behave.

The tests performed on TCP in section 5.2 was in mainly completed to test the implementation of TCP and the underlying network put together by routers and channels. The results and behavior of standard TCP are well known and were a good way to validate how the network behaved. With results that came close to the expedited output, the model showed promise.

The MPTCP implementation lacks several important features to model how a real-life implementation could look like. This taken into consideration, the simulation results are pointing to many of the known problems or obstacles that MPTCP faces. Combined with the promising results from the TCP implementation, this is, to a degree, validating the model. That being said, the results from the simulation must be seen in the light of all simplifications and imperfections that come with simulation and modeling as a whole.

Further, the data should only be used to find broader relations and should not focus on details.

Chapter 6

Conclusion

This chapter present the conclusion of the thesis including several finding in addition to suggested future work.

6.1 Research Conclusion

This thesis presents a discrete event simulator used in modeling networks with different properties to analyze MPTCP. The results are compared with an implementation of TCP over the same paths to validate and support the achieved results. To a large extent, the gathered data are just scratching the surface of bigger and more complicated issues associated with the making of MPTCP. Despite the inherent simplifications in the mode, many of the papers mentioned in chapter 3 are pointing to the same problems as the results derived from the simulator. Papers presenting results from simulation and real-world tests point to the out-of-order problem and the load imbalance problem. The results from the tests conducted in this thesis are showing the same. Both issues have the potential to reduce the usability of MPTCP significantly. In addition, both problems can amplify each other, making it imperative to mitigate the problems. Furthermore, the aforementioned DupACK problem has unfolded a significant characteristic of MPTCP that fundamentally change the way flow control and congestion control should behave. The following sections will describe the issues in detail and explain how the problems can arise, how it's detected, and how to potentially overcome the problems.

6.2 Load Imbalance Problem

The load imbalance problem is the problem of balancing the sending rate so that packets are received as ordered as possible, even if the packets come from different paths with different proprieties. If this is not handled according to the performance of each subflow, packets from one flow may arrive too early, causing the out-of-order problem. This problem becomes evident in the presented simulation when looking at the heterogeneous path

tests in subsection 5.3.4 and subsection 5.3.3. The subflow sending over the shorter path is sending significantly faster but is held back by a slower subflow caused by the lack of scheduling or buffer division. The slower subflow, on the other hand, is trying to catch up with the faster subflow but will always lag behind, causing a degradation of the potential goodput that can be achieved.

6.3 Out-of-Order Problem

The out-of-order problem is caused by receiving packets out of order. Packets must be received in order; consequently, the out-of-order packets must wait in the receiving buffer until the block is resolved. In addition, packets must be reordered at the receiver's side, making the communication more costly. Further, the packet waiting time is increased, which leads to higher latency. A bad load balancing at the sender side, or rather lack of balancing, will result in more packets received out of order, especially in heterogeneous networks. This leads to unnecessary buffer bloat and higher latency caused by the subflow which is lagging behind, ultimately reducing the potential goodput. This problem is particularly evident in the time in system and number of packets in system graphs for the heterogeneous path tests in subsection 5.3.4 and subsection 5.3.3. Comparing the time in the system graph, the number of packets in system graphs, the average time in the system W , and the average number of packets in the system L to the paths with homogeneous paths, it is clear that packets are to a higher degree received out of order. This is due to the better balancing of load when the paths are equal due to the naive 50-50 split of the receive buffer.

6.4 Redundant ACK

The redundant ACK problem explained in subsection 5.4.5 is caused by the ambiguity of ACKs in TCP. This can be seen as small spikes in the congestion window plots caused by the reduction of CWND provoked by the received ACK. Implementing MPTCP should, as a consequence, be able to distinguish the meaning of the received ACK. A concept of seen packets has been proposed by Sundararajan et al.[35] as a solution to the ambiguity of ACKs in coded TCP. This solution can be used in the context of MPTCP as well. The implementation can distinguish between ACKs representing successful sending and ACKs indicating seen packets from the other subflows if this concept is used. ACKs triggered from the other subflows, or rather information from the other subflows, will with this concept never interfere with other subflows congestion control mechanisms, effectively solving the problem. That being said, this concept does not solve the other introduces obstacles alone.

6.5 Possible Solutions

As discussed, the literature and the results gathered from the implemented simulator are pointing to the two presented problems, namely the out-of-order problem and the load imbalance problem. The given redundant ACK problem can be solved using concepts from the coded-TCP research. It turns out that using error-correcting codes can indeed solve, or at least alleviate some of the issues concerning the two problems [19] [20]. Presented in chapter 3 are two papers suggesting network coding as a solution to ordering and load balancing, i.e., the out-of-order problem and load imbalance problem. Additionally, since lost or delayed packets can be reconstructed, the need for retransmissions is reduced. For the same reason, packet reordering is less of an issue.

The Linux kernel implementation has achieved a working and deployable solution to the problems mentioned. However, this implementation which is based on RFC8684 [9], is outperformed in some papers [23] [19] showing that there is still room for improvements. From the results gathered in this thesis and the promise that MPTCP with network code shows, it seems to be a good solution to the presented problems.

6.6 Future Work

Due to lack of time, some adaptations, implementations, and experiments have been left for future work. Several different ideas would have been interesting to analyze; however, completing this was not feasible due to the time constraints. Therefore, the main focus of this thesis was to finish a working discrete event simulator, implement a versatile simulated network with interchanging possibilities, and implement TCP and MPTCP for analyzing purposes. The versatility of the simulated network was important during the early development as modeling different loss and arrival processes to the routers were planned. As the deadline came closer, though, these planes were neglected and thus left as future work. The redundant ACKs discussed in several sections earlier are probably the main priority to fix, as the results would be more precise. Also, a sending scheduler or a better pre-divining of receive-buffer is high on the priority list of improvements that could rather easily be implemented. Given more time, here are some of the ideas that could be interesting to test:

1. The queue utilization of routers on the internet is hard to model as most traffic is TCP traffic. TCP traffic has Long-Range Dependence and is adjusting sending rate according to the sending environment i.e., the path. With minor adaptations, the simulator could facilitate simulated TCP traffic over the routers using multiple TCP implementations, possibly making the model more realistic regarding queueing delay and congestion in routers. In addition, this allows for analyzing the impact MPTCP has on the performance of the router. Using results from Jackson networks and the Pollaczek-Khinchin equations described in queueing theory [12] and section 2.4.5, many essential metrics of the routers queueing systems could be analyzed, giving results about the network as a whole.

2. Another point of interest is the importance of packet sizes regarding transmission and processing delays. For this reason, it could be interesting to use a distribution to model the packet sizes and use this when modeling the transmission and processing delay. By doing this, the added benefit of having processing delay as a factor is added, which in some cases is described to reach the magnitude of long-distance propagation delay [26].
3. The current implementation of MPTCP uses a receive buffer with pre-division-based flow control. However, the implementation is not basing the split on any metric but is splitting equally amongst the available subflows. It would be of interest to test out using different metrics like RTT to separate the buffer according to the performance of the subflows. Alternatively, it could be interesting to implement a sending scheduler and compare the results to the other implementations. Looking at how this will affect the load imbalance problem and inherently the out-of-order problem is of great interest.
4. Finally, the promise of MPTCP with network codes is compelling and would be interesting to analyze and compare against the results achieved in this thesis. This is perhaps the most demanding task of the ones described; however, the results can help in the evolution of the protocol.

Nevertheless, the simulator has produced results that raise the same issues as other researchers have found. Finding reasonable solutions to the problems is, however, a far more difficult task. More research is therefore needed in the field to find the optimal solutions.

Acronyms

4G Fourth generation of broadband cellular network technology. 7, 15

5G Fifth generation of broadband cellular network technology. 1, 7

ACK Acknowledgement. 12, 13, 39–41, 50, 56, 90, 91, 94, 95

AIMD Additive-Increase/Multiplicative-Decrease. 13

ARQ Automatic Repeat Request. 12, 13, 39, 40, 90

CWND Congestion Window. 13–15, 40, 94

DupACK Duplicate Acknowledgement. 12–15, 40, 47, 60, 67, 93

FEC Forward Error Correction. 12

FIFO First In First Out. 16, 24, 25

FMTCP Fountain Code-Based MultiPath Transmission Control Protocol. 31

GB Gigabyte. 44

IP Internet Protocol. 7, 9, 10

LRD Long-Range Dependence. 36, 95

MPTCP MultiPath TCP. 1, 7–9, 15, 16, 31–33, 37, 41, 42, 44–46, 59, 60, 68, 74, 81, 90, 91, 93–96

MTU Maximum Transmission Unit. 10

NAT Network Address Translation. 16

PASTA Poisson Arrivals See Time Averages. 21, 25, 36

RTO Retransmission Timeout. 13–15, 40, 41, 47, 53, 60, 74

RTT Round Trip Time. 14, 40–42, 89, 90, 96

RWND Receiver Window. 12, 13, 50

SYN Synchronize. 12

TCP Transmission Control Protocol. 1, 9–12, 14–16, 31–33, 36, 39–42, 44–47, 59, 60, 68, 74, 75, 89–91, 93–95

UDP User Datagram Protocol. 10

VoIP Voice over Internet Protocol. 7

WiFi A family of wireless network protocols. 7, 15, 31

Bibliography

- [1] Robert T. Braden. *Requirements for Internet Hosts - Communication Layers*. RFC 1122. Oct. 1989. DOI: 10.17487/RFC1122. URL: <https://rfc-editor.org/rfc/rfc1122.txt>.
- [2] James F. Kurose and Keith W. Ross. *Computer Networking: A Top-Down Approach (6th Edition)*. 6th. Pearson, 2012. ISBN: 0132856204.
- [3] Go Hasegawa, Masayuki Murata, and Hideo Miyahara. “Fairness and stability of congestion control mechanisms of TCP.” In: *Telecommunication Systems* 15 (Nov. 2000), pp. 167–184. DOI: 10.1023/A:1019186710820.
- [4] Lorenzo Vicisano et al. *The Use of Forward Error Correction (FEC) in Reliable Multicast*. RFC 3453. Dec. 2002. DOI: 10.17487/RFC3453. URL: <https://rfc-editor.org/rfc/rfc3453.txt>.
- [5] Ying-zong Huang, Sanjeev Mehrotra, and Jin li. “A hybrid FEC-ARQ protocol for low-delay lossless sequential data streaming.” In: June 2009, pp. 718–725. DOI: 10.1109/ICME.2009.5202596.
- [6] *TCP Congestion Control*. 2017. URL: <https://yazilimcorbasi.blogspot.com/2016/12/tcp-congestion-control.html> (visited on 05/31/2021).
- [7] Ethan Blanton, Dr. Vern Paxson, and Mark Allman. *TCP Congestion Control*. RFC 5681. Sept. 2009. DOI: 10.17487/RFC5681. URL: <https://rfc-editor.org/rfc/rfc5681.txt>.
- [8] T Nishitha. “A Comparative Analysis of TCP Tahoe, Reno, New- Reno, SACK and Vegas in Homogeneous Networks.” In: 2002.
- [9] Alan Ford et al. *TCP Extensions for Multipath Operation with Multiple Addresses*. RFC 8684. Mar. 2020. DOI: 10.17487/RFC8684. URL: <https://rfc-editor.org/rfc/rfc8684.txt>.
- [10] Kae Won Choi et al. “Optimal load balancing scheduler for MPTCP-based bandwidth aggregation in heterogeneous wireless environments.” In: *Computer Communications* 112 (2017), pp. 116–130. ISSN: 0140-3664. DOI: <https://doi.org/10.1016/j.comcom.2017.08.018>. URL: <https://www.sciencedirect.com/science/article/pii/S0140366417302426>.
- [11] Costin Raiciu et al. “How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP.” In: *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. San Jose, CA: USENIX Association, Apr. 2012, pp. 399–412. URL: <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/raiciu>.

- [12] Donald Gross John F. Shortle James M Thompson and Carl M Harris. *Fundamentals Of Queueing Theory*. 5th ed. Wiley Series in Probability and Statistics. Wiley, 2018. ISBN: 9781118943526.
- [13] Ronald W. Wolff. “Poisson Arrivals See Time Averages.” In: *Operations Research* 30.2 (1982), pp. 223–231. DOI: 10.1287/opre.30.2.223. eprint: <https://doi.org/10.1287/opre.30.2.223>. URL: <https://doi.org/10.1287/opre.30.2.223>.
- [14] John D. C. Little. “A PROOF FOR THE QUEUING FORMULA: $L = \lambda W$.” In: (Nov. 1960). DOI: <http://fisherp.scripts.mit.edu/wordpress/wp-content/uploads/2015/11/ContentServer.pdf>.
- [15] D. V. Lindley. “The theory of queues with a single server.” In: *Mathematical Proceedings of the Cambridge Philosophical Society* 48.2 (1952), pp. 277–289. DOI: 10.1017/S0305004100027638.
- [16] Jerry Banks et al. *Discrete-Event System Simulation*. July 2009. ISBN: 9780136062127.
- [17] Keld Helsgaun. “Discrete Event Simulation in Java.” In: (June 2000).
- [18] et al. C. Paasch S. Barre. *MultiPath TCP - Linux Kernel implementation*. URL: <https://www.multipath-tcp.org/> (visited on 05/24/2021).
- [19] Yong Cui et al. “FMTCP: A Fountain Code-Based Multipath Transmission Control Protocol.” In: *IEEE/ACM Transactions on Networking* 23.2 (2015), pp. 465–478. DOI: 10.1109/TNET.2014.2300140.
- [20] Jason Cloud et al. “Multi-Path TCP with Network Coding for Mobile Devices in Heterogeneous Networks.” In: *2013 IEEE 78th Vehicular Technology Conference (VTC Fall)*. 2013, pp. 1–5. DOI: 10.1109/VTCFall.2013.6692295.
- [21] Vivek Adarsh, Paul Schmitt, and Elizabeth Belding. “MPTCP Performance over Heterogenous Subpaths.” In: July 2019, pp. 1–9. DOI: 10.1109/ICCCN.2019.8847086.
- [22] Vu Vu and Brenton Walker. “Redundant Multipath-TCP Scheduling with Desired Packet Latency.” In: Oct. 2019, pp. 7–12. ISBN: 978-1-4503-6933-6. DOI: 10.1145/3349625.3355440.
- [23] Jiangping Han et al. “Receive Buffer Pre-division Based Flow Control for MPTCP.” In: Mar. 2018, pp. 19–31. ISBN: 978-981-10-8889-6. DOI: 10.1007/978-981-10-8890-2_2.
- [24] Prateek Sharma. “Discrete-Event Simulation.” In: *INTERNATIONAL JOURNAL OF SCIENTIFIC & TECHNOLOGY RESEARCH* 4 (Apr. 2015).
- [25] L. Muscariello et al. “Markov models of internet traffic and a new hierarchical MMPP model.” In: *Computer Communications* 28.16 (2005), pp. 1835–1851. ISSN: 0140-3664. DOI: <https://doi.org/10.1016/j.comcom.2005.02.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0140366405000824>.
- [26] R. Ramaswamy, Ning Weng, and T. Wolf. “Characterizing network processing delay.” In: *IEEE Global Telecommunications Conference, 2004. GLOBECOM '04*. Vol. 3. 2004, 1629–1634 Vol.3. DOI: 10.1109/GLOCOM.2004.1378257.
- [27] Nicolas Hohn et al. “Bridging router performance and queuing theory.” In: vol. 32. June 2004, pp. 355–366. DOI: 10.1145/1005686.1005728.

- [28] Keyur Parikh and Junius Kim. “The Role of Network Packet Loss Modeling in Reliable Transport of Broadcast Audio.” In: (). URL: <https://www.gatesair.com/documents/papers/Parikh-K130115-Network-Modeling-Revised-02-05-2015.pdf>.
- [29] E. N. Gilbert. “Capacity of a burst-noise channel.” In: *The Bell System Technical Journal* 39.5 (1960), pp. 1253–1265. DOI: 10.1002/j.1538-7305.1960.tb03959.x.
- [30] E. O. Elliott. “Estimates of error rates for codes on burst-noise channels.” In: *The Bell System Technical Journal* 42.5 (1963), pp. 1977–1997. DOI: 10.1002/j.1538-7305.1963.tb00955.x.
- [31] Gerhard Haßlinger and Oliver Hohlfeld. “The Gilbert-Elliott Model for Packet Loss in Real Time Services on the Internet.” In: Jan. 2008, pp. 269–286.
- [32] Joseph Kang, Wayne Stark, and Alfred Hero. “Turbo Codes for Fading and Burst Channels.” In: (Feb. 1999).
- [33] Matt Sargent et al. *Computing TCP’s Retransmission Timer*. RFC 6298. June 2011. DOI: 10.17487/RFC6298. URL: <https://rfc-editor.org/rfc/rfc6298.txt>.
- [34] Rishi Sinha, Christos Papadopoulos, and John Heidemann. *Internet Packet Size Distributions: Some Observations*. Tech. rep. ISI-TR-2007-643. Originally released October 2005 as web page <http://netweb.usc.edu/%7ersinha/pkt-sizes/>. johnh: pafile: USC/Information Sciences Institute, May 2007. URL: <http://www.isi.edu/%5C%7ejohnh/PAPERS/Sinha07a.html>.
- [35] Jay Sundararajan et al. “Network Coding Meets TCP: Theory and Implementation.” In: *Proceedings of the IEEE* 99 (Apr. 2011), pp. 490–512. DOI: 10.1109/JPROC.2010.2093850.

Appendix

```
1 package org.example.data;
2
3
4 public enum Flag {
5     ACK,
6     SYN,
7     FIN;
8 }
```

```
1 package org.example.data;
2
3 public class Message implements Payload {
4
5     private String msg;
6
7     public Message(String msg) {
8         this.msg = msg;
9     }
10
11     @Override
12     public int size() {
13         return this.msg.length();
14     }
15
16     @Override
17     public String toString() {
18         return this.msg;
19     }
20
21     @Override
22     public boolean equals(Object obj) {
23         if (obj instanceof Payload) return this.toString().equals(obj.toString());
24         return false;
25     }
26
27     @Override
28     public int hashCode() {
29         return super.hashCode();
30     }
31 }
```

```

1 package org.example.data;
2
3 import org.example.network.interfaces.Endpoint;
4
5 import java.util.List;
6
7 public class Packet {
8     private final Endpoint destination;
9     private final Endpoint origin;
10    private final List<Flag> flags;
11    private final Payload payload;
12    private final int index;
13
14    private final int sequenceNumber;
15    private final int acknowledgmentNumber;
16
17
18    protected Packet(Endpoint destination, Endpoint origin, List<Flag> flags,
19        Payload payload, int sequenceNumber, int acknowledgmentNumber, int index) {
20        this.destination = destination;
21        this.origin = origin;
22        this.flags = flags;
23        this.payload = payload;
24        this.index = index;
25
26        this.sequenceNumber = sequenceNumber;
27        this.acknowledgmentNumber = acknowledgmentNumber;
28    }
29
30    public boolean hasAllFlags(Flag... flags) {
31        if (flags == null || this.flags == null) {
32            throw new NullPointerException("flags can't be null");
33        }
34
35        if (this.flags.isEmpty()) {
36            // flags given are not in the list
37            return false;
38        }
39
40        if (flags.length == 0) {
41            // no flags given implies that all flags given are in list
42            return true;
43        }
44
45        boolean hasFlag = this.flags.contains(flags[0]);
46        for (var i = 1; i < flags.length; i++) {
47            hasFlag &= this.flags.contains(flags[i]);
48        }
49        return hasFlag;
50    }
51
52    public Payload getPayload() {
53        return payload;
54    }
55
56    public Endpoint getDestination() {
57        return this.destination;
58    }
59
60    public Endpoint getOrigin() {
61        return this.origin;
62    }

```

```

62
63
64     public int getSequenceNumber() {
65         return this.sequenceNumber;
66     }
67
68     public int getIndex() {
69         return this.index;
70     }
71
72     public int getAcknowledgmentNumber() {
73         return acknowledgmentNumber;
74     }
75
76     public int size() {
77         if (this.payload == null) return 0;
78         return payload.size();
79     }
80
81     protected List<Flag> getFlags() {
82         return flags;
83     }
84
85     @Override
86     public String toString() {
87         String returnString;
88         if (this.payload == null) {
89             returnString = this.flags.toString();
90         } else {
91             returnString = "[" + this.payload.toString() + "];
92         }
93         returnString += "[seq: " + this.getSequenceNumber() + "]" + "[index: " +
94             this.getIndex() + "];
95
96         return returnString;
97     }
98
99     @Override
100    public int hashCode() {
101        return super.hashCode();
102    }
103
104    @Override
105    public boolean equals(Object o) {
106        if (this == o) return true;
107        if (o == null || getClass() != o.getClass()) return false;
108        var packet = (Packet) o;
109        return sequenceNumber == packet.sequenceNumber
110            && acknowledgmentNumber == packet.acknowledgmentNumber
111            && destination.equals(packet.destination)
112            && origin.equals(packet.origin)
113            && packet.hasAllFlags(this.flags.toArray(Flag[]::new));
114    }

```

```

1 package org.example.data;
2
3 import org.example.network.interfaces.Endpoint;
4 import org.example.protocol.Connection;
5
6 import java.util.ArrayList;
7 import java.util.List;
8
9
10 public class PacketBuilder {
11     private Endpoint destination = null;
12     private Endpoint origin = null;
13     private List<Flag> flags = new ArrayList<>();
14     private Payload payload = null;
15     private int index = 0;
16
17     private int sequenceNumber = -1;
18     private int acknowledgmentNumber = -1;
19
20     public Packet build() {
21         if (!this.hasFlag(Flag.ACK)) {
22             this.acknowledgmentNumber = -1;
23         }
24         return new Packet(this.destination, this.origin, this.flags, this.payload,
25             this.sequenceNumber, this.acknowledgmentNumber, this.index);
26     }
27
28     public Packet ackBuild(Packet packetToAck) {
29         if (packetToAck.getOrigin() == null) throw new IllegalArgumentException("no
30             origin");
31         if (packetToAck.getDestination() == null) throw new IllegalArgumentException
32             ("no destination");
33
34         this.withDestination(packetToAck.getOrigin());
35         this.withOrigin(packetToAck.getDestination());
36         this.withFlags(Flag.ACK);
37         this.withPayload(null);
38         this.withAcknowledgmentNumber(packetToAck.getSequenceNumber() + 1);
39         this.withSequenceNumber(packetToAck.getSequenceNumber());
40         this.withIndex(packetToAck.getIndex());
41         return new Packet(this.destination, this.origin, this.flags, this.payload,
42             this.sequenceNumber, this.acknowledgmentNumber, this.index);
43     }
44
45     public PacketBuilder withFlags(Flag... flags) {
46         for (Flag flag : flags) {
47             if (this.flags.contains(flag)) continue;
48             this.flags.add(flag);
49         }
50         return this;
51     }
52
53     public PacketBuilder withPayload(Payload payload) {
54         this.payload = payload;
55         return this;
56     }
57
58     public PacketBuilder withOrigin(Endpoint self) {
59         this.origin = self;
60         return this;
61     }
62 }

```

```

59     public PacketBuilder withDestination(Endpoint destination) {
60         this.destination = destination;
61         return this;
62     }
63
64     public PacketBuilder withSequenceNumber(int sequenceNumber) {
65         this.sequenceNumber = sequenceNumber;
66         return this;
67     }
68
69     public PacketBuilder withAcknowledgmentNumber(int acknowledgmentNumber) {
70         this.acknowledgmentNumber = acknowledgmentNumber;
71         return this;
72     }
73
74     public PacketBuilder withConnection(Connection connection) {
75         this.withOrigin(connection.getConnectionSource());
76         this.withDestination(connection.getConnectionedNode());
77         return this;
78     }
79
80     public PacketBuilder withIndex(int index) {
81         this.index = index;
82         return this;
83     }
84
85
86     public boolean hasFlag(Flag... flags) {
87         var hasFlag = true;
88         for (Flag flag : flags) {
89             hasFlag &= this.flags.contains(flag);
90         }
91         return hasFlag;
92     }
93
94 }

```

```

1 package org.example.data;
2
3 public interface Payload {
4
5     /**
6      * A method that returns the size of the packet.
7      * The size of an object is abstract and must be implemented according to what
8      * type
9      * Payload is created.
10     * @return the size of the packet
11     */
12     int size();
13
14     /**
15      * A method that returns a String representation of the Payload
16      *
17      * @return String representation.
18      */
19     String toString();
20
21     /**
22      * A method to check if two object are equal or not
23      *
24      * @param obj to be checked against
25      * @return true if this is considered equal to given object
26      */
27     boolean equals(Object obj);
28
29     /**
30      * A method that creates a hash number of this object
31      *
32      * @return Integer hash of the object
33      */
34     int hashCode();
35
36 }

```

```

1 package org.example.network;
2
3 import org.example.data.Packet;
4 import org.example.network.interfaces.NetworkNode;
5 import org.example.protocol.MPTCP;
6 import org.example.util.Util;
7
8 import java.util.Queue;
9 import java.util.concurrent.ArrayBlockingQueue;
10
11 public class Channel implements Comparable<Channel> {
12
13     private static final double BAD_TO_GOOD_PROB = 0.2;
14     private static final int CAPACITY = 1000;
15     private final Queue<Packet> line;
16     private final NetworkNode source;
17     private final NetworkNode destination;
18     private final int cost;
19     private final double loss;
20     private boolean goodState;
21
22
23     public Channel(NetworkNode source, NetworkNode destination, double loss, int
24         cost) {
25         this.line = new ArrayBlockingQueue<>(CAPACITY);
26         this.source = source;
27         this.destination = destination;
28         this.loss = loss;
29         this.cost = cost;
30         this.goodState = true;
31     }
32
33     public Channel(NetworkNode source, NetworkNode destination, double loss) {
34         this(source, destination, loss, Util.getNextRandomInt(100));
35     }
36
37     //loopback
38     public Channel(NetworkNode source) {
39         this(source, source, 0, 0);
40     }
41
42     public long propagationDelay() {
43         return 10 * (this.cost + 1L);
44     }
45
46     private boolean lossy() {
47         if (this.source.equals(this.destination)) return false;
48         if (this.goodState) {
49             this.goodState = Util.getNextRandomDouble() >= this.loss;
50             return Util.getNextRandomDouble() < this.loss;
51         }
52         this.goodState = Util.getNextRandomDouble() >= BAD_TO_GOOD_PROB;
53         return Util.getNextRandomDouble() < BAD_TO_GOOD_PROB;
54     }
55
56     public void channelPackage(Packet packet) {
57         if (!this.line.offer(packet)) {
58             throw new IllegalStateException("packet dropped in channel");
59         }
60     }
61
62     public NetworkNode getSource() {

```

```

62     return source;
63 }
64
65 public NetworkNode getDestination() {
66     return destination;
67 }
68
69 public int getCost() {
70     return cost;
71 }
72
73 public boolean channel() {
74     if (this.line.isEmpty()) {
75         return false;
76     }
77
78     var packet = this.line.poll();
79     if (lossy()) {
80         return false;
81     }
82     return this.destination.enqueueInputBuffer(packet);
83
84 }
85
86 @Override
87 public String toString() {
88     return this.source.toString() + " -> [" + this.cost + "] -> " + this.
89         destination.toString();
90
91 }
92
93 @Override
94 public int compareTo(Channel channel) {
95     if (this.cost > channel.getCost()) return 1;
96     if (this.cost < channel.getCost()) return -1;
97     return 0;
98 }
99
100 @Override
101 public int hashCode() {
102     return super.hashCode();
103 }
104
105 @Override
106 public boolean equals(Object obj) {
107     return super.equals(obj);
108 }
109
110 public static class ChannelBuilder {
111
112     double loss = 0;
113     int cost = Util.getNextRandomInt(100);
114
115     public ChannelBuilder withCost(int cost) {
116         if (cost < 0 || cost > 100) throw new IllegalStateException("cost is not
117             valid");
118         this.cost = cost;
119         return this;
120     }
121
122     public ChannelBuilder withLoss(double loss) {
123         this.loss = loss;
124         return this;
125     }
126 }

```



```
123
124     public void build(NetworkNode node1, NetworkNode node2) {
125         if (node1 instanceof MPTCP) node1 = ((MPTCP) node1).
            getEndpointToAddChannelTo();
126         if (node2 instanceof MPTCP) node2 = ((MPTCP) node2).
            getEndpointToAddChannelTo();
127         node1.addChannel(node2, this.loss, this.cost);
128         node2.addChannel(node1, this.loss, this.cost);
129     }
130
131 }
132
133 }
```

```

1 package org.example.network;
2
3 import org.example.data.Packet;
4 import org.example.network.address.Address;
5 import org.example.network.interfaces.NetworkNode;
6 import org.example.simulator.statistics.Stats;
7
8 import java.util.ArrayList;
9 import java.util.List;
10 import java.util.concurrent.BlockingQueue;
11
12 public abstract class Rutable implements NetworkNode {
13
14     protected final BlockingQueue<Packet> inputBuffer;
15     private final RoutingTable routingTable;
16     private final List<Channel> channels;
17     private final Address address;
18     private List<Channel> channelsUsed;
19
20     protected Rutable(BlockingQueue<Packet> inputBuffer, Address address) {
21         this.inputBuffer = inputBuffer;
22         this.channels = new ArrayList<>();
23         this.address = address;
24         this.routingTable = new RoutingTable();
25
26         this.channelsUsed = new ArrayList<>(1);
27     }
28
29     public abstract Stats getStats();
30
31     @Override
32     public void updateRoutingTable() {
33         this.routingTable.update(this);
34     }
35
36     @Override
37     public void route(Packet packet) {
38         if (packet == null) throw new IllegalStateException("Null packet can't be
39             routed");
40         NetworkNode destination = packet.getDestination();
41         var nextChannelOnPath = this.routingTable.getPath(this, destination);
42         nextChannelOnPath.channelPackage(packet);
43
44         this.channelsUsed.add(nextChannelOnPath);
45     }
46
47     @Override
48     public List<Channel> getChannelsUsed() {
49         List<Channel> used = this.channelsUsed;
50         this.channelsUsed = new ArrayList<>(1);
51         return used;
52     }
53
54     @Override
55     public long delay() {
56         return ((long) this.inputBufferSize()) * 10;
57     }
58
59     @Override
60     public List<Channel> getChannels() {
61         return this.channels;
62     }

```

```

62
63
64     @Override
65     public void addChannel(NetworkNode node, double noiseTolerance, int cost) {
66         var channel = new Channel(this, node, noiseTolerance, cost);
67         this.channels.add(channel);
68     }
69
70     @Override
71     public Address getAddress() {
72         return this.address;
73     }
74
75     @Override
76     public boolean enqueueInputBuffer(Packet packet) {
77         return this.inputBuffer.offer(packet);
78     }
79
80     @Override
81     public Packet peekInputBuffer() {
82         return this.inputBuffer.peek();
83     }
84
85     @Override
86     public Packet dequeueInputBuffer() {
87         return this.inputBuffer.poll();
88     }
89
90     @Override
91     public boolean inputBufferIsEmpty() {
92         return this.inputBuffer.isEmpty();
93     }
94
95     @Override
96     public int inputBufferSize() {
97         return this.inputBuffer.size();
98     }
99
100    @Override
101    public boolean equals(Object obj) {
102        if (obj instanceof NetworkNode) {
103            NetworkNode node = (NetworkNode) obj;
104            return this.getAddress().equals(node.getAddress());
105        }
106        return false;
107    }
108
109    @Override
110    public int hashCode() {
111        return this.getAddress().hashCode();
112    }
113
114    @Override
115    public String toString() {
116        return this.address.toString();
117    }
118
119
120 }

```

```

1 package org.example.network;
2
3 import org.example.data.Packet;
4 import org.example.network.address.UUIDAddress;
5 import org.example.network.interfaces.Endpoint;
6 import org.example.simulator.statistics.Stats;
7
8 import java.util.concurrent.ArrayBlockingQueue;
9 import java.util.concurrent.BlockingQueue;
10
11 public class RoutableEndpoint extends Routable implements Endpoint {
12
13     private final BlockingQueue<Packet> receivedPackets;
14     protected BlockingQueue<Packet> outputBuffer;
15
16     public RoutableEndpoint(BlockingQueue<Packet> inputBuffer, BlockingQueue<Packet>
17         outputBuffer) {
18         super(inputBuffer, new UUIDAddress());
19         this.outputBuffer = outputBuffer;
20         this.receivedPackets = new ArrayBlockingQueue<>(10000);
21     }
22
23     @Override
24     public Stats getStats() {
25         return null;
26     }
27
28     public Packet dequeueOutputBuffer() {
29         return this.outputBuffer.poll();
30     }
31
32     public boolean enqueueOutputBuffer(Packet packet) {
33         return this.outputBuffer.offer(packet);
34     }
35
36     @Override
37     public boolean outputBufferIsEmpty() {
38         return this.outputBuffer.isEmpty();
39     }
40
41     public int outputBufferSize() {
42         return this.outputBuffer.size();
43     }
44
45     @Override
46     public boolean isConnected() {
47         return false;
48     }
49
50     public Packet getReceivedPacket() {
51         return this.receivedPackets.poll();
52     }
53
54     @Override
55     public void run() {
56         if (this.inputBufferIsEmpty()) return;
57         Packet received = this.dequeueInputBuffer();
58         this.receivedPackets.add(received);
59     }
60
61     @Override

```

```
62     public String toString() {
63         return "Endpoint: " + super.toString();
64     }
65
66     @Override
67     public boolean equals(Object o) {
68         return super.equals(o);
69     }
70
71     @Override
72     public int hashCode() {
73         return super.hashCode();
74     }
75
76 }
```

```

1 package org.example.network;
2
3 import org.apache.commons.math3.distribution.PoissonDistribution;
4 import org.example.network.address.Address;
5 import org.example.network.address.UUIDAddress;
6 import org.example.simulator.statistics.Stats;
7 import org.example.util.Util;
8
9 import java.util.concurrent.ArrayBlockingQueue;
10
11 public class Router extends Routable {
12
13     private final double queueSizeMean;
14     private final int bufferSize;
15     private final PoissonDistribution poissonDistribution;
16     private int artificialQueueSize;
17
18     private Router(int bufferSize, Address address, double averageQueueUtilization)
19         {
20         super(new ArrayBlockingQueue<>(bufferSize), address);
21         this.bufferSize = bufferSize;
22         this.queueSizeMean = averageQueueUtilization * this.bufferSize;
23         this.poissonDistribution = Util.getPoissonDistribution(this.queueSizeMean);
24         this.setArtificialQueueSize();
25     }
26
27     @Override
28     public Stats getStats() {
29         return null;
30     }
31
32     private void setArtificialQueueSize() {
33         int queueSize = this.poissonDistribution.sample();
34         this.artificialQueueSize = queueSize;
35     }
36
37     @Override
38     public long delay() {
39         this.setArtificialQueueSize();
40         long transmissionDelay = 10;
41         return transmissionDelay + this.artificialQueueSize * transmissionDelay;
42     }
43
44     @Override
45     public void run() {
46         if (!this.inputBufferIsEmpty()) {
47             var packet = this.dequeueInputBuffer();
48             this.route(packet);
49         }
50     }
51
52     @Override
53     public boolean equals(Object obj) {
54         return super.equals(obj);
55     }
56
57     @Override
58     public int hashCode() {
59         return super.hashCode();
60     }
61
62     public static class RouterBuilder {

```

```
62
63     private int bufferSize = 100;
64     private double averageQueueUtilization = 0.85;
65     private Address address = new UUIDAddress();
66
67     public RouterBuilder withAverageQueueUtilization(double
68         averageQueueUtilization) {
69         this.averageQueueUtilization = averageQueueUtilization;
70         return this;
71     }
72
73     public RouterBuilder withAddress(Address address) {
74         this.address = address;
75         return this;
76     }
77
78     public Router build() {
79         return new Router(this.bufferSize, this.address, this.
80             averageQueueUtilization);
81     }
82 }
83
84 }
```

```

1 package org.example.network;
2
3 import org.example.network.interfaces.NetworkNode;
4
5 import java.util.*;
6
7 public class RoutingTable {
8
9     private Map<NetworkNode, Map.Entry<Channel, Integer>> table;
10
11
12     public RoutingTable() {
13         this.table = new HashMap<>();
14     }
15
16
17     private void updateTable(NetworkNode node, Channel channel) {
18         if (this.table.containsKey(node)) {
19             this.updateEntry(node, channel);
20             return;
21         }
22         this.addEntry(node, channel);
23     }
24
25     private void addEntry(NetworkNode node, Channel channel) {
26         if (channel == null) {
27             this.table.put(node, Map.entry(new Channel(node), 0));
28             return;
29         }
30
31         boolean prevNodeInTable = table.containsKey(channel.getSource());
32         int newCost = prevNodeInTable ? table.get(channel.getSource()).getValue() +
33             channel.getCost() : channel.getCost();
34
35         table.put(node, Map.entry(channel, newCost));
36     }
37
38     private void updateEntry(NetworkNode node, Channel channel) {
39         boolean prevNodeInTable = table.containsKey(channel.getSource());
40         int newCost = prevNodeInTable ? table.get(channel.getSource()).getValue() +
41             channel.getCost() : 0;
42         if (newCost < this.table.get(node).getValue()) {
43             table.replace(node, Map.entry(channel, newCost));
44         }
45     }
46
47     public void update(NetworkNode startingNode) {
48         this.addEntry(startingNode, null);
49         update(startingNode, new ArrayList<>(), new PriorityQueue<>());
50     }
51
52     private void update(NetworkNode curNode, List<NetworkNode> visited, Queue<
53         Channel> priorityQueue) {
54         visited.add(curNode);
55
56         for (Channel channel : curNode.getChannels()) {
57             priorityQueue.offer(channel);
58             this.updateTable(channel.getDestination(), channel);
59         }
60
61         while (!priorityQueue.isEmpty()) {

```



```

60         var bestChannel = priorityQueue.poll();
61         if (visited.contains(bestChannel.getDestination())) continue;
62         update(bestChannel.getDestination(), visited, priorityQueue);
63     }
64 }
65
66
67 public Channel getPath(NetworkNode source, NetworkNode destination) {
68     if (this.table.isEmpty()) throw new IllegalStateException("The routing table
69         is empty");
70     if (destination == null) throw new IllegalArgumentException("The destination
71         can't be null");
72     if (!this.table.containsKey(destination))
73         throw new IllegalArgumentException("Destination " + destination + " does
74             not exist in this routing table");
75
76     var curChannel = this.table.get(destination).getKey();
77     NetworkNode prevNode = curChannel.getSource();
78     if (source.equals(prevNode)) return curChannel;
79     return getPath(source, prevNode);
80 }
81
82 @Override
83 public String toString() {
84     var builder = new StringBuilder();
85     for (Map.Entry<NetworkNode, Map.Entry<Channel, Integer>> entry : this.table.
86         entrySet()) {
87         var networkNode = entry.getKey();
88         Map.Entry<Channel, Integer> value = entry.getValue();
89         var channel = value.getKey();
90         int cost = value.getValue();
91
92         builder.append("Node: ");
93         builder.append(networkNode);
94         builder.append(" | Channel: ");
95         builder.append(channel);
96         builder.append(" | Cost: ");
97         builder.append(cost);
98         builder.append("\n");
99     }
100     return builder.toString();
101 }

```

```

1 package org.example.network.address;
2
3 import java.util.Objects;
4
5 public abstract class Address {
6
7     public abstract String getId();
8
9     @Override
10    public boolean equals(Object o) {
11        if (this == o) return true;
12        if (o == null || getClass() != o.getClass()) return false;
13        var address = (Address) o;
14        return Objects.equals(this.getId(), address.getId());
15    }
16
17    @Override
18    public int hashCode() {
19        return Objects.hash(this.getId());
20    }
21
22    @Override
23    public String toString() {
24        return this.getId();
25    }
26
27 }

```

```

1 package org.example.network.address;
2
3 public class SimpleAddress extends Address {
4
5     private final String id;
6
7     public SimpleAddress(String id) {
8         this.id = id;
9     }
10
11    @Override
12    public String getId() {
13        return this.id;
14    }
15
16    @Override
17    public boolean equals(Object o) {
18        return super.equals(o);
19    }
20
21    @Override
22    public int hashCode() {
23        return super.hashCode();
24    }
25 }

```

```

1 package org.example.network.address;
2
3 import java.util.UUID;
4
5 public class UUIDAddress extends Address {
6
7     private final UUID id;
8
9     public UUIDAddress() {
10         this.id = UUID.randomUUID();
11     }
12
13     @Override
14     public String getId() {
15         return this.id.toString();
16     }
17
18     @Override
19     public boolean equals(Object o) {
20         return super.equals(o);
21     }
22
23     @Override
24     public int hashCode() {
25         return super.hashCode();
26     }
27
28 }

```

```

1 package org.example.network.interfaces;
2
3
4 public interface Endpoint extends NetworkNode {
5     /**
6      * A method that checks if the outputBuffer has any Packets
7      *
8      * @return True if the outputBuffer is empty
9      */
10    boolean outputBufferIsEmpty();
11
12
13    /**
14     * A method that checks is the endpoint is connected
15     *
16     * @return True if connected
17     */
18    boolean isConnected();
19
20
21 }

```

```

1 package org.example.network.interfaces;
2
3 import org.example.data.Packet;
4 import org.example.network.Channel;
5 import org.example.network.address.Address;
6
7 import java.util.List;
8
9 public interface NetworkNode {
10
11
12     /**
13      * A method that updates the routing table according to the cost
14      */
15     void updateRoutingTable();
16
17     /**
18      * A method that routes the packet to the next router on the path to it's
19      * destination
20      * @param packet to route
21      */
22     void route(Packet packet);
23
24
25     /**
26      * A method that returns the NetworkNode's associated delay
27      * @return delay
28      */
29     long delay();
30
31
32
33     /**
34      * A method that returns a List of the NetworkNode's outgoing Channels
35      * @return List of outgoing Channels from this NetworkNode
36      */
37     List<Channel> getChannels();
38
39
40
41     /**
42      * * A method that creates and adds a Channel to the list of channels.
43      * @param node NetworkNode to create channel to
44      * @param noiseTolerance
45      * @param cost the cost of the channel used in routing
46      */
47     void addChannel(NetworkNode node, double noiseTolerance, int cost);
48
49
50
51     /**
52      * A method that returns the unique Address associated with this NetworkNode
53      * @return unique Address
54      */
55     Address getAddress();
56
57
58
59     /**
60      * A method that returns the first Packet form the inputBuffer without removing
61      * it from the buffer

```

```

61     *
62     * @return the dequeued Packet
63     */
64     Packet peekInputBuffer();
65
66
67     /**
68     * A method that returns and dequeues a Packet form the inputBuffer
69     *
70     * @return the dequeued Packet
71     */
72     Packet dequeueInputBuffer();
73
74     /**
75     * A method that enqueues the given Packet to the inputBuffer
76     *
77     * @param packet to be enqueued
78     * @return True if successful
79     */
80     boolean enqueueInputBuffer(Packet packet);
81
82     /**
83     * A method that checks if the inputBuffer has any Packets
84     *
85     * @return True if the inputBuffer is empty
86     */
87     boolean inputBufferIsEmpty();
88
89     /**
90     * A method that returns the size of the inputBuffer
91     *
92     * @return size of the inputBuffer
93     */
94     int inputBufferSize();
95
96     /**
97     * A method that returns a list of all channels used during an event
98     *
99     * @return A List with Channel objects
100    */
101    List<Channel> getChannelsUsed();
102
103    /**
104    * A method to run the NetworkNode
105    */
106    void run();
107
108
109 }

```

```

1 package org.example.protocol;
2
3 import org.example.data.*;
4 import org.example.network.Channel;
5 import org.example.network.Routable;
6 import org.example.network.address.Address;
7 import org.example.network.address.UUIDAddress;
8 import org.example.network.interfaces.Endpoint;
9 import org.example.protocol.window.receiving.ReceivingWindow;
10 import org.example.protocol.window.receiving.SelectiveRepeat;
11 import org.example.protocol.window.sending.SendingWindow;
12 import org.example.protocol.window.sending.SlidingWindow;
13 import org.example.simulator.statistics.TCPStats;
14 import org.example.util.Util;
15 import org.javatuples.Pair;
16
17 import java.util.ArrayList;
18 import java.util.Comparator;
19 import java.util.List;
20 import java.util.logging.Level;
21 import java.util.logging.Logger;
22
23 public class ClassicTCP extends Routable implements TCP {
24
25     private static final Comparator<Packet> SENDING_WINDOW_PACKET_COMPARATOR =
26         Comparator.comparingInt(Packet::getSequenceNumber);
27     private final TCPStats tcpStats;
28     private final Logger logger = Logger.getLogger(ClassicTCP.class.getSimpleName());
29     ;
30     private final List<Packet> receivedPackets;
31     private final List<Pair<Integer, Payload>> payloadsToSend;
32     private final int thisReceivingWindowCapacity;
33     private final boolean isReno;
34     private final TCP mainFlow;
35     private int otherReceivingWindowCapacity;
36     private int initialSequenceNumber;
37     private long rtt;
38     private long rttStartTimer = 0;
39     private boolean timerStarted = false;
40     private boolean rttSet = false;
41     private SendingWindow sendingWindow;
42
43     private ClassicTCP(int thisReceivingWindowCapacity,
44         List<Packet> receivedPackets,
45         List<Pair<Integer, Payload>> payloadsToSend,
46         boolean isReno,
47         Address address,
48         TCP mainFlow,
49         ReceivingWindow receivingWindow) {
50         super(receivingWindow, address);
51         this.receivedPackets = receivedPackets;
52         this.payloadsToSend = payloadsToSend;
53         this.thisReceivingWindowCapacity = thisReceivingWindowCapacity;
54         this.isReno = isReno;
55
56         if (mainFlow == null) {
57             this.mainFlow = this;
58         } else {
59             this.mainFlow = mainFlow;
60         }
61
62         this.tcpStats = new TCPStats(address);

```

```

61     }
62
63
64     @Override
65     public void connect(TCP host) {
66         if (this.isConnected()) return;
67         this.initialSequenceNumber = Util.getNextRandomInt(10000);
68         Packet syn = new PacketBuilder()
69             .withDestination(host)
70             .withOrigin(this)
71             .withFlags(Flag.SYN)
72             .withSequenceNumber(this.initialSequenceNumber)
73             .withPayload(new Message(this.thisReceivingWindowCapacity + ""))
74             .build();
75         this.route(syn);
76         if (!this.timerStarted) {
77             this.rttStartTimer = Util.getTime();
78             this.timerStarted = true;
79         }
80     }
81
82     public void continueConnect(Packet synAck) {
83         Endpoint host = synAck.getOrigin();
84         if (synAck.hasAllFlags(Flag.SYN, Flag.ACK)) {
85             if (synAck.getAcknowledgmentNumber() != this.initialSequenceNumber + 1)
86                 return;
87             this.otherReceivingWindowCapacity = Integer.parseInt(synAck.getPayload()
88                 .toString());
89             int finalSeqNum = synAck.getAcknowledgmentNumber();
90             int ackNum = synAck.getSequenceNumber() + 1;
91             Packet ack = new PacketBuilder()
92                 .withDestination(host)
93                 .withOrigin(this)
94                 .withFlags(Flag.ACK)
95                 .withSequenceNumber(finalSeqNum)
96                 .withAcknowledgmentNumber(ackNum)
97                 .build();
98             this.route(ack);
99
100             this.setRTT();
101             this.setConnection(new Connection(this, host, finalSeqNum, ackNum));
102
103             this.logger.log(Level.INFO, () -> "connection established with host: " +
104                 this.getConnection());
105         }
106     }
107
108     @Override
109     public void connect(Packet syn) {
110         if (this.isConnected()) return;
111         Endpoint node = syn.getOrigin();
112         var seqNum = Util.getNextRandomInt(10000);
113         var ackNum = syn.getSequenceNumber() + 1;
114         this.otherReceivingWindowCapacity = Integer.parseInt(syn.getPayload().
115             toString());
116         Packet synAck = new PacketBuilder()
117             .withDestination(node)
118             .withOrigin(this)
119             .withFlags(Flag.SYN, Flag.ACK)
120             .withSequenceNumber(seqNum)

```

```

120         .withAcknowledgmentNumber(ackNum)
121         .withPayload(new Message(this.thisReceivingWindowCapacity + ""))
122         .build();
123     this.route(synAck);
124
125     if (!this.timerStarted) {
126         this.rttStartTimer = Util.seeTime();
127         this.timerStarted = true;
128     }
129
130     this.setConnection(new Connection(this, node, seqNum, ackNum));
131
132     this.logger.log(Level.INFO, () -> "connection established with client: " +
133         this.getConnection());
134 }
135
136 private void setRTT() {
137     if (rttSet) return;
138     this.rtt = Util.seeTime() - this.rttStartTimer;
139     this.rttSet = true;
140 }
141
142 @Override
143 public void send(Payload payload) {
144     if (this.payloadsToSend.isEmpty()) {
145         this.payloadsToSend.add(Pair.with(0, payload));
146         return;
147     }
148     int indexOfLastAdded = this.payloadsToSend.get(this.payloadsToSend.size() -
149         1).getValue0();
150     this.payloadsToSend.add(Pair.with(indexOfLastAdded + 1, payload));
151 }
152
153 @Override
154 public Packet receive() {
155     if (this.receivedPackets.isEmpty()) return null;
156     return this.receivedPackets.remove(0);
157 }
158
159 @Override
160 public Channel getChannel() {
161     return this.getChannels().get(0);
162 }
163
164 public int getThisReceivingWindowCapacity() {
165     return this.thisReceivingWindowCapacity;
166 }
167
168 public int getOtherReceivingWindowCapacity() {
169     return this.otherReceivingWindowCapacity;
170 }
171
172 @Override
173 public boolean isConnected() {
174     try {
175         return this.getSendingWindow().getConnection() != null;
176     } catch (IllegalAccessException e) {
177         //no SendingWindow means no connection
178         return false;
179     }
180 }

```



```

181
182 public Connection getConnection() {
183     try {
184         return this.getSendingWindow().getConnection();
185     } catch (IllegalAccessException e) {
186         //No SendingWindow means no connection is established
187         return null;
188     }
189 }
190
191 private void setConnection(Connection connection) {
192     this.sendingWindow = new SlidingWindow(this.otherReceivingWindowCapacity,
193         this.isReno, connection, SENDING_WINDOW_PACKET_COMPARATOR, this.
194         payloadsToSend, this.tcpStats);
195 }
196
197 @Override
198 public long getRTO() {
199     return 3 * this.rtt;
200 }
201
202 @Override
203 public Packet fastRetransmit() {
204     try {
205         return this.getSendingWindow().fastRetransmit();
206     } catch (IllegalAccessException e) {
207         //no Sending Window results in no retransmit
208         return null;
209     }
210 }
211
212 @Override
213 public long delay() {
214     return 10;
215 }
216
217 @Override
218 public TCP getMainFlow() {
219     return this.mainFlow;
220 }
221
222 @Override
223 public int getNumberOfFlows() {
224     return 1;
225 }
226
227 public SendingWindow getSendingWindow() throws IllegalAccessException {
228     if (this.sendingWindow != null) return this.sendingWindow;
229     throw new IllegalAccessException("SendingWindow is null");
230 }
231
232 public ReceivingWindow getReceivingWindow() throws IllegalAccessException {
233     if (this.inputBuffer instanceof ReceivingWindow) return (ReceivingWindow)
234         this.inputBuffer;
235     throw new IllegalAccessException("The outputBuffer is not of type
236         ReceivingWindow");
237 }
238
239 private boolean packetIsFromValidConnection(Packet packet) {
240     if (packet.hasAllFlags(Flag.SYN) && !this.isConnected()) return true;
241     if (!this.isConnected()) return false;

```

```

240     var conn = this.getConnection();
241     if (packet.hasAllFlags(Flag.ACK)) return true;
242
243
244     try {
245         boolean inWindow = this.getReceivingWindow().inReceivingWindow(packet,
246             conn);
247         if (!inWindow) {
248             this.ack(this.getReceivingWindow().ackThis(this.getSendingWindow().
249                 getConnection().getConnectedNode()));
250             return false;
251         }
252         return packet.getOrigin().equals(conn.getConnectedNode())
253             && packet.getDestination().equals(conn.getConnectionSource());
254     } catch (IllegalAccessException e) {
255         return false;
256     }
257
258     @Override
259     public boolean outputBufferIsEmpty() {
260         return this.sendingWindow.isEmpty();
261     }
262
263     @Override
264     public boolean enqueueInputBuffer(Packet packet) {
265         if (packet.isValidConnection(packet)) {
266             if (super.enqueueInputBuffer(packet)) {
267                 if (!packet.hasAllFlags(Flag.ACK) && !packet.hasAllFlags(Flag.SYN))
268                     this.tcpStats.packetArrival(packet);
269                 return true;
270             }
271             return true;
272         }
273         // return true because the packet has arrived the endpoint
274         // the packet is not added to the input buffer, but it is checked
275         return false;
276     }
277
278     private boolean unconnectedInputHandler() {
279         if (this.inputBuffer.isEmpty()) return false;
280         if (!this.peekInputBuffer().getDestination().equals(this)) return false;
281
282         var packet = this.dequeueInputBuffer();
283
284         if (packet.hasAllFlags(Flag.SYN, Flag.ACK)) {
285             this.continueConnect(packet);
286             return true;
287         }
288
289         if (packet.hasAllFlags(Flag.SYN)) {
290             this.connect(packet);
291             return true;
292         }
293         return false;
294     }
295
296     public boolean canRetransmit(Packet packet) {
297         try {
298             return this.getSendingWindow().canRetransmit(packet);
299         } catch (IllegalAccessException e) {

```

```

300         //should not retransmit without SendingWindow
301         return false;
302     }
303 }
304
305
306 private void ack(Packet packet) {
307     assert packet != null : "Packet is null";
308
309     if (packet.getSequenceNumber() == -1) return;
310
311     if (packet.getOrigin() == null) {
312         //can't call ack on packet with no origin
313         Endpoint connectedNode;
314         try {
315             connectedNode = this.getSendingWindow().getConnection().
316                 getConnectedNode();
317         } catch (IllegalAccessException e) {
318             //should not be able to ack without a SendingWindow
319             return;
320         }
321         packet = new PacketBuilder()
322             .withDestination(connectedNode)
323             .withOrigin(this)
324             .withSequenceNumber(packet.getSequenceNumber())
325             .withAcknowledgmentNumber(packet.getAcknowledgmentNumber())
326             .withPayload(packet.getPayload())
327             .build();
328     }
329     Packet ack = new PacketBuilder().ackBuild(packet);
330     this.route(ack);
331 }
332
333
334 public boolean handleIncoming() {
335     if (!isConnected()) return unconnectedInputHandler();
336
337     try {
338         var receivingWindow = this.getReceivingWindow();
339         var packetReceived = receivingWindow.receive(this.getSendingWindow());
340         if (packetReceived && receivingWindow.shouldAck()) {
341             try {
342                 this.ack(receivingWindow.ackThis(this.getSendingWindow().
343                     getConnection().getConnectedNode()));
344             } catch (IllegalArgumentException e) {
345                 return false;
346             }
347         }
348         this.setRTT();
349
350         var packetToFastRetransmit = this.fastRetransmit();
351         if (packetToFastRetransmit != null) {
352             this.route(packetToFastRetransmit);
353             this.tcpStats.packetFastRetransmit();
354         }
355         return true;
356     } catch (IllegalAccessException e) {
357         throw new IllegalStateException("TCP is connected but no ReceivingWindow
358             or SendingWindow is established");
359     }

```

```

360
361
362 @Override
363 public List<Packet> trySend() {
364     return this.trySend(new ArrayList<>());
365 }
366
367 private List<Packet> trySend(List<Packet> packetsSent) {
368     if (this.sendingWindow == null) return packetsSent;
369     if (this.sendingWindow.isQueueEmpty()) return packetsSent;
370     if (this.sendingWindow.isWaitingForAck()) return packetsSent;
371
372     var packetToSend = this.sendingWindow.send();
373     assert packetToSend != null;
374
375     this.route(packetToSend);
376     this.tcpStats.packetSend();
377
378     packetsSent.add(packetToSend);
379     return trySend(packetsSent);
380 }
381
382 @Override
383 public TCPStats getStats() {
384     return tcpStats;
385 }
386
387 @Override
388 public void run() {
389     this.handleIncoming();
390     this.trySend();
391 }
392
393 @Override
394 public boolean equals(Object o) {
395     return super.equals(o);
396 }
397
398 @Override
399 public int hashCode() {
400     return super.hashCode();
401 }
402
403
404 public static class ClassicTCPBuilder {
405
406     private int receivingWindowCapacity = 7;
407     private List<Packet> receivedPacketsContainer = new ArrayList<>();
408     private List<Pair<Integer, Payload>> payloadsToSend = new ArrayList<>();
409     private boolean isReno = true;
410     private Address address = new UUIAddress();
411     private TCP mainflow = null;
412     private ReceivingWindow receivingWindow = new SelectiveRepeat(this.
413         receivingWindowCapacity, Comparator.comparingInt(Packet::
414             getSequenceNumber), this.receivedPacketsContainer);
415
416     public ClassicTCPBuilder withReceivingWindowCapacity(int
417         receivingWindowCapacity) {
418         this.receivingWindowCapacity = receivingWindowCapacity;
419         return this;
420     }
421
422     public ClassicTCPBuilder withReceivedPacketsContainer(List<Packet>

```

```

420         receivedPacketsContainer) {
421             this.receivedPacketsContainer = receivedPacketsContainer;
422             return this;
423         }
424     public ClassicTCPBuilder withPayloadsToSend(List<Pair<Integer, Payload>>
425         payloadsToSend) {
426         this.payloadsToSend = payloadsToSend;
427         return this;
428     }
429     public ClassicTCPBuilder setTahoe() {
430         this.isReno = false;
431         return this;
432     }
433     public ClassicTCPBuilder setReno() {
434         this.isReno = true;
435         return this;
436     }
437     }
438     public ClassicTCPBuilder withAddress(Address address) {
439         this.address = address;
440         return this;
441     }
442     }
443     public ClassicTCPBuilder withMainFlow(TCP tcp) {
444         this.mainflow = tcp;
445         return this;
446     }
447     }
448     public ClassicTCPBuilder withReceivingWindow(ReceivingWindow receivingWindow
449         ) {
450         this.receivingWindow = receivingWindow;
451         return this;
452     }
453     }
454     public ClassicTCP build() {
455         return new ClassicTCP(this.receivingWindowCapacity,
456             this.receivedPacketsContainer,
457             this.payloadsToSend,
458             this.isReno,
459             this.address,
460             this.mainflow,
461             this.receivingWindow);
462     }
463 }
464
465 }

```

```

1 package org.example.protocol;
2
3 import org.example.data.Flag;
4 import org.example.data.Packet;
5 import org.example.network.interfaces.Endpoint;
6
7 import java.util.Objects;
8
9 public class Connection {
10
11     private Endpoint self;
12     private Endpoint other;
13     private int sequenceNumber;
14     private int acknowledgementNumber;
15
16
17     public Connection(Endpoint self, Endpoint other, int sequenceNumber, int
18         acknowledgementNumber) {
19         this.self = self;
20         this.other = other;
21         this.sequenceNumber = sequenceNumber;
22         this.acknowledgementNumber = acknowledgementNumber;
23     }
24
25     public void update(Packet packet) {
26         if (packet.hasAllFlags(Flag.ACK)) {
27             this.sequenceNumber = packet.getAcknowledgmentNumber();
28             return;
29         }
30         this.acknowledgementNumber = packet.getSequenceNumber() + 1;
31     }
32
33     public int getNextSequenceNumber() {
34         return sequenceNumber;
35     }
36
37     public int getNextAcknowledgementNumber() {
38         return acknowledgementNumber;
39     }
40
41     public Endpoint getConnectedNode() {
42         return other;
43     }
44
45     public Endpoint getConnectionSource() {
46         return self;
47     }
48
49     @Override
50     public boolean equals(Object o) {
51         if (this == o) return true;
52         if (o == null || getClass() != o.getClass()) return false;
53         Connection that = (Connection) o;
54         return Objects.equals(self, that.self) &&
55             Objects.equals(other, that.other);
56     }
57
58     @Override
59     public int hashCode() {
60         return Objects.hash(self, other);
61     }

```

```
62 |
63 |     @Override
64 |     public String toString() {
65 |         return "Connection[" + other + "];
66 |     }
67 | }
```

```

1 package org.example.protocol;
2
3 import org.example.data.Packet;
4 import org.example.data.Payload;
5 import org.example.network.Channel;
6 import org.example.network.address.Address;
7 import org.example.network.address.SimpleAddress;
8 import org.example.network.address.UUIDAddress;
9 import org.example.network.interfaces.Endpoint;
10 import org.example.network.interfaces.NetworkNode;
11 import org.example.protocol.window.receiving.ReceivingWindow;
12 import org.example.protocol.window.receiving.SelectiveRepeat;
13 import org.example.simulator.statistics.TCPStats;
14 import org.javatuples.Pair;
15
16 import java.util.ArrayList;
17 import java.util.Comparator;
18 import java.util.List;
19 import java.util.logging.Level;
20 import java.util.logging.Logger;
21
22 public class MPTCP implements TCP {
23
24     private static final String DEPRECATED_STRING = "DEPRECATED";
25     private static final Comparator<Packet> PACKET_INDEX_COMPARATOR = Comparator.
        comparingInt(Packet::getIndex);
26     private final List<Packet> receivedPackets;
27     private final List<Pair<Integer, Payload>> payloadsToSend;
28     private final TCP[] subflows;
29     private final Address address;
30     private final ReceivingWindow receivingWindow;
31
32     private MPTCP(int numberOfSubflows, int receivingWindowCapacity, Address address
33         ) {
34         this.receivedPackets = new ArrayList<>();
35         this.payloadsToSend = new ArrayList<>();
36         this.subflows = new TCP[numberOfSubflows];
37         this.address = address;
38         this.receivingWindow = new SelectiveRepeat(receivingWindowCapacity,
39             PACKET_INDEX_COMPARATOR, this.receivedPackets);
40
41         for (var i = 0; i < numberOfSubflows; i++) {
42             TCP tcp = new ClassicTCP.ClassicTCPBuilder()
43                 .withReceivingWindowCapacity(receivingWindowCapacity /
44                     numberOfSubflows)
45                 .withAddress(new SimpleAddress("Subflow " + i + " " + this.
46                     address))
47                 .withReceivedPacketsContainer(this.receivedPackets)
48                 .withPayloadsToSend(this.payloadsToSend)
49                 .withReceivingWindow(this.receivingWindow)
50                 .withMainFlow(this)
51                 .setReno()
52                 .build();
53             this.subflows[i] = tcp;
54         }
55     }
56
57     @Override
58     public TCP getMainFlow() {
59         return this;
60     }
61 }

```



```

58     @Override
59     public int getNumberOfFlows() {
60         return this.subflows.length;
61     }
62
63     @Override
64     public boolean outputBufferIsEmpty() {
65         return this.payloadsToSend.isEmpty();
66     }
67
68     @Override
69     public void updateRoutingTable() {
70         for (TCP subflow : this.subflows) {
71             subflow.updateRoutingTable();
72         }
73     }
74
75     @Override
76     public void route(Packet packet) {
77         throw new IllegalStateException(DEPRECATED_STRING);
78     }
79
80     @Override
81     public long delay() {
82         return 10;
83     }
84
85     @Override
86     public List<Channel> getChannels() {
87         List<Channel> channels = new ArrayList<>();
88         for (TCP subflow : this.subflows) {
89             for (Channel c : subflow.getChannels()) {
90                 channels.add(c);
91             }
92         }
93         return channels;
94     }
95
96     public Endpoint getEndpointToAddChannelTo() {
97         for (Endpoint endpoint : this.subflows) {
98             if (endpoint.getChannels().isEmpty()) return endpoint;
99         }
100        throw new IllegalStateException("no endpoints available for channel adding");
101    }
102
103
104     @Override
105     public void addChannel(NetworkNode node, double noiseTolerance, int cost) {
106         for (TCP subflow : this.subflows) {
107             if (subflow.getChannels().isEmpty()) { // only allowing one channel per
108                 subflow.addChannel(node, noiseTolerance, cost);
109                 return;
110             }
111         }
112     }
113
114     @Override
115     public Address getAddress() {
116         throw new IllegalStateException(DEPRECATED_STRING);
117     }
118

```

```

119     @Override
120     public Packet peekInputBuffer() {
121         throw new IllegalStateException(DEPRECATED_STRING);
122     }
123
124     @Override
125     public Packet dequeueInputBuffer() {
126         throw new IllegalStateException(DEPRECATED_STRING);
127     }
128
129     @Override
130     public boolean enqueueInputBuffer(Packet packet) {
131         throw new IllegalStateException(DEPRECATED_STRING);
132     }
133
134     @Override
135     public boolean inputBufferIsEmpty() {
136         return this.receivingWindow.isEmpty();
137     }
138
139     @Override
140     public int inputBufferSize() {
141         return this.receivingWindow.size();
142     }
143
144
145     @Override
146     public void run() {
147         throw new IllegalStateException(DEPRECATED_STRING);
148     }
149
150     public TCP[] getSubflows() {
151         return this.subflows;
152     }
153
154     @Override
155     public void connect(TCP host) {
156         if (!(host instanceof MPTCP)) {
157             this.subflows[0].connect(host);
158             return;
159         }
160         var mptcpHost = (MPTCP) host;
161         TCP[] hostSubflows = mptcpHost.getSubflows();
162
163         for (var i = 0; i < this.subflows.length; i++) {
164             TCP cFlow = this.subflows[i];
165             if (cFlow.isConnected()) continue;
166             for (var j = i; j < hostSubflows.length; j++) {
167                 TCP hFlow = hostSubflows[j];
168                 if (hFlow.isConnected()) continue;
169                 try {
170                     cFlow.connect(hFlow);
171                     break;
172                 } catch (IllegalArgumentException e) {
173                     Logger.getLogger(this.getClass().getSimpleName()).log(Level.INFO,
174                         "connection failed");
175                 }
176             }
177         }
178
179     @Override
180     public void connect(Packet syn) {

```

```

181         throw new IllegalStateException(DEPRECATED_STRING);
182     }
183
184     @Override
185     public void send(Payload payload) {
186         // avoiding duplicate code
187         // adds to the payloadsToSend list that is shared among all subflows
188         // therefore, it does not matter which subflow does this
189         this.subflows[0].send(payload);
190     }
191
192     @Override
193     public Packet receive() {
194         if (this.receivedPackets.isEmpty()) return null;
195         return this.receivedPackets.remove(0);
196     }
197
198     @Override
199     public boolean isConnected() {
200         for (TCP subflow : this.subflows) {
201             if (!subflow.isConnected()) return false;
202         }
203         return true;
204     }
205
206     @Override
207     public Channel getChannel() {
208         throw new IllegalStateException(DEPRECATED_STRING);
209     }
210
211     @Override
212     public List<Channel> getChannelsUsed() {
213         List<Channel> usedChannels = new ArrayList<>(this.subflows.length);
214         for (TCP subflow : this.subflows) {
215             for (Channel channel : subflow.getChannelsUsed()) {
216                 usedChannels.add(channel);
217             }
218         }
219         return usedChannels;
220     }
221
222     @Override
223     public long getRTO() {
224         throw new IllegalStateException(DEPRECATED_STRING);
225     }
226
227
228     @Override
229     public boolean handleIncoming() {
230         for (var i = 0; i < this.subflows.length; i++) {
231             for (TCP subflow : this.subflows) {
232                 try {
233                     subflow.handleIncoming();
234                 } catch (IllegalArgumentException e) {
235                     //do nothing, continue loop
236                 }
237             }
238         }
239         return false;
240     }
241
242     @Override
243

```

```

244     public List<Packet> trySend() {
245         List<Packet> packetsSent = new ArrayList<>();
246         for (TCP subflow : this.subflows) {
247             for (Packet packet : subflow.trySend()) {
248                 packetsSent.add(packet);
249             }
250         }
251         return packetsSent;
252     }
253
254     @Override
255     public boolean canRetransmit(Packet packet) {
256         for (TCP subflow : this.subflows) {
257             if (subflow.canRetransmit(packet)) return true;
258         }
259         return false;
260     }
261
262     public TCPStats[] getTcpStats() {
263         var tcpStats = new TCPStats[this.subflows.length];
264         for (var i = 0; i < this.subflows.length; i++) {
265             tcpStats[i] = ((ClassicTCP) this.subflows[i]).getStats();
266         }
267         return tcpStats;
268     }
269
270     @Override
271     public Packet fastRetransmit() {
272         throw new IllegalStateException(DEPRECATED_STRING);
273     }
274
275     public static class MPTCPBuilder {
276
277         private int numberOfSubflows = 2;
278         private int receivingWindowCapacity = 20;
279         private Address address = new UIDAddress();
280
281         public MPTCPBuilder withNumberOfSubflows(int numberOfSubflows) {
282             this.numberOfSubflows = numberOfSubflows;
283             return this;
284         }
285
286         public MPTCPBuilder withReceivingWindowCapacity(int capacity) {
287             this.receivingWindowCapacity = capacity;
288             return this;
289         }
290
291         public MPTCPBuilder withAddress(Address address) {
292             this.address = address;
293             return this;
294         }
295
296         public MPTCP build() {
297             return new MPTCP(this.numberOfSubflows, this.receivingWindowCapacity,
298                 this.address);
299         }
300     }
301 }

```

```

1 package org.example.protocol;
2
3 import org.example.data.Packet;
4 import org.example.data.Payload;
5 import org.example.network.Channel;
6 import org.example.network.interfaces.Endpoint;
7
8 import java.util.List;
9
10 public interface TCP extends Endpoint {
11
12     /**
13      * A method that initiates connection with a host
14      * <p>
15      * 1. Send SYN with random number A
16      * 2. Receive SYN-ACK with A+1 and sequence number random B (the server chooses
17      *    B)
18      * 3. Send ACK back to server. The ACK is now A+2 and sequence number is B+1
19      *
20      * @param host to connect to
21      */
22     void connect(TCP host);
23
24     /**
25      * A method that handles incoming connections
26      * <p>
27      * 1. Receive SYN Packet with random number A that indicates that a client wants
28      *    to connect
29      * 2. Send SYN-ACK with A+1 and sequence number random B
30      * 3. Receive ACK with ACK-number A+2 and sequence number is B+1
31      *
32      * @param syn Packet to start incoming connection from a client
33      */
34     void connect(Packet syn);
35
36     /**
37      * Creates a packet with Payload and enqueues the new Packet to the output-
38      *    buffer
39      *
40      * @param payload
41      */
42     void send(Payload payload);
43
44     /**
45      * Dequeues the Packet from the received queue
46      * All packets in the received queue are acknowledged and in correct order
47      *
48      * @return next Packet
49      */
50     Packet receive();
51
52     /**
53      * A method that to determine if TCP has an open connection
54      *
55      * @return true if TCP has an active connection
56      */
57     boolean isConnected();
58
59     /**
60      * A method that returns the available channel to route packets through
61      *
62      */

```

```

60     * @return an available Channel
61     */
62     Channel getChannel();
63
64     /**
65     * A method that returns the retransmission timeout value
66     *
67     * @return the RTO for this connection
68     */
69     long getRTO();
70
71     /**
72     * A method to handle the incoming packets
73     *
74     * @return returns true if a packet should be ACKed
75     */
76     boolean handleIncoming();
77
78     /**
79     * A method to handle the sending process of TCP
80     *
81     * @return a List of the sent packets
82     */
83     List<Packet> trySend();
84
85     /**
86     * A method that checks if a packet should be retransmitted or not
87     *
88     * @param packet that should be checked for possible retransmission
89     * @return true if the packet should be retransmitted
90     */
91     boolean canRetransmit(Packet packet);
92
93     /**
94     * A method that returns a packet that should be fast retransmitted
95     *
96     * @return Packet if there is a packet to fast-retransmit, else null
97     */
98     Packet fastRetransmit();
99
100    /**
101    * A method that returns the main flow or controller of this TCP
102    *
103    * @return the main TCP flow
104    */
105    TCP getMainFlow();
106
107    /**
108    * A method that returns the number of subflows available
109    *
110    * @return the number of subflows
111    */
112    int getNumberOfFlows();
113
114
115 }

```

```
1 package org.example.protocol.window;
2
3
4 public interface IWindow {
5
6     /**
7      * A method that uses the given packet-index to determine
8      * if the packet is inside the window
9      *
10     * @param packetIndex
11     * @return true if the packet is in the window
12     */
13     boolean inWindow(int packetIndex);
14
15     /**
16     * A method that returns the window capacity
17     *
18     * @return the window capacity
19     */
20     int getWindowCapacity();
21
22 }
```

```

1 package org.example.protocol.window;
2
3 import org.example.data.Packet;
4 import org.example.util.BoundedPriorityBlockingQueue;
5
6 import java.util.Comparator;
7
8 public abstract class Window extends BoundedPriorityBlockingQueue<Packet> implements
   IWindow {
9
10    protected Window(int windowCapacity, Comparator<Packet> comparator) {
11        super(windowCapacity, comparator);
12    }
13
14    @Override
15    public boolean inWindow(int packetIndex) {
16        return packetIndex < this.getWindowCapacity() && packetIndex >= 0;
17    }
18
19    @Override
20    public int getWindowCapacity() {
21        return this.bound();
22    }
23
24    @Override
25    public String toString() {
26        var stringBuilder = new StringBuilder();
27        for (Packet packet : this) {
28            stringBuilder.append("[");
29            stringBuilder.append(packet);
30            stringBuilder.append("]");
31            stringBuilder.append("\n");
32        }
33        return stringBuilder.toString();
34    }
35 }

```



```

1 package org.example.protocol.window.receiving;
2
3 import org.example.data.Packet;
4 import org.example.network.interfaces.Endpoint;
5 import org.example.protocol.Connection;
6 import org.example.protocol.window.IWindow;
7 import org.example.protocol.window.sending.SendingWindow;
8 import org.example.util.BoundedQueue;
9
10 public interface ReceivingWindow extends IWindow, BoundedQueue<Packet> {
11
12     /**
13      * A method that receives incoming packets and returns true
14      * if a packet was added to the received packets queue
15      *
16      * @param sendingWindow
17      * @return true if a non ACK packet was received
18      */
19     boolean receive(SendingWindow sendingWindow);
20
21     /**
22      * A method that returns the packet to ACK
23      *
24      * @param endpointToReceiveAck the endpoint to receive the ACK
25      * @return the packet to ACK
26      */
27     Packet ackThis(Endpoint endpointToReceiveAck);
28
29     /**
30      * A method that checks if a ACK should be sent or not
31      *
32      * @return
33      */
34     boolean shouldAck();
35
36     /**
37      * A method that calculates the receiving window packet index
38      * of the given packet
39      *
40      * @param packet to calculate packet index of
41      * @param connection to use in the calculation
42      * @return the packet index of the given packet
43      */
44     int receivingPacketIndex(Packet packet, Connection connection);
45
46     /**
47      * A method that checks if a packet is inside the receiving window
48      *
49      * @param packet packet to check
50      * @param connection to use in calculation
51      * @return true if the packet is inside the receiving window
52      */
53     boolean inReceivingWindow(Packet packet, Connection connection);
54
55 }

```

```

1 package org.example.protocol.window.receiving;
2
3 import org.example.data.Flag;
4 import org.example.data.Packet;
5 import org.example.data.PacketBuilder;
6 import org.example.network.interfaces.Endpoint;
7 import org.example.protocol.Connection;
8 import org.example.protocol.window.Window;
9 import org.example.protocol.window.sending.SendingWindow;
10
11 import java.util.Comparator;
12 import java.util.HashMap;
13 import java.util.List;
14 import java.util.Map;
15
16 public class SelectiveRepeat extends Window implements ReceivingWindow {
17
18     private final List<Packet> received;
19     private int packetCount;
20     private Map<Endpoint, Packet> ackThisMap;
21
22     public SelectiveRepeat(int windowSize, Comparator<Packet> comparator, List<
23         Packet> receivedContainer) {
24         super(windowSize, comparator);
25         this.received = receivedContainer;
26         this.packetCount = 0;
27         this.ackThisMap = new HashMap<>();
28     }
29
30     private void receive(Packet packet) {
31         if (this.received.contains(packet)) return;
32         boolean added = this.received.add(packet);
33         if (!added) throw new IllegalStateException("Packet was not added to the
34             received queue");
35     }
36
37     @Override
38     public boolean receive(SendingWindow sendingWindow) {
39         if (this.isEmpty()) return false;
40         if (this.peek().hasAllFlags(Flag.SYN)) return false;
41
42         var connection = sendingWindow.getConnection();
43
44         if (this.peek().hasAllFlags(Flag.ACK)) {
45             if (this.peek().getOrigin().equals(connection.getConnectedNode())) {
46                 sendingWindow.ackReceived(this.peek());
47                 this.poll();
48             }
49             return false;
50         }
51
52         var packetToUpdateWith = this.ackThisMap.get(connection.getConnectedNode());
53         if (packetToUpdateWith != null) connection.update(packetToUpdateWith);
54
55         if (this.inReceivingWindow(this.peek(), connection)) {
56             while (receivingPacketIndex(this.peek(), connection) == 0 && this.peek()
57                 .getIndex() == this.packetCount) {
58                 connection.update(this.peek());
59                 this.ackThisMap.put(this.peek().getOrigin(), this.peek());
60                 this.receive(this.peek());
61                 var packetRemoved = this.remove();
62                 if (packetRemoved == null) throw new IllegalStateException("removing

```

```

60         null packet");
61         sendingWindow.getStats().packetDeparture(packetRemoved);
62         this.packetCount++;
63         if (this.isEmpty()) return true;
64     }
65     return true; // true so that duplicate AKCs are sent
66 }
67
68
69 @Override
70 public boolean shouldAck() {
71     return true;
72 }
73
74 @Override
75 public Packet ackThis(Endpoint endpointToReceiveAck) {
76     return this.ackThisMap.getOrDefault(endpointToReceiveAck, new PacketBuilder
77         ().build());
78 }
79
80 @Override
81 public int receivingPacketIndex(Packet packet, Connection connection) {
82     int seqNum = packet.getSequenceNumber();
83     int ackNum = connection.getNextAcknowledgementNumber();
84     return seqNum - ackNum;
85 }
86
87 @Override
88 public boolean inReceivingWindow(Packet packet, Connection connection) {
89     int packetIndex = receivingPacketIndex(packet, connection);
90     return inWindow(packetIndex);
91 }
92
93 @Override
94 public boolean offer(Packet packet) {
95     if (this.isFull()) {
96         return false;
97     }
98     if (this.contains(packet) && !packet.hasAllFlags(Flag.ACK)) {
99         return false;
100     }
101     return super.offer(packet);
102 }
103 }

```

```

1 package org.example.protocol.window.sending;
2
3 import org.example.data.Packet;
4 import org.example.protocol.Connection;
5 import org.example.protocol.window.IWindow;
6 import org.example.simulator.statistics.TCPStats;
7 import org.example.util.BoundedQueue;
8
9 public interface SendingWindow extends IWindow, BoundedQueue<Packet> {
10
11     /**
12      * A method that checks if the sendigwindow should wait for ACK or not
13      *
14      * @return true if the sendingwindow should wait
15      */
16     boolean isWaitingForAck();
17
18     /**
19      * A method that handles received ACKs
20      *
21      * @param ack
22      */
23     void ackReceived(Packet ack);
24
25     /**
26      * A method that returns the packet that should be sent
27      *
28      * @return packet to be sent
29      */
30     Packet send();
31
32     /**
33      * A method that checks if a packet should be retransmitted or not
34      *
35      * @param packet that should be checked for possible retransmission
36      * @return true if the packet should be retransmitted
37      */
38     boolean canRetransmit(Packet packet);
39
40     /**
41      * A method that returns a packet that should be fast retransmitted
42      *
43      * @return Packet if there is a packet to fast-retransmit, else null
44      */
45     Packet fastRetransmit();
46
47     /**
48      * A method that is called to increase the size of the sendigwindow.
49      */
50     void increase();
51
52     /**
53      * A method that is called to decrease the size of the sendigwindow.
54      */
55     void decrease();
56
57     /**
58      * A method that calculates the sending window packet index
59      * of the given packet
60      *
61      * @param packet to calculate packet index of
62      * @return the packet index of the given packet

```

```
63     */
64     int sendingPacketIndex(Packet packet);
65
66     /**
67      * A method that returns the connection
68      *
69      * @return connection
70      */
71     Connection getConnection();
72
73     /**
74      * A method that checks if the sending window is empty or not
75      *
76      * @return true if the sending window is empty
77      */
78     boolean isQueueEmpty();
79
80     /**
81      * A method that returns the stats of the sending window
82      *
83      * @return the TCPStats
84      */
85     TCPStats getStats();
86
87 }
```

```

1 package org.example.protocol.window.sending;
2
3 import org.example.data.Packet;
4 import org.example.data.PacketBuilder;
5 import org.example.data.Payload;
6 import org.example.protocol.Connection;
7 import org.example.protocol.window.Window;
8 import org.example.simulator.statistics.TCPStats;
9 import org.example.util.BoundedQueue;
10 import org.javatuples.Pair;
11
12 import java.util.Comparator;
13 import java.util.List;
14
15 public class SlidingWindow extends Window implements SendingWindow, BoundedQueue<
    Packet> {
16
17     private static final int DEFAULT_CONGESTION_WINDOW_CAPACITY = 1;
18     private final int numDupAckFastRetransmitTrigger;
19     private final List<Pair<Integer, Payload>> payloadsToSend;
20     private final int receiverWindowSize;
21     private final boolean isReno;
22     private final Connection connection;
23     private final TCPStats stats;
24     private boolean seriousLossDetected = false;
25     private int numPacketsReceivedWithoutIncreasingWindow;
26     private int ssthresh;
27     private int dupAckCount;
28     private boolean fastRetransmitted;
29
30     public SlidingWindow(int receiverWindowCapacity, boolean isReno, Connection
        connection, Comparator<Packet> comparator, List<Pair<Integer, Payload>>
        payloadsToSend, TCPStats stats) {
31         super(DEFAULT_CONGESTION_WINDOW_CAPACITY, comparator);
32         this.payloadsToSend = payloadsToSend;
33         this.receiverWindowSize = receiverWindowCapacity;
34         this.numPacketsReceivedWithoutIncreasingWindow = 0;
35         this.ssthresh = receiverWindowCapacity; //initial value is essentially no
            ssthresh
36         this.isReno = isReno;
37         this.dupAckCount = 0;
38         this.fastRetransmitted = false;
39         this.connection = connection;
40         this.numDupAckFastRetransmitTrigger = 3;
41         this.stats = stats;
42
43     }
44
45     @Override
46     public boolean isWaitingForAck() {
47         return this.isFull();
48     }
49
50     @Override
51     public void ackReceived(Packet ack) {
52         this.stats.ackReceived();
53         int ackIndex = this.sendingPacketIndex(ack);
54
55         boolean dupAck = ackIndex == -1;
56         if (dupAck) {
57             this.dupAckCount++;
58             return;

```

```

59     }
60     this.dupAckCount = 0;
61
62     for (var i = 0; i <= ackIndex; i++) {
63         if (this.isEmpty()) break;
64         this.poll();
65         this.increase();
66         this.seriousLossDetected = false;
67     }
68     this.connection.update(ack);
69 }
70
71 @Override
72 public Packet send() {
73     int nextPacketSeqNum = this.connection.getNextSequenceNumber() + this.size()
74     ;
75     Pair<Integer, Payload> indexAndPayload = this.payloadsToSend.remove(0);
76
77     var packet = new PacketBuilder()
78         .withConnection(this.connection)
79         .withPayload(indexAndPayload.getValue1())
80         .withSequenceNumber(nextPacketSeqNum)
81         .withIndex(indexAndPayload.getValue0())
82         .build();
83
84     if (this.contains(packet)) throw new IllegalStateException("can't add same
85         packet twice");
86     if (super.offer(packet)) {
87         return packet;
88     }
89     return null;
90 }
91
92 @Override
93 public boolean canRetransmit(Packet packet) {
94     if (this.contains(packet)) {
95         if (this.sendingPacketIndex(packet) == 0) {
96             if (this.fastRetransmitted) this.fastRetransmitted = false;
97             else this.decrease();
98         }
99         if (this.sendingPacketIndex(packet) >= this.getWindowCapacity() - 1)
100             this.seriousLossDetected = true;
101         this.dupAckCount = 0;
102         return true;
103     }
104     return false;
105 }
106
107 @Override
108 public Packet fastRetransmit() {
109     if (this.dupAckCount > this.numDupAckFastRetransmitTrigger) {
110         this.dupAckCount = 0;
111         this.decrease(false);
112         this.fastRetransmitted = true;
113         return this.peek();
114     }
115     return null;
116 }
117
118 @Override
119 public void increase() {
120     if (this.getWindowCapacity() >= this.receiverWindowSize) return;

```

```

119         if (this.ssthresh > this.getWindowCapacity()) slowStart();
120         else congestionAvoidance();
121     }
122
123     private void slowStart() {
124         this.setBound(this.getWindowCapacity() + 1);
125         this.stats.trackCwnd(this.getWindowCapacity());
126     }
127
128     private void congestionAvoidance() {
129         boolean allPacketsInOneWindowReceived =
130             this.numPacketsReceivedWithoutIncreasingWindow % this.
131                 getWindowCapacity() == 0
132                 && this.numPacketsReceivedWithoutIncreasingWindow != 0;
133         if (allPacketsInOneWindowReceived) {
134             this.numPacketsReceivedWithoutIncreasingWindow = 0;
135             this.setBound(this.getWindowCapacity() + 1);
136             this.stats.trackCwnd(this.getWindowCapacity());
137             return;
138         }
139         this.numPacketsReceivedWithoutIncreasingWindow++;
140     }
141
142     @Override
143     public void decrease() {
144         this.decrease(true);
145     }
146
147     public void decrease(boolean timeout) {
148         this.stats.trackCwnd(this.getWindowCapacity());
149         this.ssthresh = this.findNewSSThresh();
150         int newWindowSize = this.findNewWindowSize(timeout);
151         this.setBound(newWindowSize);
152         this.stats.trackCwnd(this.getWindowCapacity());
153     }
154
155     private int findNewWindowSize(boolean timeout) {
156         if (this.isReno && !timeout) return this.ssthresh;
157         return DEFAULT_CONGESTION_WINDOW_CAPACITY;
158     }
159
160     private int findNewSSThresh() {
161         return Math.max((int) (this.getWindowCapacity() / 2.0),
162             DEFAULT_CONGESTION_WINDOW_CAPACITY);
163     }
164
165     @Override
166     public int sendingPacketIndex(Packet packet) {
167         int packetSeqNum = packet.getSequenceNumber();
168         int connSeqNum = this.connection.getNextSequenceNumber();
169         return packetSeqNum - connSeqNum;
170     }
171
172     @Override
173     public Connection getConnection() {
174         return connection;
175     }
176
177     @Override
178     public boolean isQueueEmpty() {
179         return this.payloadsToSend.isEmpty();

```



```
180     }
181
182     @Override
183     public TCPStats getStats() {
184         return this.stats;
185     }
186 }
```

```

1 package org.example.simulator;
2
3 import org.example.simulator.events.Event;
4 import org.example.util.Util;
5
6 import java.util.PriorityQueue;
7 import java.util.Queue;
8
9 public class EventHandler {
10
11     private final Queue<Event> events;
12
13     public EventHandler() {
14         this.events = new PriorityQueue<>();
15     }
16
17     public void addEvent(Event event) {
18         this.events.add(event);
19     }
20
21     public Event peekEvent() {
22         return this.events.peek();
23     }
24
25     public Queue<Event> getEvents() {
26         return events;
27     }
28
29     public int getNumberOfEvents() {
30         return this.events.size();
31     }
32
33     public boolean singleRun() {
34         var event = this.events.poll();
35         if (event == null) {
36             if (!this.events.isEmpty())
37                 throw new IllegalStateException("get null event when events queue
38                     are nonempty!");
39             return false;
40         }
41         Util.tickTime(event);
42         event.run();
43         event.generateNextEvent(this.events);
44         return true;
45     }
46
47     public void run() {
48         while (singleRun()) ;
49     }

```

```

1 package org.example.simulator.events;
2
3 import org.example.network.Channel;
4 import org.example.network.interfaces.Endpoint;
5 import org.example.network.interfaces.NetworkNode;
6 import org.example.protocol.TCP;
7 import org.example.simulator.events.run.RunEndpointEvent;
8 import org.example.simulator.events.run.RunNetworkNodeEvent;
9 import org.example.simulator.events.tcp.RunTCPEvent;
10
11 import java.util.Queue;
12
13 public class ChannelEvent extends Event {
14
15     private final Channel channel;
16     private boolean channelSuccess;
17
18     public ChannelEvent(Channel channel) {
19         super(channel);
20         this.channel = channel;
21         this.channelSuccess = false;
22     }
23
24     @Override
25     public void run() {
26         this.channelSuccess = this.channel.channel();
27     }
28
29     @Override
30     public void generateNextEvent(Queue<Event> events) {
31         if (this.channelSuccess) {
32             NetworkNode nextNode = this.channel.getDestination();
33             if (nextNode instanceof TCP) {
34                 var tcp = ((TCP) nextNode).getMainFlow();
35                 events.add(new RunTCPEvent(tcp));
36                 return;
37             }
38             assert !nextNode.inputBufferIsEmpty() : "The next NetworkNode has no
39                 packet in the input buffer";
40             if (nextNode instanceof Endpoint) {
41                 events.add(new RunEndpointEvent((Endpoint) nextNode));
42                 return;
43             }
44             events.add(new RunNetworkNodeEvent(nextNode));
45             return;
46         }
47         NetworkNode nextNode = this.channel.getDestination();
48         if (nextNode instanceof TCP) {
49             var tcp = ((TCP) nextNode).getMainFlow();
50             if (!tcp.inputBufferIsEmpty()) return;
51             events.add(new RunTCPEvent(tcp));
52         }
53     }
54
55     @Override
56     public boolean equals(Object obj) {
57         return super.equals(obj);
58     }
59
60     @Override
61     public int hashCode() {

```

```
62     return super.hashCode();
63 }
64 }
```

```

1 package org.example.simulator.events;
2
3 import org.example.network.Channel;
4 import org.example.network.interfaces.NetworkNode;
5 import org.example.util.Util;
6
7 import java.util.Queue;
8
9 public abstract class Event implements Comparable<Event> {
10
11     private final long instant;
12
13     protected Event(long delay) {
14         this.instant = this.findInstant(delay);
15     }
16
17     protected Event() {
18         this.instant = findInstant(0);
19     }
20
21     protected Event(NetworkNode node) {
22         if (node == null) throw new IllegalArgumentException("node is null");
23         this.instant = this.findInstant(node.delay());
24     }
25
26     protected Event(Channel channel) {
27         if (channel == null) throw new IllegalArgumentException("channel is null");
28         this.instant = this.findInstant(channel.propagationDelay());
29     }
30
31     private long findInstant(long delay) {
32         return Util.getTime() + delay;
33     }
34
35     public abstract void run();
36
37     public abstract void generateNextEvent(Queue<Event> events);
38
39     public long getInstant() {
40         return this.instant;
41     }
42
43     @Override
44     public int compareTo(Event o) {
45         return Long.compare(this.getInstant(), o.getInstant());
46     }
47
48     @Override
49     public boolean equals(Object obj) {
50         return super.equals(obj);
51     }
52
53     @Override
54     public int hashCode() {
55         return super.hashCode();
56     }
57
58     @Override
59     public String toString() {
60         return this.getClass().getSimpleName();
61     }
62 }

```

```
1 package org.example.simulator.events;
2
3 public abstract class EventGenerator extends Event {
4
5     protected EventGenerator(long delay) {
6         super(delay);
7     }
8
9     @Override
10    public void run() {
11        //EventGenerators should only generate new events and not run
12    }
13 }
```

```

1 package org.example.simulator.events;
2
3 import org.example.data.Packet;
4 import org.example.network.Channel;
5 import org.example.network.interfaces.Endpoint;
6 import org.example.protocol.MPTCP;
7
8 import java.util.List;
9 import java.util.Queue;
10
11 public class RouteEvent extends Event {
12
13     private final Packet packet;
14     private final Endpoint endpoint;
15
16     public RouteEvent(Endpoint endpoint, Packet packet) {
17         super();
18         if (endpoint instanceof MPTCP) throw new IllegalArgumentException("Should
19             only handle Regular tcp or subflows");
20         this.endpoint = endpoint;
21         this.packet = packet;
22     }
23
24     @Override
25     public void run() {
26         this.endpoint.route(this.packet);
27     }
28
29     @Override
30     public void generateNextEvent(Queue<Event> events) {
31         List<Channel> channelsUsed = this.endpoint.getChannelsUsed();
32         var channel = channelsUsed.get(0);
33         if (channel == null) return;
34         events.add(new ChannelEvent(channel));
35     }
36
37     @Override
38     public boolean equals(Object obj) {
39         return super.equals(obj);
40     }
41
42     @Override
43     public int hashCode() {
44         return super.hashCode();
45     }
46 }

```

```

1 package org.example.simulator.events.run;
2
3 import org.example.network.interfaces.Endpoint;
4 import org.example.simulator.events.Event;
5
6 import java.util.Queue;
7
8
9 public class RunEndpointEvent extends Event {
10
11     private final Endpoint endpoint;
12
13     public RunEndpointEvent(Endpoint endpoint) {
14         super(endpoint);
15         this.endpoint = endpoint;
16     }
17
18     @Override
19     public void run() {
20         this.endpoint.run();
21     }
22
23     @Override
24     public void generateNextEvent(Queue<Event> events) {
25         // do nothing because an endpoint without tcp functionality is not doing
26         // anything
27     }
28
29     @Override
30     public boolean equals(Object obj) {
31         return super.equals(obj);
32     }
33
34     @Override
35     public int hashCode() {
36         return super.hashCode();
37     }
38 }

```



```

1 package org.example.simulator.events.run;
2
3 import org.example.network.Channel;
4 import org.example.network.interfaces.NetworkNode;
5 import org.example.simulator.events.ChannelEvent;
6 import org.example.simulator.events.Event;
7
8 import java.util.List;
9 import java.util.Queue;
10
11 public class RunNetworkNodeEvent extends Event {
12
13     private final NetworkNode node;
14
15     public RunNetworkNodeEvent(NetworkNode node) {
16         super(node);
17         this.node = node;
18     }
19
20     @Override
21     public void run() {
22         assert !this.node.inputBufferIsEmpty() : "RunNetworkNodeEvent added, but no
23             packet to be sent";
24         this.node.run();
25     }
26
27     @Override
28     public void generateNextEvent(Queue<Event> events) {
29         List<Channel> channelsUsed = this.node.getChannelsUsed();
30         assert channelsUsed.size() == 1 : "Multiple channels used in one
31             NetworkNodeEvent";
32         var channel = channelsUsed.get(0);
33         events.add(new ChannelEvent(channel));
34     }
35
36     @Override
37     public boolean equals(Object obj) {
38         return super.equals(obj);
39     }
40
41     @Override
42     public int hashCode() {
43         return super.hashCode();
44     }
45 }

```

```

1 package org.example.simulator.events.tcp;
2
3 import org.example.data.Packet;
4 import org.example.network.Channel;
5 import org.example.protocol.MPTCP;
6 import org.example.protocol.TCP;
7 import org.example.simulator.events.ChannelEvent;
8 import org.example.simulator.events.Event;
9
10 import java.util.List;
11 import java.util.Queue;
12
13 public class RunTCPEvent extends Event {
14
15     private final TCP tcp;
16     private List<Packet> packetsSent;
17     private boolean scheduleFirstSend;
18
19     public RunTCPEvent(TCP tcp) {
20         super(tcp);
21         this.tcp = tcp;
22         this.scheduleFirstSend = false;
23     }
24
25     public RunTCPEvent(TCP tcp, long delay) {
26         super(delay);
27         this.tcp = tcp;
28         this.scheduleFirstSend = false;
29     }
30
31     @Override
32     public void run() {
33         if (this.tcp instanceof MPTCP) {
34             for (TCP subflow : ((MPTCP) this.tcp).getSubflows()) {
35                 if (subflow.isConnected()) {
36                     this.tcp.handleIncoming();
37                     this.packetsSent = this.tcp.trySend();
38                     return;
39                 }
40             }
41         } else if (this.tcp.isConnected()) {
42             this.tcp.handleIncoming();
43             this.packetsSent = this.tcp.trySend();
44             return;
45         }
46         this.tcp.handleIncoming();
47         this.scheduleFirstSend = true;
48     }
49
50     @Override
51     public void generateNextEvent(Queue<Event> events) {
52         for (Channel channel : this.tcp.getChannelsUsed()) {
53             events.add(new ChannelEvent(channel));
54         }
55
56         if (this.scheduleFirstSend) {
57             //add delay before sending first packet
58             //this is important if we have a "downloading situation"
59             events.add(new RunTCPEvent(this.tcp, 1000000));
60             return;
61         }
62     }

```

```
63
64     if (packetsSent == null) return;
65     for (Packet packet : this.packetsSent) {
66         events.add(new TCPRetransmitEventGenerator(packet, 0));
67     }
68 }
69
70 @Override
71 public boolean equals(Object obj) {
72     return super.equals(obj);
73 }
74
75 @Override
76 public int hashCode() {
77     return super.hashCode();
78 }
79 }
```

```

1 package org.example.simulator.events.tcp;
2
3 import org.example.network.Channel;
4 import org.example.protocol.TCP;
5 import org.example.simulator.events.ChannelEvent;
6 import org.example.simulator.events.Event;
7
8 import java.util.Queue;
9
10 public class TCPConnectEvent extends Event {
11
12     private final TCP client;
13     private final TCP host;
14     private final int numAttempts;
15
16     private TCPConnectEvent(int delay, TCP client, TCP host, int numAttempts) {
17         super(delay);
18         this.client = client;
19         this.host = host;
20         this.numAttempts = numAttempts;
21     }
22
23     public TCPConnectEvent(TCP client, TCP host) {
24         super();
25         this.client = client;
26         this.host = host;
27         this.numAttempts = 0;
28     }
29
30     @Override
31     public void run() {
32         if (this.client.isConnected()) return;
33         this.client.connect(this.host);
34     }
35
36     @Override
37     public void generateNextEvent(Queue<Event> events) {
38         if (this.numAttempts > this.client.getNumberOfFlows() * 3) return;
39         if (this.client.isConnected()) return;
40
41         for (Channel channel : this.client.getChannelsUsed()) {
42             events.add(new ChannelEvent(channel));
43         }
44         events.add(new TCPConnectEvent(10000000, this.client, this.host, this.
45             numAttempts + 1));
46     }
47
48     @Override
49     public boolean equals(Object obj) {
50         return super.equals(obj);
51     }
52
53     @Override
54     public int hashCode() {
55         return super.hashCode();
56     }
57 }

```

```

1 package org.example.simulator.events.tcp;
2
3 import org.example.data.Packet;
4 import org.example.protocol.ClassicTCP;
5 import org.example.protocol.TCP;
6 import org.example.simulator.events.Event;
7 import org.example.simulator.events.EventGenerator;
8 import org.example.simulator.events.RouteEvent;
9
10 import java.util.Queue;
11
12 public class TCPRetransmitEventGenerator extends EventGenerator {
13
14     private final TCP tcp;
15     private final Packet packet;
16     private final int numAttempts;
17
18     public TCPRetransmitEventGenerator(Packet packet, int numAttempts) {
19         super(((TCP) packet.getOrigin()).getRTO());
20         this.tcp = (TCP) packet.getOrigin();
21         this.packet = packet;
22         this.numAttempts = numAttempts;
23     }
24
25     @Override
26     public void generateNextEvent(Queue<Event> events) {
27         if (this.numAttempts > 3) return;
28         if (!this.tcp.isConnected()) return;
29         if (this.tcp.canRetransmit(this.packet)) {
30             ((ClassicTCP) this.tcp).getStats().packetRetransmit();
31             events.add(new RouteEvent(this.tcp, this.packet));
32             events.add(new TCPRetransmitEventGenerator(this.packet, this.numAttempts
33                 + 1));
34         }
35
36         @Override
37         public boolean equals(Object obj) {
38             return super.equals(obj);
39         }
40
41         @Override
42         public int hashCode() {
43             return super.hashCode();
44         }
45
46
47 }

```

```

1 package org.example.simulator.statistics;
2
3 import com.fasterxml.jackson.core.JsonProcessingException;
4 import com.fasterxml.jackson.databind.ObjectMapper;
5 import org.example.data.Packet;
6 import org.example.util.Util;
7 import org.knowm.xchart.VectorGraphicsEncoder;
8 import org.knowm.xchart.XYChart;
9 import org.knowm.xchart.XYChartBuilder;
10 import org.knowm.xchart.XYSeries;
11 import org.knowm.xchart.style.Styler;
12
13 import java.awt.*;
14 import java.io.IOException;
15 import java.util.ArrayList;
16 import java.util.Collections;
17 import java.util.List;
18
19
20 public abstract class Stats {
21
22     protected static final String DIR = "./charts/";
23     protected static final Styler.ChartTheme theme = Styler.ChartTheme.Matlab;
24     protected static final int CHART_WIDTH = 1500;
25     protected static final int CHART_HEIGHT = 600;
26     protected static final Font TITLE_FONT = new Font("titleFont", 0, 30);
27     protected static final Font AXIS_FONT = new Font("axisFont", 0, 25);
28     protected static final Font TICK_FONT = new Font("tickFont", 0, 15);
29
30     protected static final int TIMESCALE = 1000;
31
32     // Arrival
33     protected final ArrayList<Packet> arrivalPacket = new ArrayList<>();
34     protected final ArrayList<Double> arrivalTime = new ArrayList<>();
35     protected final ArrayList<Integer> arrivalNum = new ArrayList<>();
36
37     // Departure
38     protected final ArrayList<Packet> departurePacket = new ArrayList<>();
39     protected final ArrayList<Double> departureTime = new ArrayList<>();
40     protected final ArrayList<Integer> departureNum = new ArrayList<>();
41
42     // Inter arrival times
43     protected final ArrayList<Double> interArrivalTimes = new ArrayList<>();
44
45     // Time in system
46     protected final ArrayList<Double> timeInSystem = new ArrayList<>();
47
48     // number of packets in system
49     protected final ArrayList<Integer> numPacketsInSystem = new ArrayList<>();
50
51     private double meanArrivalRate;
52     private double meanTimeInSystem;
53     private double meanNumPacketsInSystem;
54
55     protected abstract String fileName();
56
57     protected abstract void additionalCalculations();
58
59     private void doCalculations() {
60         if (this.interArrivalTimes.isEmpty()) return;
61         this.meanArrivalRate = this.arrivalNum.size() / (this.arrivalTime.get(this.
            arrivalTime.size() - 1));

```

```

62
63     this.meanTimeInSystem = this.timeInSystem.stream().mapToDouble(Double::
        doubleValue).average().getAsDouble();
64
65     //Little's law
66     // L = 1/E[A] * W
67     this.meanNumPacketsInSystem = this.meanArrivalRate * this.meanTimeInSystem;
68 }
69
70 public void packetArrival(Packet packet) {
71     this.arrivalPacket.add(packet);
72     this.arrivalNum.add(this.arrivalNum.size());
73     this.arrivalTime.add((double) Util.seeTime() / TIMESCALE);
74
75     // number of packets in system
76     this.addPacketsInSystem();
77
78     // inter arrival time calculation
79     this.addInterArrivalTime();
80 }
81
82 public void packetDeparture(Packet packet) {
83     this.departurePacket.add(packet);
84     this.departureNum.add(this.arrivalNum.size());
85     this.departureTime.add((double) Util.seeTime() / TIMESCALE);
86
87     // number of packets in system
88     this.removePacketsInSystem();
89
90     // Time in system calculation
91     this.addTimeInSystem(packet);
92 }
93
94 private void addTimeInSystem(Packet packet) {
95     int packetArrivalIndex = this.arrivalIndexof(packet);
96     int packetDepartureIndex = this.departureIndexof(packet);
97
98     if (packetArrivalIndex == -1 || packetDepartureIndex == -1) {
99         return;
100     }
101
102     double arrival = this.arrivalTime.get(packetArrivalIndex);
103     double departure = this.departureTime.get(packetDepartureIndex);
104     double inSystem = departure - arrival;
105
106     double minTime = (double) 10 / (double) TIMESCALE;
107     if (inSystem < minTime) inSystem = minTime;
108     if (inSystem < 0) throw new IllegalStateException("packet cannot be in
        system a negative amount of time");
109     this.timeInSystem.add(inSystem);
110 }
111
112 private void addInterArrivalTime() {
113     int n = this.arrivalNum.size() - 1;
114     if (n > 1) {
115         double interArrivalTime = this.arrivalTime.get(n) - this.arrivalTime.get
            (n - 1);
116         if (interArrivalTime < 0) throw new IllegalStateException("interarrival
            time is less than 0");
117         this.interArrivalTimes.add(interArrivalTime);
118     } else {
119         this.interArrivalTimes.add(0.0);
120     }

```

```

121     }
122
123     private void addPacketsInSystem() {
124         if (this.numPacketsInSystem.isEmpty()) {
125             this.numPacketsInSystem.add(1);
126         } else {
127             this.numPacketsInSystem.add(this.numPacketsInSystem.get(this.
                numPacketsInSystem.size() - 1) + 1);
128         }
129     }
130
131     private void removePacketsInSystem() {
132         if (!this.numPacketsInSystem.isEmpty()) {
133             this.numPacketsInSystem.add(this.numPacketsInSystem.get(this.
                numPacketsInSystem.size() - 1) - 1);
134         }
135     }
136
137     private int arrivalIndexof(Packet packet) {
138         return this.arrivalPacket.indexOf(packet);
139     }
140
141     private int departureIndexof(Packet packet) {
142         return this.departurePacket.indexOf(packet);
143     }
144
145
146     public void createArrivalChart() {
147         if (this.arrivalTime.isEmpty() || this.arrivalNum.isEmpty()) return;
148         if (this.arrivalTime.size() != this.arrivalNum.size())
149             throw new IllegalStateException("the arrays must be of equal length");
150
151         XYChart chart = new XYChartBuilder()
152             .width(CHART_WIDTH)
153             .height(CHART_HEIGHT)
154             .xAxisTitle("Packet Arrival-Time")
155             .yAxisTitle("Packet")
156             .title("Packet Arrivals")
157             .theme(theme)
158             .build();
159         chart.getStyler().setChartTitleFont(TITLE_FONT);
160         chart.getStyler().setAxisTitleFont(AXIS_FONT);
161         chart.getStyler().setAxisTickLabelsFont(TICK_FONT);
162         chart.getStyler().setLegendVisible(false);
163         chart.getStyler().setXAxisLabelRotation(45);
164         chart.addSeries("Packet Arrivals", this.arrivalTime, this.arrivalNum);
165         chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
            Scatter);
166         saveChart(chart, "ArrivalChart_");
167     }
168
169     public void createDepartureChart() {
170         if (this.departureTime.isEmpty() || this.departureNum.isEmpty()) return;
171         if (this.departureTime.size() != this.departureNum.size())
172             throw new IllegalStateException("the arrays must be of equal length");
173
174         XYChart chart = new XYChartBuilder()
175             .width(CHART_WIDTH)
176             .height(CHART_HEIGHT)
177             .xAxisTitle("Packet Departure-Time")
178             .yAxisTitle("Packet")
179             .title("Packet Departures")
180             .theme(theme)

```



```

181         .build();
182     chart.getStyler().setChartTitleFont(TITLE_FONT);
183     chart.getStyler().setAxisTitleFont(AXIS_FONT);
184     chart.getStyler().setAxisTickLabelsFont(TICK_FONT);
185     chart.getStyler().setLegendVisible(false);
186     chart.getStyler().setXAxisLabelRotation(45);
187     chart.addSeries("Packet Departures", this.departureTime, this.departureNum);
188     chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
        Scatter);
189     saveChart(chart, "DepartureChart_");
190 }
191
192 public void createInterArrivalChart() {
193     XYChart chart = new XYChartBuilder()
194         .width(CHART_WIDTH)
195         .height(CHART_HEIGHT)
196         .xAxisTitle("Arrival Time")
197         .yAxisTitle("Interarrival Time")
198         .title("Packet Interarrival-Time")
199         .theme(theme)
200         .build();
201     chart.getStyler().setChartTitleFont(TITLE_FONT);
202     chart.getStyler().setAxisTitleFont(AXIS_FONT);
203     chart.getStyler().setAxisTickLabelsFont(TICK_FONT);
204     chart.getStyler().setLegendVisible(false);
205     chart.getStyler().setXAxisLabelRotation(45);
206     chart.addSeries("Interarrival Times", this.arrivalTime, this.
        interArrivalTimes);
207     chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
        Scatter);
208     saveChart(chart, "InterarrivalTime_");
209 }
210
211 public void createTimeInSystemChart() {
212     XYChart chart = new XYChartBuilder()
213         .width(CHART_WIDTH)
214         .height(CHART_HEIGHT)
215         .xAxisTitle("Time")
216         .yAxisTitle("Time In System")
217         .title("Packet Time In System")
218         .theme(theme)
219         .build();
220     chart.getStyler().setChartTitleFont(TITLE_FONT);
221     chart.getStyler().setAxisTitleFont(AXIS_FONT);
222     chart.getStyler().setAxisTickLabelsFont(TICK_FONT);
223     chart.getStyler().setLegendVisible(false);
224     chart.getStyler().setXAxisLabelRotation(45);
225     chart.addSeries("Time in system", this.departureTime, this.timeInSystem);
226     chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
        Scatter);
227     saveChart(chart, "TimeInSystem_");
228 }
229
230 public void createNumberOfPacketsInSystemChart() {
231     XYChart chart = new XYChartBuilder()
232         .width(CHART_WIDTH)
233         .height(CHART_HEIGHT)
234         .xAxisTitle("Time")
235         .yAxisTitle("Number of packets in system")
236         .title("number of packets in system")
237         .theme(theme)
238         .build();
239     chart.getStyler().setChartTitleFont(TITLE_FONT);

```

```

240     chart.getStyler().setAxisTitleFont (AXIS_FONT);
241     chart.getStyler().setAxisTickLabelsFont (TICK_FONT);
242     chart.getStyler().setLegendVisible(false);
243     chart.getStyler().setXAxisLabelRotation(45);
244     chart.addSeries("Packets in system", this.combine(this.arrivalTime, this.
        departureTime), this.numPacketsInSystem);
245     chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
        Scatter);
246     saveChart(chart, "PacketsInSystem_");
247 }
248
249 private List<Double> combine(List<Double> l1, List<Double> l2) {
250     List<Double> result = new ArrayList<>();
251     result.addAll(l1);
252     result.addAll(l2);
253     Collections.sort(result);
254     return result;
255 }
256
257 public double getMeanArrivalRate() {
258     return meanArrivalRate;
259 }
260
261 public double getMeanTimeInSystem() {
262     return meanTimeInSystem;
263 }
264
265 public double getMeanNumPacketsInSystem() {
266     return meanNumPacketsInSystem;
267 }
268
269
270 protected void saveChart(XYChart chart, String chartName) {
271     try {
272         VectorGraphicsEncoder.saveVectorGraphic(chart, DIR + chartName +
            fileName(), VectorGraphicsEncoder.VectorGraphicsFormat.SVG);
273     } catch (IOException e) {
274     }
275 }
276
277 @Override
278 public String toString() {
279     this.doCalculations();
280     this.additionalCalculations();
281     var mapper = new ObjectMapper();
282     try {
283         var formattedString = mapper.writeValueAsString(this)
284             .replace("{", "{\n        ")
285             .replace("}", "\n}")
286             .replace(",", ",\n        ");
287         return this.fileName() + " " + formattedString;
288     } catch (JsonProcessingException e) {
289         return "fail";
290     }
291 }
292 }

```

```

1 package org.example.simulator.statistics;
2
3 import org.example.data.Packet;
4 import org.example.network.address.Address;
5 import org.example.util.Util;
6 import org.knowm.xchart.XYChart;
7 import org.knowm.xchart.XYChartBuilder;
8 import org.knowm.xchart.XYSeries;
9 import org.knowm.xchart.style.markers.SeriesMarkers;
10
11 import java.util.ArrayList;
12
13 public class TCPStats extends Stats {
14
15     private final String filename;
16     // CWND
17     private final ArrayList<Integer> congestionWindowCapacities = new ArrayList<>();
18     private final ArrayList<Double> congestionWindowTime = new ArrayList<>();
19     private int numberOfPacketsSent; //total number of packets sent (both normal and
20         retransmissions)
21     private int numberOfPacketsRetransmitted; //total number of packets
22         retransmitted
23     private int numberOfPacketsFastRetransmitted; // total number of packets dropped
24     private int numberOfPacketsArrived; //total number of packets that are enqueued
25     private int numberOfAcksReceived; //total number of ACKs received
26     private int numberOfPacketsReceived; //total number of packets received
27
28     private double goodput;
29     private double lossRate;
30
31     public TCPStats(Address address) {
32         this.filename = address.toString();
33     }
34
35     @Override
36     protected void additionalCalculations() {
37         this.setGoodput();
38         this.setLossRate();
39     }
40
41     @Override
42     protected String fileName() {
43         return this.filename;
44     }
45
46     public void packetSend() {
47         this.numberOfPacketsSent++;
48     }
49
50     public void packetRetransmit() {
51         this.numberOfPacketsRetransmitted++;
52     }
53
54     public void packetFastRetransmit() {
55         this.numberOfPacketsFastRetransmitted++;
56     }
57
58     @Override
59     public void packetArrival(Packet packet) {
60         this.numberOfPacketsArrived++;
61         super.packetArrival(packet);
62     }

```

```

61
62     @Override
63     public void packetDeparture(Packet packet) {
64         this.numberOfPacketsReceived++;
65         super.packetDeparture(packet);
66     }
67
68     public void ackReceived() {
69         this.numberOfAcksReceived++;
70     }
71
72     public int getNumberOfPacketsSent() {
73         return this.numberOfPacketsSent;
74     }
75
76     public int getNumberOfPacketsRetransmitted() {
77         return this.numberOfPacketsRetransmitted;
78     }
79
80     public int getNumberOfPacketsFastRetransmitted() {
81         return this.numberOfPacketsFastRetransmitted;
82     }
83
84     public int getNumberOfAcksReceived() {
85         return this.numberOfAcksReceived;
86     }
87
88     public int getNumberOfPacketsArrived() {
89         return numberOfPacketsArrived;
90     }
91
92     public int getNumberOfPacketsReceived() {
93         return numberOfPacketsReceived;
94     }
95
96     public double getGoodput() {
97         return goodput;
98     }
99
100    public double getLossRate() {
101        return lossRate;
102    }
103
104    private void setGoodput() {
105        if (congestionWindowTime.isEmpty()) return;
106        this.goodput = (double) this.numberOfPacketsSent / (this.
            congestionWindowTime.get(this.congestionWindowTime.size() - 1) / (double
            ) TIMESCALE);
107    }
108
109    private void setLossRate() {
110        if (this.numberOfPacketsSent == 0) {
111            this.lossRate = 0;
112            return;
113        }
114
115        this.lossRate = (this.numberOfPacketsFastRetransmitted + this.
            numberOfPacketsRetransmitted) / (double) this.numberOfPacketsSent;
116    }
117
118    public void trackCwnd(int cwnd) {
119        congestionWindowCapacities.add(cwnd);
120        congestionWindowTime.add((double) Util.seeTime());

```

```

121     }
122
123     public void createCWNDChart() {
124         if (this.congestionWindowTime.isEmpty() || this.congestionWindowCapacities.
125             isEmpty()) return;
126         if (this.congestionWindowTime.size() != this.congestionWindowCapacities.size
127             ())
128             throw new IllegalStateException("the arrays must be of equal length");
129
130         XYChart chart = new XYChartBuilder()
131             .width(CHART_WIDTH)
132             .height(CHART_HEIGHT)
133             .xAxisTitle("Time")
134             .yAxisTitle("CWND Size")
135             .title("Congestion Window Capacity")
136             .theme(theme)
137             .build();
138         chart.getStyler().setChartTitleFont(TITLE_FONT);
139         chart.getStyler().setAxisTitleFont(AXIS_FONT);
140         chart.getStyler().setAxisTickLabelsFont(TICK_FONT);
141         chart.getStyler().setLegendVisible(false);
142         chart.getStyler().setXAxisLabelRotation(45);
143         chart.addSeries("CWND", congestionWindowTime, congestionWindowCapacities).
144             setMarker(SeriesMarkers.NONE);
145         chart.getStyler().setDefaultSeriesRenderStyle(XYSeries.XYSeriesRenderStyle.
146             Area);
147         this.saveChart(chart, "CWNDChart_");
148     }
149 }

```

```

1 package org.example.util;
2
3 import java.util.Collection;
4 import java.util.Comparator;
5 import java.util.Iterator;
6 import java.util.concurrent.PriorityBlockingQueue;
7 import java.util.concurrent.TimeUnit;
8
9 public class BoundedPriorityBlockingQueue<T> implements BoundedQueue<T> {
10
11
12     private PriorityBlockingQueue<T> pbq;
13     private int bound;
14
15     public BoundedPriorityBlockingQueue(int bound, Comparator<T> comparator) {
16         this.pbq = new PriorityBlockingQueue<>(bound, comparator);
17         this.bound = bound;
18     }
19
20     public BoundedPriorityBlockingQueue(int bound) {
21         this.pbq = new PriorityBlockingQueue<>(bound);
22         this.bound = bound;
23     }
24
25     @Override
26     public boolean isFull() {
27         return this.pbq.size() >= this.bound;
28     }
29
30     @Override
31     public boolean offer(T t) {
32         if (isFull()) return false;
33         return this.pbq.offer(t);
34     }
35
36     @Override
37     public void put(T t) throws InterruptedException {
38         if (isEmpty()) return;
39         this.pbq.put(t);
40     }
41
42     @Override
43     public boolean offer(T t, long l, TimeUnit timeUnit) throws InterruptedException
44     {
45         if (isFull()) return false;
46         return this.pbq.offer(t);
47     }
48
49     @Override
50     public T take() throws InterruptedException {
51         return this.pbq.take();
52     }
53
54     @Override
55     public T poll(long l, TimeUnit timeUnit) throws InterruptedException {
56         return this.pbq.poll();
57     }
58
59     @Override
60     public int remainingCapacity() {
61         return this.bound - this.pbq.size();
62     }

```

```

62
63
64     @Override
65     public int bound() {
66         return this.bound;
67     }
68
69     @Override
70     public void setBound(int bound) {
71         this.bound = bound;
72     }
73
74     /**
75      * @deprecated because it may violate the bound constraint of this class.
76      * New implementation that follows the capacity constraint is needed
77      */
78     @Override
79     @Deprecated(since = "4.5", forRemoval = true)
80     public int drainTo(Collection<? super T> collection) {
81         return this.pbq.drainTo(collection);
82     }
83
84     /**
85      * @deprecated because it may violate the bound constraint of this class.
86      * New implementation that follows the capacity constraint is needed
87      */
88     @Override
89     @Deprecated(since = "4.5", forRemoval = true)
90     public int drainTo(Collection<? super T> collection, int i) {
91         return this.pbq.drainTo(collection, i);
92     }
93
94     @Override
95     public T remove() {
96         return this.pbq.remove();
97     }
98
99     @Override
100    public T poll() {
101        return this.pbq.poll();
102    }
103
104    @Override
105    public T element() {
106        return this.pbq.element();
107    }
108
109    @Override
110    public T peek() {
111        return this.pbq.peek();
112    }
113
114    @Override
115    public int size() {
116        return this.pbq.size();
117    }
118
119    @Override
120    public boolean isEmpty() {
121        return this.pbq.isEmpty();
122    }
123
124    @Override

```

```

125     public boolean contains(Object o) {
126         return this.pbq.contains(o);
127     }
128
129     @Override
130     public Object[] toArray() {
131         return this.pbq.toArray();
132     }
133
134     @Override
135     public <E> E[] toArray(E[] ts) {
136         return this.pbq.toArray(ts);
137     }
138
139     @Override
140     public boolean add(T t) {
141         if (isFull()) return false;
142         return this.pbq.add(t);
143     }
144
145     @Override
146     public boolean remove(Object o) {
147         return this.pbq.remove(o);
148     }
149
150     @Override
151     public boolean containsAll(Collection<?> collection) {
152         return this.pbq.containsAll(collection);
153     }
154
155     /**
156      * @deprecated because it may violate the bound constraint of this class.
157      * New implementation that follows the capacity constraint is needed
158      */
159     @Override
160     @Deprecated(since = "4.5", forRemoval = true)
161     public boolean addAll(Collection<? extends T> collection) {
162         return this.pbq.addAll(collection);
163     }
164
165     @Override
166     public boolean removeAll(Collection<?> collection) {
167         return this.pbq.removeAll(collection);
168     }
169
170     /**
171      * @deprecated because it may violate the bound constraint of this class.
172      * New implementation that follows the capacity constraint is needed
173      */
174     @Override
175     @Deprecated(since = "4.5", forRemoval = true)
176     public boolean retainAll(Collection<?> collection) {
177         return this.pbq.retainAll(collection);
178     }
179
180     @Override
181     public void clear() {
182         this.pbq.clear();
183     }
184
185     @Override
186     public Iterator<T> iterator() {
187         return this.pbq.iterator();

```



```
188     }
189 }
```

```
1 package org.example.util;
2
3 import java.util.Collection;
4 import java.util.Queue;
5 import java.util.concurrent.BlockingQueue;
6
7 public interface BoundedQueue<T> extends Iterable<T>, Collection<T>, BlockingQueue<T>, Queue<T> {
8
9     /**
10      * A method to check is the queue is full
11      *
12      * @return True if the queue is full.
13      */
14     boolean isFull();
15
16     /**
17      * A method that returns the bound of the BoundedQueue
18      *
19      * @return the bound
20      */
21     int bound();
22
23     /**
24      * A method to set the bound of the BoundedQueue
25      *
26      * @param bound
27      */
28     void setBound(int bound);
29 }
```

```

1 package org.example.util;
2
3 import org.apache.commons.math3.distribution.PoissonDistribution;
4 import org.apache.commons.math3.random.RandomGeneratorFactory;
5 import org.example.simulator.events.Event;
6
7 import java.util.Random;
8
9 public class Util {
10
11     public static final Random random = new Random(1337);
12     private static long time = 0;
13
14     private Util() {
15         throw new IllegalStateException("Utility class");
16     }
17
18     public static void setSeed(int seed) {
19         random.setSeed(seed);
20     }
21
22     public static double getNextRandomDouble() {
23         return random.nextDouble();
24     }
25
26     public static int getNextRandomInt() {
27         return random.nextInt();
28     }
29
30     public static int getNextRandomInt(int bound) {
31         return random.nextInt(bound);
32     }
33
34     public static double getNextGaussian() {
35         return random.nextGaussian();
36     }
37
38     public static void tickTime(Event event) {
39         time = event.getInstant();
40     }
41
42     public static long getTime() {
43         return time;
44     }
45
46     public static long seeTime() {
47         return time;
48     }
49
50     public static void resetTime() {
51         time = 0;
52     }
53
54     public static PoissonDistribution getPoissonDistribution(double mean) {
55         return new PoissonDistribution(RandomGeneratorFactory.createRandomGenerator(
56             random), mean, PoissonDistribution.DEFAULT_EPSILON, PoissonDistribution.
57             DEFAULT_MAX_ITERATIONS);
58     }
59 }

```

```

1 import org.example.data.Message;
2 import org.example.data.Packet;
3 import org.example.network.Channel;
4 import org.example.network.Routable;
5 import org.example.network.Router;
6 import org.example.network.address.SimpleAddress;
7 import org.example.protocol.ClassicTCP;
8 import org.example.protocol.MPTCP;
9 import org.example.protocol.TCP;
10 import org.example.simulator.EventHandler;
11 import org.example.simulator.events.tcp.TCPConnectEvent;
12 import org.example.simulator.statistics.TCPStats;
13 import org.example.util.Util;
14 import org.junit.Assert;
15 import org.junit.Before;
16 import org.junit.Test;
17
18 public class Simulator {
19
20     private int numPacketsToSend = 2000;
21
22     @Before
23     public void setUp() {
24         Util.resetTime();
25         Util.setSeed(1);
26     }
27
28     private void numPacketsToSend(TCP tcpToSend, int numPacketsToSend) {
29         for (int i = 1; i <= numPacketsToSend; i++) {
30             Message msg = new Message("test " + i);
31             tcpToSend.send(msg);
32         }
33     }
34
35     private EventHandler connectAndRun(TCP client, TCP server) {
36         EventHandler eventHandler = new EventHandler();
37         eventHandler.addEvent(new TCPConnectEvent(client, server));
38         eventHandler.run();
39         return eventHandler;
40     }
41
42     private void allReceived(TCP receiver, int numPacketsToSend) {
43         for (int i = 1; i <= numPacketsToSend; i++) {
44             Message msg = new Message("test " + i);
45             Packet received = receiver.receive();
46             Assert.assertNotNull(received);
47             Assert.assertEquals(msg, received.getPayload());
48         }
49     }
50
51     @Test
52     public void shortPathLowLoss() {
53         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
54             withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
55                 "Client-ShortPathLowLoss")).build();
56         Routable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
57             withAddress(new SimpleAddress("Router 1")).build();
58         Routable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
59             withAddress(new SimpleAddress("Router 2")).build();
60         Routable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
61             withAddress(new SimpleAddress("Router 3")).build();
62         Routable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).

```

```

58         withAddress(new SimpleAddress("Router 4")).build();
ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
            "Server-ShortPathLowLoss")).build();
59
60     new Channel.ChannelBuilder().withCost(1).build(client, r1);
61     new Channel.ChannelBuilder().withCost(10).build(r1, r2);
62     new Channel.ChannelBuilder().withCost(10).withLoss(0.001).build(r2, r3);
63     new Channel.ChannelBuilder().withCost(10).build(r3, r4);
64     new Channel.ChannelBuilder().withCost(1).build(r4, server);
65
66     client.updateRoutingTable();
67     r1.updateRoutingTable();
68     r2.updateRoutingTable();
69     r3.updateRoutingTable();
70     r4.updateRoutingTable();
71     server.updateRoutingTable();
72
73     this.numPacketsToSend(server, this.numPacketsToSend);
74     EventHandler eventHandler = this.connectAndRun(client, server);
75
76     Assert.assertTrue(client.isConnected());
77     Assert.assertTrue(server.isConnected());
78     Assert.assertTrue(client.inputBufferIsEmpty());
79     Assert.assertTrue(server.inputBufferIsEmpty());
80     Assert.assertTrue(client.outputBufferIsEmpty());
81     Assert.assertTrue(server.outputBufferIsEmpty());
82
83     this.allReceived(client, this.numPacketsToSend);
84
85     TCPStats stat = client.getStats();
86     System.out.println(stat);
87     stat.createArrivalChart();
88     stat.createDepartureChart();
89     stat.createInterArrivalChart();
90     stat.createTimeInSystemChart();
91     stat.createNumberOfPacketsInSystemChart();
92
93     System.out.println(server.getStats().toString());
94     server.getStats().createCWNDChart();
95     server.getStats().createDepartureChart();
96 }
97
98 @Test
99 public void shortPathHighLoss() {
100     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
            "Client-ShortPathHighLoss")).build();
101     Routable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
        withAddress(new SimpleAddress("Router 1")).build();
102     Routable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
        withAddress(new SimpleAddress("Router 2")).build();
103     Routable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
        withAddress(new SimpleAddress("Router 3")).build();
104     Routable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
        withAddress(new SimpleAddress("Router 4")).build();
105     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
            "Server-ShortPathHighLoss")).build();
106
107     new Channel.ChannelBuilder().withCost(1).build(client, r1);
108     new Channel.ChannelBuilder().withCost(10).build(r1, r2);
109

```

```

110     new Channel.ChannelBuilder().withCost(10).withLoss(0.01).build(r2, r3);
111     new Channel.ChannelBuilder().withCost(10).build(r3, r4);
112     new Channel.ChannelBuilder().withCost(1).build(r4, server);
113
114     client.updateRoutingTable();
115     r1.updateRoutingTable();
116     r2.updateRoutingTable();
117     r3.updateRoutingTable();
118     r4.updateRoutingTable();
119     server.updateRoutingTable();
120
121     this.numPacketsToSend(server, this.numPacketsToSend);
122     EventHandler eventHandler = this.connectAndRun(client, server);
123
124     Assert.assertTrue(client.isConnected());
125     Assert.assertTrue(server.isConnected());
126     Assert.assertTrue(client.inputBufferIsEmpty());
127     Assert.assertTrue(server.inputBufferIsEmpty());
128     Assert.assertTrue(client.outputBufferIsEmpty());
129     Assert.assertTrue(server.outputBufferIsEmpty());
130
131     this.allReceived(client, this.numPacketsToSend);
132
133     TCPStats stat = client.getStats();
134     System.out.println(stat);
135     stat.createArrivalChart();
136     stat.createDepartureChart();
137     stat.createInterArrivalChart();
138     stat.createTimeInSystemChart();
139     stat.createNumberOfPacketsInSystemChart();
140
141     System.out.println(server.getStats().toString());
142     server.getStats().createCWNDChart();
143     server.getStats().createDepartureChart();
144 }
145
146 @Test
147 public void longPathLowLoss() {
148     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
149         withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
150             "Client-LongPathLowLoss")).build();
151     Rutable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
152         withAddress(new SimpleAddress("Router 1")).build();
153     Rutable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
154         withAddress(new SimpleAddress("Router 2")).build();
155     Rutable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
156         withAddress(new SimpleAddress("Router 3")).build();
157     Rutable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
158         withAddress(new SimpleAddress("Router 4")).build();
159     Rutable r5 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
160         withAddress(new SimpleAddress("Router 5")).build();
161     Rutable r6 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
162         withAddress(new SimpleAddress("Router 6")).build();
163     Rutable r7 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
164         withAddress(new SimpleAddress("Router 7")).build();
165     Rutable r8 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
166         withAddress(new SimpleAddress("Router 8")).build();
167     Rutable r9 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
168         withAddress(new SimpleAddress("Router 9")).build();
169     Rutable r10 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
170         withAddress(new SimpleAddress("Router 10")).build();
171     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
172         withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(

```

```

160         "Server-LongPathLowLoss")).build();
161     new Channel.ChannelBuilder().withCost(1).build(client, r1);
162     new Channel.ChannelBuilder().withCost(10).build(r1, r2);
163     new Channel.ChannelBuilder().withCost(50).build(r2, r3);
164     new Channel.ChannelBuilder().withCost(50).build(r3, r4);
165     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r4, r5);
166     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r5, r6);
167     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r6, r7);
168     new Channel.ChannelBuilder().withCost(50).build(r7, r8);
169     new Channel.ChannelBuilder().withCost(50).build(r8, r9);
170     new Channel.ChannelBuilder().withCost(10).build(r9, r10);
171     new Channel.ChannelBuilder().withCost(1).build(r10, server);
172
173     client.updateRoutingTable();
174     r1.updateRoutingTable();
175     r2.updateRoutingTable();
176     r3.updateRoutingTable();
177     r4.updateRoutingTable();
178     r5.updateRoutingTable();
179     r6.updateRoutingTable();
180     r7.updateRoutingTable();
181     r8.updateRoutingTable();
182     r9.updateRoutingTable();
183     r10.updateRoutingTable();
184     server.updateRoutingTable();
185
186     this.numPacketsToSend(server, this.numPacketsToSend);
187     EventHandler eventHandler = this.connectAndRun(client, server);
188
189     Assert.assertTrue(client.isConnected());
190     Assert.assertTrue(server.isConnected());
191     Assert.assertTrue(client.inputBufferIsEmpty());
192     Assert.assertTrue(server.inputBufferIsEmpty());
193     Assert.assertTrue(client.outputBufferIsEmpty());
194     Assert.assertTrue(server.outputBufferIsEmpty());
195
196     this.allReceived(client, this.numPacketsToSend);
197
198     TCPStats stat = client.getStats();
199     System.out.println(stat);
200     stat.createArrivalChart();
201     stat.createDepartureChart();
202     stat.createInterArrivalChart();
203     stat.createTimeInSystemChart();
204     stat.createNumberOfPacketsInSystemChart();
205
206     System.out.println(server.getStats().toString());
207     server.getStats().createCWNDChart();
208 }
209
210 @Test
211 public void longPathHighLoss() {
212     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
213         withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
214             "Client-LongPathHighLoss")).build();
215     Rutable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
216         withAddress(new SimpleAddress("Router 1")).build();
217     Rutable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
218         withAddress(new SimpleAddress("Router 2")).build();
219     Rutable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
220         withAddress(new SimpleAddress("Router 3")).build();
221     Rutable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).

```

```

217         withAddress(new SimpleAddress("Router 4")).build();
218     Routable r5 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
219         withAddress(new SimpleAddress("Router 5")).build();
220     Routable r6 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
221         withAddress(new SimpleAddress("Router 6")).build();
222     Routable r7 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
223         withAddress(new SimpleAddress("Router 7")).build();
224     Routable r8 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
225         withAddress(new SimpleAddress("Router 8")).build();
226     Routable r9 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
227         withAddress(new SimpleAddress("Router 9")).build();
228     Routable r10 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
229         withAddress(new SimpleAddress("Router 10")).build();
230     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
231         withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
232             "Server-LongPathHighLoss")).build();
233
234     new Channel.ChannelBuilder().withCost(1).build(client, r1);
235     new Channel.ChannelBuilder().withCost(10).build(r1, r2);
236     new Channel.ChannelBuilder().withCost(50).build(r2, r3);
237     new Channel.ChannelBuilder().withCost(50).build(r3, r4);
238     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r4, r5);
239     new Channel.ChannelBuilder().withCost(50).withLoss(0.01).build(r5, r6);
240     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r6, r7);
241     new Channel.ChannelBuilder().withCost(50).build(r7, r8);
242     new Channel.ChannelBuilder().withCost(50).build(r8, r9);
243     new Channel.ChannelBuilder().withCost(10).build(r9, r10);
244     new Channel.ChannelBuilder().withCost(1).build(r10, server);
245
246     client.updateRoutingTable();
247     r1.updateRoutingTable();
248     r2.updateRoutingTable();
249     r3.updateRoutingTable();
250     r4.updateRoutingTable();
251     r5.updateRoutingTable();
252     r6.updateRoutingTable();
253     r7.updateRoutingTable();
254     r8.updateRoutingTable();
255     r9.updateRoutingTable();
256     r10.updateRoutingTable();
257     server.updateRoutingTable();
258
259     this.numPacketsToSend(server, this.numPacketsToSend);
260     EventHandler eventHandler = this.connectAndRun(client, server);
261
262     Assert.assertTrue(client.isConnected());
263     Assert.assertTrue(server.isConnected());
264     Assert.assertTrue(client.inputBufferIsEmpty());
265     Assert.assertTrue(server.inputBufferIsEmpty());
266     Assert.assertTrue(client.outputBufferIsEmpty());
267     Assert.assertTrue(server.outputBufferIsEmpty());
268
269     this.allReceived(client, this.numPacketsToSend);
270
271     TCPStats stat = client.getStats();
272     System.out.println(stat);
273     stat.createArrivalChart();
274     stat.createDepartureChart();
275     stat.createInterArrivalChart();
276     stat.createTimeInSystemChart();
277     stat.createNumberOfPacketsInSystemChart();
278
279     System.out.println(server.getStats().toString());

```

```

271     server.getStats().createCWNDChart();
272 }
273
274
275 private void getStats(MPTCP sender, MPTCP receiver) {
276     //receiver
277     for (TCPStats stat : receiver.getTcpStats()) {
278         System.out.println(stat.toString());
279         stat.createArrivalChart();
280         stat.createDepartureChart();
281         stat.createInterArrivalChart();
282         stat.createTimeInSystemChart();
283         stat.createNumberOfPacketsInSystemChart();
284     }
285
286     //sender
287     for (TCPStats stat : sender.getTcpStats()) {
288         System.out.println(stat.toString());
289         stat.createCWNDChart();
290     }
291 }
292
293 @Test
294 public void MPTCP_HomoShortPathLowLoss() {
295     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
296         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
297             Client-ShortPathLowLoss")).build();
298
299     Routable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
300         withAddress(new SimpleAddress("Router 1")).build();
301     Routable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
302         withAddress(new SimpleAddress("Router 2")).build();
303     Routable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
304         withAddress(new SimpleAddress("Router 3")).build();
305     Routable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
306         withAddress(new SimpleAddress("Router 4")).build();
307
308     Routable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
309         withAddress(new SimpleAddress("Router 1")).build();
310     Routable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
311         withAddress(new SimpleAddress("Router 2")).build();
312     Routable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
313         withAddress(new SimpleAddress("Router 3")).build();
314     Routable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
315         withAddress(new SimpleAddress("Router 4")).build();
316
317     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
318         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
319             Server-ShortPathLowLoss")).build();
320
321     //path one
322     new Channel.ChannelBuilder().withCost(1).build(client, r11);
323     new Channel.ChannelBuilder().withCost(10).build(r11, r12);
324     new Channel.ChannelBuilder().withCost(10).withLoss(0.001).build(r12, r13);
325     new Channel.ChannelBuilder().withCost(10).build(r13, r14);
326     new Channel.ChannelBuilder().withCost(1).build(r14, server);
327
328     //path two
329     new Channel.ChannelBuilder().withCost(1).build(client, r21);
330     new Channel.ChannelBuilder().withCost(10).build(r21, r22);
331     new Channel.ChannelBuilder().withCost(10).withLoss(0.001).build(r22, r23);
332     new Channel.ChannelBuilder().withCost(10).build(r23, r24);
333     new Channel.ChannelBuilder().withCost(1).build(r24, server);

```



```

322     client.updateRoutingTable();
323     r11.updateRoutingTable();
324     r12.updateRoutingTable();
325     r13.updateRoutingTable();
326     r14.updateRoutingTable();
327     r21.updateRoutingTable();
328     r22.updateRoutingTable();
329     r23.updateRoutingTable();
330     r24.updateRoutingTable();
331     server.updateRoutingTable();
332
333     this.numPacketsToSend(server, this.numPacketsToSend);
334     EventHandler eventHandler = this.connectAndRun(server, client);
335
336     Assert.assertTrue("client still has packets to send", client.
337         outputBufferIsEmpty());
338     Assert.assertTrue(server.outputBufferIsEmpty());
339     Assert.assertTrue(client.inputBufferIsEmpty());
340     Assert.assertTrue(server.inputBufferIsEmpty());
341     Assert.assertTrue(r11.inputBufferIsEmpty());
342     Assert.assertTrue(r12.inputBufferIsEmpty());
343     Assert.assertTrue(r13.inputBufferIsEmpty());
344     Assert.assertTrue(r14.inputBufferIsEmpty());
345     Assert.assertTrue(r21.inputBufferIsEmpty());
346     Assert.assertTrue(r22.inputBufferIsEmpty());
347     Assert.assertTrue(r23.inputBufferIsEmpty());
348     Assert.assertTrue(r24.inputBufferIsEmpty());
349
350     this.allReceived(client, this.numPacketsToSend);
351
352     Assert.assertNull(server.receive());
353     this.getStats(server, client);
354 }
355
356
357 @Test
358 public void MPTCP_HomoShortPathHighLoss() {
359     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
360         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
361             Client-ShortPathHighLoss")).build();
362
363     Routable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
364         withAddress(new SimpleAddress("Router 1")).build();
365     Routable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
366         withAddress(new SimpleAddress("Router 2")).build();
367     Routable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
368         withAddress(new SimpleAddress("Router 3")).build();
369     Routable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
370         withAddress(new SimpleAddress("Router 4")).build();
371
372     Routable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
373         withAddress(new SimpleAddress("Router 1")).build();
374     Routable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
375         withAddress(new SimpleAddress("Router 2")).build();
376     Routable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
377         withAddress(new SimpleAddress("Router 3")).build();
378     Routable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
379         withAddress(new SimpleAddress("Router 4")).build();
380
381     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
382         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
383             Server-ShortPathHighLoss")).build();

```

```

372
373 //path one
374 new Channel.ChannelBuilder().withCost(1).build(client, r11);
375 new Channel.ChannelBuilder().withCost(10).build(r11, r12);
376 new Channel.ChannelBuilder().withCost(10).withLoss(0.01).build(r12, r13);
377 new Channel.ChannelBuilder().withCost(10).build(r13, r14);
378 new Channel.ChannelBuilder().withCost(1).build(r14, server);
379
380 //path two
381 new Channel.ChannelBuilder().withCost(1).build(client, r21);
382 new Channel.ChannelBuilder().withCost(10).build(r21, r22);
383 new Channel.ChannelBuilder().withCost(10).withLoss(0.01).build(r22, r23);
384 new Channel.ChannelBuilder().withCost(10).build(r23, r24);
385 new Channel.ChannelBuilder().withCost(1).build(r24, server);
386
387 client.updateRoutingTable();
388 r11.updateRoutingTable();
389 r12.updateRoutingTable();
390 r13.updateRoutingTable();
391 r14.updateRoutingTable();
392 r21.updateRoutingTable();
393 r22.updateRoutingTable();
394 r23.updateRoutingTable();
395 r24.updateRoutingTable();
396 server.updateRoutingTable();
397
398 this.numPacketsToSend(server, this.numPacketsToSend);
399 EventHandler eventHandler = this.connectAndRun(server, client);
400
401 Assert.assertTrue("client still has packets to send", client.
402     outputBufferIsEmpty());
403 Assert.assertTrue(server.outputBufferIsEmpty());
404 Assert.assertTrue(client.inputBufferIsEmpty());
405 Assert.assertTrue(server.inputBufferIsEmpty());
406 Assert.assertTrue(r11.inputBufferIsEmpty());
407 Assert.assertTrue(r12.inputBufferIsEmpty());
408 Assert.assertTrue(r13.inputBufferIsEmpty());
409 Assert.assertTrue(r14.inputBufferIsEmpty());
410 Assert.assertTrue(r21.inputBufferIsEmpty());
411 Assert.assertTrue(r22.inputBufferIsEmpty());
412 Assert.assertTrue(r23.inputBufferIsEmpty());
413 Assert.assertTrue(r24.inputBufferIsEmpty());
414
415 this.allReceived(client, this.numPacketsToSend);
416
417 Assert.assertNull(server.receive());
418 this.getStats(server, client);
419 }
420
421 @Test
422 public void MPTCP_HomoLongPathLowLoss() {
423     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
424         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
425             Client-LongPathLowLoss")).build();
426
427     Routable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
428         withAddress(new SimpleAddress("Router 11")).build();
429     Routable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
430         withAddress(new SimpleAddress("Router 12")).build();
431     Routable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
432         withAddress(new SimpleAddress("Router 13")).build();
433     Routable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).

```

```

    withAddress(new SimpleAddress("Router 14")).build();
429 Rutable r15 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 15")).build();
430 Rutable r16 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 16")).build();
431 Rutable r17 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 17")).build();
432 Rutable r18 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 18")).build();
433 Rutable r19 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 19")).build();
434 Rutable r110 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
    withAddress(new SimpleAddress("Router 110")).build();
435
436 Rutable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
    withAddress(new SimpleAddress("Router 21")).build();
437 Rutable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 22")).build();
438 Rutable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 23")).build();
439 Rutable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 24")).build();
440 Rutable r25 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 25")).build();
441 Rutable r26 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 26")).build();
442 Rutable r27 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 27")).build();
443 Rutable r28 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 28")).build();
444 Rutable r29 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
    withAddress(new SimpleAddress("Router 29")).build();
445 Rutable r210 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
    withAddress(new SimpleAddress("Router 210")).build();
446
447 MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
    withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
    Server-LongPathLowLoss")).build();
448
449 //path one
450 new Channel.ChannelBuilder().withCost(1).build(client, r11);
451 new Channel.ChannelBuilder().withCost(10).build(r11, r12);
452 new Channel.ChannelBuilder().withCost(50).build(r12, r13);
453 new Channel.ChannelBuilder().withCost(50).build(r13, r14);
454 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r14, r15);
455 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r15, r16);
456 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r16, r17);
457 new Channel.ChannelBuilder().withCost(50).build(r17, r18);
458 new Channel.ChannelBuilder().withCost(50).build(r18, r19);
459 new Channel.ChannelBuilder().withCost(10).build(r19, r110);
460 new Channel.ChannelBuilder().withCost(1).build(r110, server);
461
462 //path two
463 new Channel.ChannelBuilder().withCost(1).build(client, r21);
464 new Channel.ChannelBuilder().withCost(10).build(r21, r22);
465 new Channel.ChannelBuilder().withCost(50).build(r22, r23);
466 new Channel.ChannelBuilder().withCost(50).build(r23, r24);
467 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r24, r25);
468 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r25, r26);
469 new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r26, r27);
470 new Channel.ChannelBuilder().withCost(50).build(r27, r28);
471 new Channel.ChannelBuilder().withCost(50).build(r28, r29);
472 new Channel.ChannelBuilder().withCost(10).build(r29, r210);

```

```

473     new Channel.ChannelBuilder().withCost(1).build(r210, server);
474
475     client.updateRoutingTable();
476
477     r11.updateRoutingTable();
478     r12.updateRoutingTable();
479     r13.updateRoutingTable();
480     r14.updateRoutingTable();
481     r15.updateRoutingTable();
482     r16.updateRoutingTable();
483     r17.updateRoutingTable();
484     r18.updateRoutingTable();
485     r19.updateRoutingTable();
486     r110.updateRoutingTable();
487
488     r21.updateRoutingTable();
489     r22.updateRoutingTable();
490     r23.updateRoutingTable();
491     r24.updateRoutingTable();
492     r25.updateRoutingTable();
493     r26.updateRoutingTable();
494     r27.updateRoutingTable();
495     r28.updateRoutingTable();
496     r29.updateRoutingTable();
497     r210.updateRoutingTable();
498
499     server.updateRoutingTable();
500
501     this.numPacketsToSend(server, this.numPacketsToSend);
502     EventHandler eventHandler = this.connectAndRun(server, client);
503
504     Assert.assertTrue("client still has packets to send", client.
505         outputBufferIsEmpty());
506     Assert.assertTrue(server.outputBufferIsEmpty());
507     Assert.assertTrue(client.inputBufferIsEmpty());
508     Assert.assertTrue(server.inputBufferIsEmpty());
509     Assert.assertTrue(r11.inputBufferIsEmpty());
510     Assert.assertTrue(r12.inputBufferIsEmpty());
511     Assert.assertTrue(r13.inputBufferIsEmpty());
512     Assert.assertTrue(r14.inputBufferIsEmpty());
513     Assert.assertTrue(r15.inputBufferIsEmpty());
514     Assert.assertTrue(r16.inputBufferIsEmpty());
515     Assert.assertTrue(r17.inputBufferIsEmpty());
516     Assert.assertTrue(r18.inputBufferIsEmpty());
517     Assert.assertTrue(r19.inputBufferIsEmpty());
518     Assert.assertTrue(r110.inputBufferIsEmpty());
519
520     Assert.assertTrue(r21.inputBufferIsEmpty());
521     Assert.assertTrue(r22.inputBufferIsEmpty());
522     Assert.assertTrue(r23.inputBufferIsEmpty());
523     Assert.assertTrue(r24.inputBufferIsEmpty());
524     Assert.assertTrue(r25.inputBufferIsEmpty());
525     Assert.assertTrue(r26.inputBufferIsEmpty());
526     Assert.assertTrue(r27.inputBufferIsEmpty());
527     Assert.assertTrue(r28.inputBufferIsEmpty());
528     Assert.assertTrue(r29.inputBufferIsEmpty());
529     Assert.assertTrue(r210.inputBufferIsEmpty());
530
531     this.allReceived(client, this.numPacketsToSend);
532
533     Assert.assertNull(server.receive());
534     this.getStats(server, client);

```

```

535     }
536
537     @Test
538     public void MPTCP_HomoLongPathHighLoss() {
539         MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
540             withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
541                 Client-LongPathHighLoss")).build();
542
543         Rutable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
544             withAddress(new SimpleAddress("Router 11")).build();
545         Rutable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
546             withAddress(new SimpleAddress("Router 12")).build();
547         Rutable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
548             withAddress(new SimpleAddress("Router 13")).build();
549         Rutable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
550             withAddress(new SimpleAddress("Router 14")).build();
551         Rutable r15 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
552             withAddress(new SimpleAddress("Router 15")).build();
553         Rutable r16 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
554             withAddress(new SimpleAddress("Router 16")).build();
555         Rutable r17 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
556             withAddress(new SimpleAddress("Router 17")).build();
557         Rutable r18 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
558             withAddress(new SimpleAddress("Router 18")).build();
559         Rutable r19 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
560             withAddress(new SimpleAddress("Router 19")).build();
561         Rutable r110 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
562             withAddress(new SimpleAddress("Router 110")).build();
563
564         Rutable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
565             withAddress(new SimpleAddress("Router 21")).build();
566         Rutable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
567             withAddress(new SimpleAddress("Router 22")).build();
568         Rutable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
569             withAddress(new SimpleAddress("Router 23")).build();
570         Rutable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
571             withAddress(new SimpleAddress("Router 24")).build();
572         Rutable r25 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
573             withAddress(new SimpleAddress("Router 25")).build();
574         Rutable r26 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
575             withAddress(new SimpleAddress("Router 26")).build();
576         Rutable r27 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
577             withAddress(new SimpleAddress("Router 27")).build();
578         Rutable r28 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
579             withAddress(new SimpleAddress("Router 28")).build();
580         Rutable r29 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
581             withAddress(new SimpleAddress("Router 29")).build();
582         Rutable r210 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
583             withAddress(new SimpleAddress("Router 210")).build();
584
585         MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
586             withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
587                 Server-LongPathHighLoss")).build();
588
589         //path one
590         new Channel.ChannelBuilder().withCost(1).build(client, r11);
591         new Channel.ChannelBuilder().withCost(10).build(r11, r12);
592         new Channel.ChannelBuilder().withCost(50).build(r12, r13);
593         new Channel.ChannelBuilder().withCost(50).build(r13, r14);
594         new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r14, r15);
595         new Channel.ChannelBuilder().withCost(50).withLoss(0.01).build(r15, r16);
596         new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r16, r17);
597         new Channel.ChannelBuilder().withCost(50).build(r17, r18);

```

```

574     new Channel.ChannelBuilder().withCost(50).build(r18, r19);
575     new Channel.ChannelBuilder().withCost(10).build(r19, r110);
576     new Channel.ChannelBuilder().withCost(1).build(r110, server);
577
578     //path two
579     new Channel.ChannelBuilder().withCost(1).build(client, r21);
580     new Channel.ChannelBuilder().withCost(10).build(r21, r22);
581     new Channel.ChannelBuilder().withCost(50).build(r22, r23);
582     new Channel.ChannelBuilder().withCost(50).build(r23, r24);
583     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r24, r25);
584     new Channel.ChannelBuilder().withCost(50).withLoss(0.01).build(r25, r26);
585     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r26, r27);
586     new Channel.ChannelBuilder().withCost(50).build(r27, r28);
587     new Channel.ChannelBuilder().withCost(50).build(r28, r29);
588     new Channel.ChannelBuilder().withCost(10).build(r29, r210);
589     new Channel.ChannelBuilder().withCost(1).build(r210, server);
590
591     client.updateRoutingTable();
592
593     r11.updateRoutingTable();
594     r12.updateRoutingTable();
595     r13.updateRoutingTable();
596     r14.updateRoutingTable();
597     r15.updateRoutingTable();
598     r16.updateRoutingTable();
599     r17.updateRoutingTable();
600     r18.updateRoutingTable();
601     r19.updateRoutingTable();
602     r110.updateRoutingTable();
603
604     r21.updateRoutingTable();
605     r22.updateRoutingTable();
606     r23.updateRoutingTable();
607     r24.updateRoutingTable();
608     r25.updateRoutingTable();
609     r26.updateRoutingTable();
610     r27.updateRoutingTable();
611     r28.updateRoutingTable();
612     r29.updateRoutingTable();
613     r210.updateRoutingTable();
614
615     server.updateRoutingTable();
616
617     this.numPacketsToSend(server, this.numPacketsToSend);
618     EventHandler eventHandler = this.connectAndRun(server, client);
619
620     Assert.assertTrue("client still has packets to send", client.
        outputBufferIsEmpty());
621     Assert.assertTrue(server.outputBufferIsEmpty());
622     Assert.assertTrue(client.inputBufferIsEmpty());
623     Assert.assertTrue(server.inputBufferIsEmpty());
624     Assert.assertTrue(r11.inputBufferIsEmpty());
625     Assert.assertTrue(r12.inputBufferIsEmpty());
626     Assert.assertTrue(r13.inputBufferIsEmpty());
627     Assert.assertTrue(r14.inputBufferIsEmpty());
628     Assert.assertTrue(r15.inputBufferIsEmpty());
629     Assert.assertTrue(r16.inputBufferIsEmpty());
630     Assert.assertTrue(r17.inputBufferIsEmpty());
631     Assert.assertTrue(r18.inputBufferIsEmpty());
632     Assert.assertTrue(r19.inputBufferIsEmpty());
633     Assert.assertTrue(r110.inputBufferIsEmpty());
634
635     Assert.assertTrue(r21.inputBufferIsEmpty());

```

```

636     Assert.assertTrue(r22.inputBufferIsEmpty());
637     Assert.assertTrue(r23.inputBufferIsEmpty());
638     Assert.assertTrue(r24.inputBufferIsEmpty());
639     Assert.assertTrue(r25.inputBufferIsEmpty());
640     Assert.assertTrue(r26.inputBufferIsEmpty());
641     Assert.assertTrue(r27.inputBufferIsEmpty());
642     Assert.assertTrue(r28.inputBufferIsEmpty());
643     Assert.assertTrue(r29.inputBufferIsEmpty());
644     Assert.assertTrue(r210.inputBufferIsEmpty());
645
646     this.allReceived(client, this.numPacketsToSend);
647
648     Assert.assertNull(server.receive());
649     this.getStats(server, client);
650
651 }
652
653 @Test
654 public void MPTCP_HeteroLowLoss() {
655     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
656         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
657             Client-HeteroLowLoss")).build();
658
659     Routable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
660         withAddress(new SimpleAddress("Router 11")).build();
661     Routable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
662         withAddress(new SimpleAddress("Router 12")).build();
663     Routable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
664         withAddress(new SimpleAddress("Router 13")).build();
665     Routable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
666         withAddress(new SimpleAddress("Router 14")).build();
667     Routable r15 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
668         withAddress(new SimpleAddress("Router 15")).build();
669     Routable r16 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
670         withAddress(new SimpleAddress("Router 16")).build();
671     Routable r17 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
672         withAddress(new SimpleAddress("Router 17")).build();
673     Routable r18 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
674         withAddress(new SimpleAddress("Router 18")).build();
675     Routable r19 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
676         withAddress(new SimpleAddress("Router 19")).build();
677     Routable r110 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
678         withAddress(new SimpleAddress("Router 110")).build();
679
680     Routable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
681         withAddress(new SimpleAddress("Router 1")).build();
682     Routable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
683         withAddress(new SimpleAddress("Router 2")).build();
684     Routable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
685         withAddress(new SimpleAddress("Router 3")).build();
686     Routable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
687         withAddress(new SimpleAddress("Router 4")).build();
688
689     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
690         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
691             Server-HeteroLowLoss")).build();
692
693     // long path
694     new Channel.ChannelBuilder().withCost(1).build(client, r11);
695     new Channel.ChannelBuilder().withCost(10).build(r11, r12);
696     new Channel.ChannelBuilder().withCost(50).build(r12, r13);
697     new Channel.ChannelBuilder().withCost(50).build(r13, r14);
698     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r14, r15);

```

```

681     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r15, r16);
682     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r16, r17);
683     new Channel.ChannelBuilder().withCost(50).build(r17, r18);
684     new Channel.ChannelBuilder().withCost(50).build(r18, r19);
685     new Channel.ChannelBuilder().withCost(10).build(r19, r110);
686     new Channel.ChannelBuilder().withCost(1).build(r110, server);
687
688     // short path
689     new Channel.ChannelBuilder().withCost(1).build(client, r21);
690     new Channel.ChannelBuilder().withCost(10).build(r21, r22);
691     new Channel.ChannelBuilder().withCost(10).withLoss(0.001).build(r22, r23);
692     new Channel.ChannelBuilder().withCost(10).build(r23, r24);
693     new Channel.ChannelBuilder().withCost(1).build(r24, server);
694
695     client.updateRoutingTable();
696
697     r11.updateRoutingTable();
698     r12.updateRoutingTable();
699     r13.updateRoutingTable();
700     r14.updateRoutingTable();
701     r15.updateRoutingTable();
702     r16.updateRoutingTable();
703     r17.updateRoutingTable();
704     r18.updateRoutingTable();
705     r19.updateRoutingTable();
706     r110.updateRoutingTable();
707
708     r21.updateRoutingTable();
709     r22.updateRoutingTable();
710     r23.updateRoutingTable();
711     r24.updateRoutingTable();
712
713     server.updateRoutingTable();
714
715     this.numPacketsToSend(server, this.numPacketsToSend);
716     EventHandler eventHandler = this.connectAndRun(server, client);
717
718
719     Assert.assertTrue("client still has packets to send", client.
720         outputBufferIsEmpty());
721     Assert.assertTrue(server.outputBufferIsEmpty());
722     Assert.assertTrue(client.inputBufferIsEmpty());
723     Assert.assertTrue(server.inputBufferIsEmpty());
724     Assert.assertTrue(r11.inputBufferIsEmpty());
725     Assert.assertTrue(r12.inputBufferIsEmpty());
726     Assert.assertTrue(r13.inputBufferIsEmpty());
727     Assert.assertTrue(r14.inputBufferIsEmpty());
728     Assert.assertTrue(r15.inputBufferIsEmpty());
729     Assert.assertTrue(r16.inputBufferIsEmpty());
730     Assert.assertTrue(r17.inputBufferIsEmpty());
731     Assert.assertTrue(r18.inputBufferIsEmpty());
732     Assert.assertTrue(r19.inputBufferIsEmpty());
733     Assert.assertTrue(r110.inputBufferIsEmpty());
734
735     Assert.assertTrue(r21.inputBufferIsEmpty());
736     Assert.assertTrue(r22.inputBufferIsEmpty());
737     Assert.assertTrue(r23.inputBufferIsEmpty());
738     Assert.assertTrue(r24.inputBufferIsEmpty());
739
740     this.allReceived(client, this.numPacketsToSend);
741
742     Assert.assertNull(server.receive());
743     this.getStats(server, client);

```



```

743 }
744 }
745
746 @Test
747 public void MPTCP_HeteroHighLoss () {
748     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
749         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
750             Client-HeteroHighLoss")).build();
751
752     Rutable r11 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
753         withAddress(new SimpleAddress("Router 11")).build();
754     Rutable r12 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
755         withAddress(new SimpleAddress("Router 12")).build();
756     Rutable r13 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
757         withAddress(new SimpleAddress("Router 13")).build();
758     Rutable r14 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
759         withAddress(new SimpleAddress("Router 14")).build();
760     Rutable r15 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
761         withAddress(new SimpleAddress("Router 15")).build();
762     Rutable r16 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
763         withAddress(new SimpleAddress("Router 16")).build();
764     Rutable r17 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
765         withAddress(new SimpleAddress("Router 17")).build();
766     Rutable r18 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
767         withAddress(new SimpleAddress("Router 18")).build();
768     Rutable r19 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
769         withAddress(new SimpleAddress("Router 19")).build();
770     Rutable r110 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
771         withAddress(new SimpleAddress("Router 110")).build();
772
773     Rutable r21 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
774         withAddress(new SimpleAddress("Router 1")).build();
775     Rutable r22 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
776         withAddress(new SimpleAddress("Router 2")).build();
777     Rutable r23 = new Router.RouterBuilder().withAverageQueueUtilization(0.98).
778         withAddress(new SimpleAddress("Router 3")).build();
779     Rutable r24 = new Router.RouterBuilder().withAverageQueueUtilization(0.5).
780         withAddress(new SimpleAddress("Router 4")).build();
781
782     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
783         withReceivingWindowCapacity(30).withAddress(new SimpleAddress("MPTCP-
784             Server-HeteroHighLoss")).build();
785
786     // long path
787     new Channel.ChannelBuilder().withCost(1).build(client, r11);
788     new Channel.ChannelBuilder().withCost(10).build(r11, r12);
789     new Channel.ChannelBuilder().withCost(50).build(r12, r13);
790     new Channel.ChannelBuilder().withCost(50).build(r13, r14);
791     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r14, r15);
792     new Channel.ChannelBuilder().withCost(50).withLoss(0.01).build(r15, r16);
793     new Channel.ChannelBuilder().withCost(50).withLoss(0.001).build(r16, r17);
794     new Channel.ChannelBuilder().withCost(50).build(r17, r18);
795     new Channel.ChannelBuilder().withCost(50).build(r18, r19);
796     new Channel.ChannelBuilder().withCost(10).build(r19, r110);
797     new Channel.ChannelBuilder().withCost(1).build(r110, server);
798
799     // short path
800     new Channel.ChannelBuilder().withCost(1).build(client, r21);
801     new Channel.ChannelBuilder().withCost(10).build(r21, r22);
802     new Channel.ChannelBuilder().withCost(10).withLoss(0.01).build(r22, r23);
803     new Channel.ChannelBuilder().withCost(10).build(r23, r24);
804     new Channel.ChannelBuilder().withCost(1).build(r24, server);
805
806 }

```

```

788     client.updateRoutingTable();
789
790     r11.updateRoutingTable();
791     r12.updateRoutingTable();
792     r13.updateRoutingTable();
793     r14.updateRoutingTable();
794     r15.updateRoutingTable();
795     r16.updateRoutingTable();
796     r17.updateRoutingTable();
797     r18.updateRoutingTable();
798     r19.updateRoutingTable();
799     r110.updateRoutingTable();
800
801     r21.updateRoutingTable();
802     r22.updateRoutingTable();
803     r23.updateRoutingTable();
804     r24.updateRoutingTable();
805
806     server.updateRoutingTable();
807
808     this.numPacketsToSend(server, this.numPacketsToSend);
809     EventHandler eventHandler = this.connectAndRun(server, client);
810
811
812     Assert.assertTrue("client still has packets to send", client.
        outputBufferIsEmpty());
813     Assert.assertTrue(server.outputBufferIsEmpty());
814     Assert.assertTrue(client.inputBufferIsEmpty());
815     Assert.assertTrue(server.inputBufferIsEmpty());
816     Assert.assertTrue(r11.inputBufferIsEmpty());
817     Assert.assertTrue(r12.inputBufferIsEmpty());
818     Assert.assertTrue(r13.inputBufferIsEmpty());
819     Assert.assertTrue(r14.inputBufferIsEmpty());
820     Assert.assertTrue(r15.inputBufferIsEmpty());
821     Assert.assertTrue(r16.inputBufferIsEmpty());
822     Assert.assertTrue(r17.inputBufferIsEmpty());
823     Assert.assertTrue(r18.inputBufferIsEmpty());
824     Assert.assertTrue(r19.inputBufferIsEmpty());
825     Assert.assertTrue(r110.inputBufferIsEmpty());
826
827     Assert.assertTrue(r21.inputBufferIsEmpty());
828     Assert.assertTrue(r22.inputBufferIsEmpty());
829     Assert.assertTrue(r23.inputBufferIsEmpty());
830     Assert.assertTrue(r24.inputBufferIsEmpty());
831
832     this.allReceived(client, this.numPacketsToSend);
833
834     Assert.assertNull(server.receive());
835     this.getStats(server, client);
836
837 }
838
839 }

```

```

1 package org.example.data;
2
3 import org.junit.Assert;
4 import org.junit.Test;
5
6 public class MessageTest {
7
8     @Test
9     public void createMessageWorksTest() {
10         Message message = new Message("test");
11         Assert.assertTrue(message instanceof Message);
12     }
13
14     @Test
15     public void stringInMessageIsCorrectTest() {
16         String s1 = "test";
17         Message message1 = new Message(s1);
18         Assert.assertEquals(s1, message1.toString());
19
20         String s2 = "test12345";
21         Message message2 = new Message(s2);
22         Assert.assertEquals(s2, message2.toString());
23     }
24
25     @Test
26     public void sizeOfMessageIsCorrectTest() {
27         String s1 = "test";
28         Message message1 = new Message(s1);
29         Assert.assertEquals(s1.length(), message1.size());
30
31         String s2 = "test12345";
32         Message message2 = new Message(s2);
33         Assert.assertEquals(s2.length(), message2.size());
34     }
35
36
37     @Test
38     public void messagesWithEqualStringsAreEqualTest() {
39         String s = "test";
40         Message message1 = new Message(s);
41         Message message2 = new Message(s);
42         Assert.assertEquals(message1, message2);
43     }
44 }

```

```

1 package org.example.data;
2
3 import org.example.network.interfaces.Endpoint;
4 import org.example.protocol.ClassicTCP;
5 import org.junit.Assert;
6 import org.junit.Test;
7
8 import java.util.List;
9
10 public class PacketTest {
11
12     @Test
13     public void buildPacketBuildsPacketTest() {
14         Assert.assertTrue(new PacketBuilder().build() instanceof Packet);
15     }
16
17     @Test
18     public void buildPacketWithDestinationBuildsPacketWithDestinationTest() {
19         Endpoint endpoint = new ClassicTCP.ClassicTCPBuilder().
20             withReceivingWindowCapacity(7).build();
21         Packet packet = new PacketBuilder().withDestination(endpoint).build();
22         Assert.assertEquals(endpoint, packet.getDestination());
23     }
24
25     @Test
26     public void buildPacketWithOriginBuildsPacketWithOriginTest() {
27         Endpoint endpoint = new ClassicTCP.ClassicTCPBuilder().
28             withReceivingWindowCapacity(7).build();
29         Packet packet = new PacketBuilder().withOrigin(endpoint).build();
30         Assert.assertEquals(endpoint, packet.getOrigin());
31     }
32
33     @Test
34     public void buildPacketWithMsgBuildsPacketMsgOriginTest() {
35         Message msg = new Message("Test");
36         Packet packet = new PacketBuilder().withPayload(msg).build();
37         Assert.assertEquals(msg, packet.getPayload());
38     }
39
40     @Test
41     public void buildPacketWithSeqNumBuildsPacketSeqNumOriginTest() {
42         int seqNum = 123;
43         Packet packet = new PacketBuilder().withSequenceNumber(seqNum).build();
44         Assert.assertEquals(seqNum, packet.getSequenceNumber());
45     }
46
47     @Test
48     public void buildPacketWithAckNumBuildsPacketAckNumOriginTest() {
49         int ackNum = 321;
50         Packet packet = new PacketBuilder().withAcknowledgmentNumber(ackNum).
51             withFlags(Flag.ACK).build();
52         Assert.assertEquals(ackNum, packet.getAcknowledgmentNumber());
53     }
54
55     @Test
56     public void buildPacketWithFlagBuildsPacketThatHasFlagTest() {
57         Flag flag = Flag.ACK;
58         Packet packet = new PacketBuilder().withFlags(flag).build();
59         Assert.assertTrue(packet.hasAllFlags(flag));
60     }
61
62     @Test

```

```

60     public void buildPacketWithFlagsBuildsPacketThatHasFlagsTest () {
61         Flag ack = Flag.ACK;
62         Flag syn = Flag.SYN;
63         Flag fin = Flag.FIN;
64         Packet packet = new PacketBuilder().withFlags(ack, syn, fin).build();
65         Assert.assertTrue(packet.hasAllFlags(syn, fin));
66         Assert.assertTrue(packet.hasAllFlags(ack));
67     }
68
69     @Test
70     public void buildPacketWithoutFlagBuildsPacketThatDoesNotHaveThatFlagTest () {
71         Flag ack = Flag.ACK;
72         Flag syn = Flag.SYN;
73         Flag fin = Flag.FIN;
74         Packet packet = new PacketBuilder().withFlags(syn, fin).build();
75         Assert.assertFalse(packet.hasAllFlags(syn, fin, ack));
76         Assert.assertFalse(packet.hasAllFlags(ack));
77         Assert.assertFalse(packet.hasAllFlags(syn, ack));
78         Assert.assertTrue(packet.hasAllFlags(syn, fin));
79     }
80
81     @Test
82     public void buildPacketWithDuplicateFlagsBuildsPacketThatHasOneOfThatFlagTest ()
83     {
84         Flag ack = Flag.ACK;
85         Flag syn = Flag.SYN;
86         Flag fin = Flag.FIN;
87         Packet packet = new PacketBuilder().withFlags(ack, syn, syn, fin, fin).build
88         ();
89         List<Flag> flags = packet.getFlags();
90         Assert.assertEquals(3, flags.size());
91         Assert.assertTrue(packet.hasAllFlags(syn, fin, ack));
92     }
93
94     @Test
95     public void buildPacketWithPayloadHasCorrectSizeTest () {
96         Message message = new Message("test");
97         Packet packet = new PacketBuilder()
98             .withPayload(message)
99             .build();
100         Assert.assertEquals(message.size(), packet.size());
101     }
102
103     @Test
104     public void buildPacketWithoutPayloadHasZeroSizeTest () {
105         Packet packet = new PacketBuilder()
106             .build();
107         Assert.assertEquals(0, packet.size());
108     }

```

```

1 package org.example.network;
2
3 import org.example.network.address.Address;
4 import org.example.network.address.UUIDAddress;
5 import org.junit.Assert;
6 import org.junit.Test;
7
8 public class AddressTest {
9
10     @Test
11     public void createAddressWorksTest() {
12         Address address = new UUIDAddress();
13         Assert.assertTrue(address instanceof Address);
14     }
15
16     @Test
17     public void twoUUIDAddressesAreNotEqual() {
18         Address address1 = new UUIDAddress();
19         Address address2 = new UUIDAddress();
20         Assert.assertNotEquals(address1, address2);
21     }
22
23     @Test
24     public void nullAddressesAreNotEqualToAddress() {
25         Address address1 = new UUIDAddress();
26         Address address2 = null;
27         Assert.assertNotEquals(address1, address2);
28     }
29
30 }

```

```

1 package org.example.network;
2
3 import org.example.data.*;
4 import org.example.network.interfaces.Endpoint;
5 import org.example.protocol.ClassicTCP;
6 import org.junit.Assert;
7 import org.junit.Test;
8
9 import java.util.PriorityQueue;
10
11 public class ChannelTest {
12
13     private Endpoint source = new ClassicTCP.ClassicTCPBuilder().
14         withReceivingWindowCapacity(7).build();
15     private Endpoint destination = new ClassicTCP.ClassicTCPBuilder().
16         withReceivingWindowCapacity(7).build();
17
18     @Test
19     public void channelConstructor1Test() {
20         Channel channel = new Channel(source, destination, 100);
21         Assert.assertTrue(channel instanceof Channel);
22         Assert.assertEquals(source, channel.getSource());
23         Assert.assertEquals(destination, channel.getDestination());
24     }
25
26     @Test
27     public void channelConstructor2Test() {
28         Channel channel = new Channel(source);
29         Assert.assertTrue(channel instanceof Channel);
30         Assert.assertEquals(source, channel.getSource());
31         Assert.assertEquals(source, channel.getDestination());
32         Assert.assertEquals(0, channel.getCost());
33     }
34
35     @Test
36     public void channelWithConstructor1CostIsPositiveTest() {
37         for (int i = 0; i < 10000; i++) {
38             Channel channel = new Channel(source, destination, 100);
39             Assert.assertTrue(channel.getCost() >= 0);
40         }
41     }
42
43     @Test
44     public void channelWithConstructor2CostIsZeroTest() {
45         for (int i = 0; i < 10000; i++) {
46             Channel channel = new Channel(source);
47             Assert.assertEquals(0, channel.getCost());
48         }
49     }
50
51     @Test
52     public void channelPacketDeliversPacketToDestinationNodeTest() {
53         Channel channel = new Channel(source, destination, 0);
54         Payload payload = new Message("Test");
55         Packet packet = new PacketBuilder()
56             .withOrigin(source)
57             .withDestination(destination)
58             .withPayload(payload)
59             .withFlags(Flag.SYN) // hack to overcome connection check in the
60                 endpoints
61             .build();
62         channel.channelPackage(packet);

```

```

60     Assert.assertTrue(channel.channel());
61     Packet receivedPacket = destination.dequeueInputBuffer();
62     Assert.assertEquals(packet, receivedPacket);
63 }
64
65
66 @Test
67 public void lossyChannelDropPacketTest() {
68     Channel channel = new Channel(source, destination, 0);
69     Payload payload = new Message("Test");
70     Packet packet = new PacketBuilder()
71         .withOrigin(source)
72         .withDestination(destination)
73         .withPayload(payload)
74         .withFlags(Flag.SYN) // hack to overcome connection check in the
75             endpoints
76         .build();
77     channel.channelPackage(packet);
78     Packet receivedPacket = destination.dequeueInputBuffer();
79     Assert.assertEquals(null, receivedPacket);
80 }
81
82 @Test
83 public void aLittleLossyChannelWillDropPacketAfterAWhileTest() {
84     for (int i = 0; i < 10000; i++) {
85         Channel channel = new Channel(source, destination, 3);
86         Payload payload = new Message("Test");
87         Packet packet = new PacketBuilder()
88             .withOrigin(source)
89             .withDestination(destination)
90             .withPayload(payload)
91             .withFlags(Flag.SYN) // hack to overcome connection check in the
92                 endpoints
93             .build();
94         channel.channelPackage(packet);
95         Packet receivedPacket = destination.dequeueInputBuffer();
96         if (receivedPacket == null) {
97             Assert.assertTrue(true);
98             return;
99         }
100     }
101     Assert.assertFalse(true);
102 }
103
104 @Test
105 public void channelCompareToWorksInPriorityQueueTest() {
106     int size = 1000;
107     PriorityQueue<Channel> pq = new PriorityQueue<>(size);
108     for (int i = 0; i < size; i++) {
109         Channel c = new Channel(source, destination, 0);
110         pq.add(c);
111     }
112     while (pq.size() > 1) {
113         Assert.assertTrue(pq.poll().getCost() <= pq.peek().getCost());
114     }
115 }
116
117 @Test
118 public void differentChannelsAreNotEqualTest() {
119     Endpoint source1 = new ClassicTCP.ClassicTCPBuilder().
120         withReceivingWindowCapacity(7).build();
121     Endpoint destination1 = new ClassicTCP.ClassicTCPBuilder().

```



```
120         withReceivingWindowCapacity(7).build();
121     Endpoint source2 = new ClassicTCP.ClassicTCPBuilder().
122         withReceivingWindowCapacity(7).build();
123     Endpoint destination2 = new ClassicTCP.ClassicTCPBuilder().
124         withReceivingWindowCapacity(7).build();
125     Channel channel1 = new Channel(source1, destination1, 100);
126     Channel channel2 = new Channel(source2, destination2, 100);
127     Assert.assertNotEquals(channel1, channel2);
128 }
129
130
131 }
```

```

1 package org.example.network;
2
3 import org.example.data.*;
4 import org.example.network.interfaces.Endpoint;
5 import org.example.network.interfaces.NetworkNode;
6 import org.example.simulator.EventHandler;
7 import org.example.simulator.events.RouteEvent;
8 import org.junit.Assert;
9 import org.junit.Test;
10
11 import java.util.concurrent.ArrayBlockingQueue;
12
13 public class RutableTest {
14
15     @Test
16     public void routableWithRandomAddressAreNotEqualTest () {
17         NetworkNode node1 = new Router.RouterBuilder().build();
18         NetworkNode node2 = new Router.RouterBuilder().build();
19         Assert.assertNotEquals(node1, node2);
20     }
21
22
23     @Test
24     public void routingPacketRoutsItToItsDestinationStraitLine () {
25         Endpoint client = new RutableEndpoint(new ArrayBlockingQueue<>(100), new
26             ArrayBlockingQueue<>(100));
27         Router r1 = new Router.RouterBuilder().build();
28         Router r2 = new Router.RouterBuilder().build();
29         Router r3 = new Router.RouterBuilder().build();
30         Router r4 = new Router.RouterBuilder().build();
31         Endpoint server = new RutableEndpoint(new ArrayBlockingQueue<>(100), new
32             ArrayBlockingQueue<>(100));
33
34         new Channel.ChannelBuilder().build(client, r1);
35         new Channel.ChannelBuilder().build(r1, r2);
36         new Channel.ChannelBuilder().build(r2, r3);
37         new Channel.ChannelBuilder().build(r3, r4);
38         new Channel.ChannelBuilder().build(r4, server);
39
40         client.updateRoutingTable();
41         r1.updateRoutingTable();
42         r2.updateRoutingTable();
43         r3.updateRoutingTable();
44         r4.updateRoutingTable();
45         server.updateRoutingTable();
46
47         Message msg = new Message("Test");
48         Packet packet = new PacketBuilder()
49             .withPayload(msg)
50             .withDestination(server)
51             .build();
52
53         EventHandler eventHandler = new EventHandler();
54         eventHandler.addEvent(new RouteEvent(client, packet));
55         eventHandler.run();
56
57         Packet receivedPacket = ((RutableEndpoint) server).getReceivedPacket();
58         Payload receivedPayload = receivedPacket.getPayload();
59
60         Assert.assertEquals(receivedPayload, msg);
61     }
62 }

```

```

61
62     @Test
63     public void routingPacketRoutsItToItsDestinationCircleGraph() {
64         Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
65             ArrayBlockingQueue<>(100));
66         Router r1 = new Router.RouterBuilder().build();
67         Router r2 = new Router.RouterBuilder().build();
68         Router r3 = new Router.RouterBuilder().build();
69         Router r4 = new Router.RouterBuilder().build();
70         Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
71             ArrayBlockingQueue<>(100));
72
73         new Channel.ChannelBuilder().build(client, r1);
74         new Channel.ChannelBuilder().build(client, r2);
75         new Channel.ChannelBuilder().build(r1, r3);
76         new Channel.ChannelBuilder().build(r2, r4);
77         new Channel.ChannelBuilder().build(server, r3);
78         new Channel.ChannelBuilder().build(server, r4);
79
80         client.updateRoutingTable();
81         r1.updateRoutingTable();
82         r2.updateRoutingTable();
83         r3.updateRoutingTable();
84         r4.updateRoutingTable();
85         server.updateRoutingTable();
86
87         Message msg = new Message("Test");
88         Packet packet = new PacketBuilder()
89             .withPayload(msg)
90             .withDestination(server)
91             .withFlags(Flag.SYN) // hack to overcome connection check in the
92                 endpoints
93             .build();
94
95         EventHandler eventHandler = new EventHandler();
96         eventHandler.addEvent(new RouteEvent(client, packet));
97         eventHandler.run();
98
99         Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
100        Payload receivedPayload = receivedPacket.getPayload();
101
102        Assert.assertEquals(receivedPayload, msg);
103    }
104
105     @Test
106     public void routingPacketRoutsItToItsDestinationWithCycle() {
107         Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
108             ArrayBlockingQueue<>(100));
109         Router r1 = new Router.RouterBuilder().build();
110         Router r2 = new Router.RouterBuilder().build();
111         Router r3 = new Router.RouterBuilder().build();
112         Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
113             ArrayBlockingQueue<>(100));
114
115         new Channel.ChannelBuilder().build(client, r1);
116         new Channel.ChannelBuilder().build(r1, r2);
117         new Channel.ChannelBuilder().build(r1, r3);
118         new Channel.ChannelBuilder().build(r2, r3);
119         new Channel.ChannelBuilder().build(server, r3);
120
121         client.updateRoutingTable();

```

```

119     server.updateRoutingTable();
120     r1.updateRoutingTable();
121     r2.updateRoutingTable();
122     r3.updateRoutingTable();
123
124
125     Message msg = new Message("Test");
126     Packet packet = new PacketBuilder()
127         .withPayload(msg)
128         .withDestination(server)
129         .build();
130
131     EventHandler eventHandler = new EventHandler();
132     eventHandler.addEvent(new RouteEvent(client, packet));
133     eventHandler.run();
134
135     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
136     Payload receivedPayload = receivedPacket.getPayload();
137
138     Assert.assertEquals(receivedPayload, msg);
139 }
140
141 @Test
142 public void routingPacketRoutsItToItsDestinationWithDeadEnd() {
143     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
144         ArrayBlockingQueue<>(100));
145     Router r1 = new Router.RouterBuilder().build();
146     Router r2 = new Router.RouterBuilder().build();
147     Router r3 = new Router.RouterBuilder().build();
148     Router r4 = new Router.RouterBuilder().build();
149     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
150         ArrayBlockingQueue<>(100));
151
152     new Channel.ChannelBuilder().build(client, r1);
153     new Channel.ChannelBuilder().build(r1, r2);
154     new Channel.ChannelBuilder().build(r1, r4);
155     new Channel.ChannelBuilder().build(r1, r3);
156     new Channel.ChannelBuilder().build(server, r4);
157
158     client.updateRoutingTable();
159     r1.updateRoutingTable();
160     r2.updateRoutingTable();
161     r3.updateRoutingTable();
162     r4.updateRoutingTable();
163     server.updateRoutingTable();
164
165     Message msg = new Message("Test");
166     Packet packet = new PacketBuilder()
167         .withPayload(msg)
168         .withDestination(server)
169         .withFlags(Flag.SYN) // hack to overcome connection check in the
170             endpoints
171         .build();
172
173     EventHandler eventHandler = new EventHandler();
174     eventHandler.addEvent(new RouteEvent(client, packet));
175     eventHandler.run();
176
177     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
178     Payload receivedPayload = receivedPacket.getPayload();
179
180     Assert.assertEquals(msg, receivedPayload);
181 }

```

```

179
180
181 @Test
182 public void routingPacketRoutsItToItsDestinationForrest() {
183     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
184         ArrayBlockingQueue<>(100));
185     Router r1 = new Router.RouterBuilder().build();
186     Router r2 = new Router.RouterBuilder().build();
187     Router r3 = new Router.RouterBuilder().build();
188     Router r4 = new Router.RouterBuilder().build();
189     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
190         ArrayBlockingQueue<>(100));
191
192     //tree one
193     new Channel.ChannelBuilder().build(client, r1);
194     new Channel.ChannelBuilder().build(r1, server);
195
196     //tree two
197     new Channel.ChannelBuilder().build(r2, r3);
198     new Channel.ChannelBuilder().build(r3, r4);
199
200     client.updateRoutingTable();
201     r1.updateRoutingTable();
202     r2.updateRoutingTable();
203     r3.updateRoutingTable();
204     r4.updateRoutingTable();
205     server.updateRoutingTable();
206
207     Message msg = new Message("Test");
208     Packet packet = new PacketBuilder()
209         .withPayload(msg)
210         .withDestination(server)
211         .build();
212
213     EventHandler eventHandler = new EventHandler();
214     eventHandler.addEvent(new RouteEvent(client, packet));
215     eventHandler.run();
216
217     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
218     Payload receivedPayload = receivedPacket.getPayload();
219
220     Assert.assertEquals(msg, receivedPayload);
221 }
222
223 @Test
224 public void routingPacketRoutsItToItsDestinationWithUnconnectedNode() {
225     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
226         ArrayBlockingQueue<>(100));
227     Router r1 = new Router.RouterBuilder().build();
228     Router r2 = new Router.RouterBuilder().build();
229     Router r3 = new Router.RouterBuilder().build();
230     Router r4 = new Router.RouterBuilder().build();
231     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
232         ArrayBlockingQueue<>(100));
233
234     new Channel.ChannelBuilder().build(client, r1);
235     new Channel.ChannelBuilder().build(r1, r2);
236     new Channel.ChannelBuilder().build(r2, r3);
237     new Channel.ChannelBuilder().build(server, r3);
238
239     client.updateRoutingTable();

```

```

238     r1.updateRoutingTable();
239     r2.updateRoutingTable();
240     r3.updateRoutingTable();
241     r4.updateRoutingTable();
242     server.updateRoutingTable();
243
244     Message msg = new Message("Test");
245     Packet packet = new PacketBuilder()
246         .withPayload(msg)
247         .withDestination(server)
248         .build();
249
250     EventHandler eventHandler = new EventHandler();
251     eventHandler.addEvent(new RouteEvent(client, packet));
252     eventHandler.run();
253
254     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
255     Payload receivedPayload = receivedPacket.getPayload();
256
257     Assert.assertEquals(msg, receivedPayload);
258 }
259
260
261 @Test
262 public void routingPacketRoutsItToItsDestinationCrazyGraph() {
263     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
264         ArrayBlockingQueue<>(100));
265     Router r1 = new Router.RouterBuilder().build();
266     Router r2 = new Router.RouterBuilder().build();
267     Router r3 = new Router.RouterBuilder().build();
268     Router r4 = new Router.RouterBuilder().build();
269     Router r5 = new Router.RouterBuilder().build();
270     Router r6 = new Router.RouterBuilder().build();
271     Router r7 = new Router.RouterBuilder().build();
272     Router r8 = new Router.RouterBuilder().build();
273     Router r9 = new Router.RouterBuilder().build();
274     Router r10 = new Router.RouterBuilder().build();
275     Router r11 = new Router.RouterBuilder().build();
276     Router r12 = new Router.RouterBuilder().build();
277     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
278         ArrayBlockingQueue<>(100));
279
280     new Channel.ChannelBuilder().build(client, r1);
281     new Channel.ChannelBuilder().build(client, r2);
282     new Channel.ChannelBuilder().build(r1, r3);
283     new Channel.ChannelBuilder().build(r2, r3);
284     new Channel.ChannelBuilder().build(r3, r4);
285     new Channel.ChannelBuilder().build(r3, r9);
286     new Channel.ChannelBuilder().build(r4, r5);
287     new Channel.ChannelBuilder().build(r4, r6);
288     new Channel.ChannelBuilder().build(r5, r6);
289     new Channel.ChannelBuilder().build(r6, r9);
290     new Channel.ChannelBuilder().build(r6, r7);
291     new Channel.ChannelBuilder().build(r7, r8);
292     new Channel.ChannelBuilder().build(r9, r10);
293     new Channel.ChannelBuilder().build(r10, r11);
294     new Channel.ChannelBuilder().build(r11, r12);
295     new Channel.ChannelBuilder().build(r12, server);
296
297     client.updateRoutingTable();
298     r1.updateRoutingTable();
299     r2.updateRoutingTable();
300     r3.updateRoutingTable();

```

```

299     r4.updateRoutingTable();
300     r5.updateRoutingTable();
301     r6.updateRoutingTable();
302     r7.updateRoutingTable();
303     r8.updateRoutingTable();
304     r9.updateRoutingTable();
305     r10.updateRoutingTable();
306     r11.updateRoutingTable();
307     r12.updateRoutingTable();
308     server.updateRoutingTable();
309
310     Message msg = new Message("Test");
311     Packet packet = new PacketBuilder()
312         .withPayload(msg)
313         .withDestination(server)
314         .build();
315
316     EventHandler eventHandler = new EventHandler();
317     eventHandler.addEvent(new RouteEvent(client, packet));
318     eventHandler.run();
319
320     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
321     Payload receivedPayload = receivedPacket.getPayload();
322
323     Assert.assertEquals(msg, receivedPayload);
324 }
325
326
327 @Test(expected = IllegalArgumentException.class)
328 public void unconnectedClientCantRoutePacketToDestination() {
329     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
330         ArrayBlockingQueue<>(100));
331     Router r1 = new Router.RouterBuilder().build();
332     Router r2 = new Router.RouterBuilder().build();
333     Router r3 = new Router.RouterBuilder().build();
334     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
335         ArrayBlockingQueue<>(100));
336
337     new Channel.ChannelBuilder().build(r1, r2);
338     new Channel.ChannelBuilder().build(r2, r3);
339     new Channel.ChannelBuilder().build(server, r3);
340
341     client.updateRoutingTable();
342     r1.updateRoutingTable();
343     r2.updateRoutingTable();
344     r3.updateRoutingTable();
345     server.updateRoutingTable();
346
347     Message msg = new Message("Test");
348     client.route(new PacketBuilder()
349         .withPayload(msg)
350         .withDestination(server)
351         .build());
352 }
353
354 @Test(expected = IllegalArgumentException.class)
355 public void unconnectedTreesCantRoutePacket() {
356     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
357         ArrayBlockingQueue<>(100));
358     Router r1 = new Router.RouterBuilder().build();
359     Router r2 = new Router.RouterBuilder().build();
360     Router r3 = new Router.RouterBuilder().build();

```

```

359 Router r4 = new Router.RouterBuilder().build();
360 Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
    ArrayBlockingQueue<>(100));
361
362 new Channel.ChannelBuilder().build(client, r1);
363 new Channel.ChannelBuilder().build(r1, r2);
364
365 new Channel.ChannelBuilder().build(r3, r4);
366 new Channel.ChannelBuilder().build(server, r4);
367
368 client.updateRoutingTable();
369 r1.updateRoutingTable();
370 r2.updateRoutingTable();
371 r3.updateRoutingTable();
372 server.updateRoutingTable();
373
374 Message msg = new Message("Test");
375 client.route(new PacketBuilder()
376     .withPayload(msg)
377     .withDestination(server)
378     .build());
379 }
380
381
382 @Test
383 public void faultyChannelsDropPacket() {
384     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
        ArrayBlockingQueue<>(100));
385     Router r1 = new Router.RouterBuilder().build();
386     Router r2 = new Router.RouterBuilder().build();
387     Router r3 = new Router.RouterBuilder().build();
388     Router r4 = new Router.RouterBuilder().build();
389     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(100), new
        ArrayBlockingQueue<>(100));
390
391     new Channel.ChannelBuilder().build(client, r1);
392     new Channel.ChannelBuilder().withLoss(100).build(r1, r2);
393     new Channel.ChannelBuilder().build(r2, r3);
394     new Channel.ChannelBuilder().build(r3, r4);
395     new Channel.ChannelBuilder().build(server, r4);
396
397     client.updateRoutingTable();
398     r1.updateRoutingTable();
399     r2.updateRoutingTable();
400     r3.updateRoutingTable();
401     r4.updateRoutingTable();
402     server.updateRoutingTable();
403
404
405     Message msg = new Message("Test");
406     Packet packet = new PacketBuilder()
407         .withPayload(msg)
408         .withDestination(server)
409         .build();
410
411     EventHandler eventHandler = new EventHandler();
412     eventHandler.addEvent(new RouteEvent(client, packet));
413     eventHandler.run();
414
415     Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
416
417     Assert.assertNull(receivedPacket);
418 }

```



```

419
420
421 @Test
422 public void not100PercentLossyRoutersAreLoosingPacketIfEnoughPacketsAreSent() {
423     double noiseTolerance = 2.5;
424     Endpoint client = new RoutableEndpoint(new ArrayBlockingQueue<>(1000), new
         ArrayBlockingQueue<>(1000));
425     Router r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.1).
         build();
426     Router r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.1).
         build();
427     Router r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.1).
         build();
428     Router r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.1).
         build();
429     Endpoint server = new RoutableEndpoint(new ArrayBlockingQueue<>(1000), new
         ArrayBlockingQueue<>(1000));
430
431     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(client, r1);
432     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(r1, r2);
433     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(r2, r3);
434     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(r3, r4);
435     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(server, r4);
436
437     client.updateRoutingTable();
438     r1.updateRoutingTable();
439     r2.updateRoutingTable();
440     r3.updateRoutingTable();
441     r4.updateRoutingTable();
442     server.updateRoutingTable();
443
444
445     EventHandler eventHandler = new EventHandler();
446
447     for (int i = 0; i < 1000; i++) {
448         Message msg = new Message(i + "");
449         Packet packet = new PacketBuilder()
450             .withPayload(msg)
451             .withDestination(server)
452             .build();
453         eventHandler.addEvent(new RouteEvent(client, packet));
454     }
455
456     eventHandler.run();
457
458     for (int i = 0; i < 1000; i++) {
459         Packet receivedPacket = ((RoutableEndpoint) server).getReceivedPacket();
460         if (receivedPacket == null) {
461             return;
462         }
463     }
464     Assert.fail();
465
466 }
467
468
469 }

```

```

1 package org.example.network;
2
3 import org.example.data.PacketBuilder;
4 import org.example.network.interfaces.Endpoint;
5 import org.example.network.interfaces.NetworkNode;
6 import org.example.protocol.ClassicTCP;
7 import org.junit.Assert;
8 import org.junit.Test;
9
10 public class RoutingTableTest {
11
12     @Test(expected = IllegalStateException.class)
13     public void routingTableThrowsIllegalStateExceptionIfNotUpdated() {
14         Endpoint r1 = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
15             (7).build();
16         NetworkNode r2 = new Router.RouterBuilder().build();
17         NetworkNode r3 = new Router.RouterBuilder().build();
18         Endpoint r4 = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
19             (7).build();
20
21         new Channel.ChannelBuilder().build(r1, r2);
22         new Channel.ChannelBuilder().build(r2, r3);
23         new Channel.ChannelBuilder().build(r3, r4);
24
25         r1.route(new PacketBuilder().withDestination(r4).build());
26     }
27
28     @Test(expected = IllegalArgumentException.class)
29     public void routingTableThrowsIllegalArgumentExceptionIfDestinationIsNull() {
30         NetworkNode r1 = new Router.RouterBuilder().build();
31         NetworkNode r2 = new Router.RouterBuilder().build();
32         NetworkNode r3 = new Router.RouterBuilder().build();
33         NetworkNode r4 = new Router.RouterBuilder().build();
34
35         new Channel.ChannelBuilder().build(r1, r2);
36         new Channel.ChannelBuilder().build(r2, r3);
37         new Channel.ChannelBuilder().build(r3, r4);
38
39         r1.updateRoutingTable();
40         r2.updateRoutingTable();
41         r3.updateRoutingTable();
42         r4.updateRoutingTable();
43
44         r1.route(new PacketBuilder().withDestination(null).build());
45     }
46
47     @Test(expected = IllegalArgumentException.class)
48     public void routingTableThrowsIllegalArgumentExceptionIfDestinationIsUnconnected
49         () {
50         Endpoint r1 = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
51             (7).build();
52         NetworkNode r2 = new Router.RouterBuilder().build();
53         NetworkNode r3 = new Router.RouterBuilder().build();
54         Endpoint r4 = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
55             (7).build();
56
57         new Channel.ChannelBuilder().build(r1, r2);
58         new Channel.ChannelBuilder().build(r2, r3);
59
60         r1.updateRoutingTable();
61         r2.updateRoutingTable();

```

```

58     r3.updateRoutingTable();
59     r4.updateRoutingTable();
60
61     r1.route(new PacketBuilder().withDestination(r4).build());
62 }
63
64
65 private Channel getChannelFromTo(NetworkNode from, NetworkNode destination) {
66     for (Channel c : from.getChannels()) {
67         if (c.getDestination().equals(destination)) return c;
68     }
69     Assert.assertFalse(true);
70     return null;
71 }
72
73
74 @Test
75 public void getPathChoosesShortestPathTest() {
76     for (int i = 0; i < 100; i++) {
77         Endpoint r1 = new ClassicTCP.ClassicTCPBuilder().
78             withReceivingWindowCapacity(7).build();
79         NetworkNode r2 = new Router.RouterBuilder().build();
80         NetworkNode r3 = new Router.RouterBuilder().build();
81         Endpoint r4 = new ClassicTCP.ClassicTCPBuilder().
82             withReceivingWindowCapacity(7).build();
83
84         new Channel.ChannelBuilder().build(r1, r2);
85         new Channel.ChannelBuilder().build(r1, r3);
86         new Channel.ChannelBuilder().build(r2, r4);
87         new Channel.ChannelBuilder().build(r3, r4);
88
89         r1.updateRoutingTable();
90         r2.updateRoutingTable();
91         r3.updateRoutingTable();
92         r4.updateRoutingTable();
93
94         Channel channelOnePathOne = getChannelFromTo(r1, r2);
95         Channel channelTwoPathOne = getChannelFromTo(r2, r4);
96
97         Channel channelOnePathTwo = getChannelFromTo(r1, r3);
98         Channel channelTwoPathTwo = getChannelFromTo(r3, r4);
99
100        int pathOne = channelOnePathOne.getCost() + channelTwoPathOne.getCost();
101        int pathTwo = channelOnePathTwo.getCost() + channelTwoPathTwo.getCost();
102        if (pathOne == pathTwo) continue;
103
104        Channel usedChannel = pathOne < pathTwo ? channelOnePathOne :
105            channelOnePathTwo;
106        Channel notUsedChannel = pathOne > pathTwo ? channelOnePathOne :
107            channelOnePathTwo;
108
109        r1.route(new PacketBuilder().withOrigin(r1).withDestination(r4).build())
110            ;
111
112        Assert.assertTrue(usedChannel.channel());
113        Assert.assertFalse(notUsedChannel.channel());
114
115        Assert.assertTrue(!usedChannel.getDestination().inputBufferIsEmpty());
116        Assert.assertNull(notUsedChannel.getDestination().dequeueInputBuffer());
117    }
118 }

```

```
116     @Test
117     public void toStringTest() {
118         RoutingTable routingTable = new RoutingTable();
119         routingTable.update(new Router.RouterBuilder().build());
120         routingTable.toString();
121         Assert.assertTrue(routingTable.toString() instanceof String);
122     }
123
124     @Test(expected = IllegalArgumentException.class)
125     public void getPathWhenNoPathIsThereTest() {
126         RoutingTable routingTable = new RoutingTable();
127         routingTable.update(new Router.RouterBuilder().build());
128         routingTable.getPath(new Router.RouterBuilder().build(), new Router.
129             RouterBuilder().build());
130     }
131 }
```

```

1 package org.example.protocol;
2
3 import org.example.data.Message;
4 import org.example.data.Packet;
5 import org.example.data.PacketBuilder;
6 import org.example.network.Channel;
7 import org.example.network.Routable;
8 import org.example.network.Router;
9 import org.example.protocol.window.receiving.ReceivingWindow;
10 import org.example.simulator.EventHandler;
11 import org.example.simulator.events.RouteEvent;
12 import org.example.simulator.events.tcp.RunTCPEvent;
13 import org.example.simulator.events.tcp.TCPConnectEvent;
14 import org.example.util.Util;
15 import org.junit.Assert;
16 import org.junit.Before;
17 import org.junit.Rule;
18 import org.junit.Test;
19 import org.junit.rules.Timeout;
20
21 import java.util.concurrent.TimeUnit;
22
23 public class ClassicTCPTest {
24
25     @Rule
26     public Timeout globalTimeout = new Timeout(60, TimeUnit.SECONDS);
27
28
29     @Before
30     public void setup() {
31         Util.setSeed(1996);
32         Util.resetTime();
33     }
34
35     @Test
36     public void connectToEndpointTest() {
37         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().build();
38         Routable router = new Router.RouterBuilder().build();
39         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().build();
40
41         new Channel.ChannelBuilder().build(client, router);
42         new Channel.ChannelBuilder().build(router, server);
43
44         client.updateRoutingTable();
45         router.updateRoutingTable();
46         server.updateRoutingTable();
47
48         EventHandler eventHandler = new EventHandler();
49         eventHandler.addEvent(new TCPConnectEvent(client, server));
50         eventHandler.run();
51
52         Assert.assertEquals(server, client.getConnection().getConnectedNode());
53         Assert.assertEquals(client, server.getConnection().getConnectedNode());
54
55         Assert.assertEquals(server.getConnection().getNextSequenceNumber(), client.
56             getConnection().getNextAcknowledgementNumber());
57     }
58
59     @Test
60     public void connectToEndpointShouldResultInCorrectReceivingAndSendingWindowCapacityTest
61     () {

```

```

60     int clientReceivingWindow = 3;
61     int serverReceivingWindow = 7;
62     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(clientReceivingWindow).build();
63     Routable router = new Router.RouterBuilder().build();
64     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(serverReceivingWindow).build();
65
66     new Channel.ChannelBuilder().build(client, router);
67     new Channel.ChannelBuilder().build(router, server);
68
69     client.updateRoutingTable();
70     router.updateRoutingTable();
71     server.updateRoutingTable();
72
73     EventHandler eventHandler = new EventHandler();
74     eventHandler.addEvent(new TCPConnectEvent(client, server));
75     eventHandler.run();
76
77     Assert.assertEquals(server, client.getConnection().getConnectedNode());
78     Assert.assertEquals(client, server.getConnection().getConnectedNode());
79
80     Assert.assertEquals(server.getConnection().getNextSequenceNumber(), client.
        getConnection().getNextAcknowledgementNumber());
81
82     Assert.assertEquals(clientReceivingWindow, client.
        getThisReceivingWindowCapacity());
83     Assert.assertEquals(serverReceivingWindow, server.
        getThisReceivingWindowCapacity());
84
85     Assert.assertEquals(serverReceivingWindow, client.
        getOtherReceivingWindowCapacity());
86     Assert.assertEquals(clientReceivingWindow, server.
        getOtherReceivingWindowCapacity());
87 }
88
89 @Test
90 public void connectThenSendMsgWorksTest() {
91     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(7).build();
92     Routable router = new Router.RouterBuilder().build();
93     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
        withReceivingWindowCapacity(7).build();
94
95     new Channel.ChannelBuilder().build(client, router);
96     new Channel.ChannelBuilder().build(router, server);
97
98     client.updateRoutingTable();
99     router.updateRoutingTable();
100    server.updateRoutingTable();
101
102    EventHandler eventHandler = new EventHandler();
103    eventHandler.addEvent(new TCPConnectEvent(client, server));
104    eventHandler.run();
105
106    Message msg = new Message("Test");
107
108    client.send(msg);
109    eventHandler.addEvent(new RunTCPEvent(client));
110    eventHandler.run();
111
112    Assert.assertEquals(msg, server.receive().getPayload());
113

```

```

114     }
115
116
117     @Test
118     public void connectThenSendMsgOverMultipleNodesLineWorksTest () {
119         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
120             withReceivingWindowCapacity(7).build();
121         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
122             withReceivingWindowCapacity(7).build();
123         Router r1 = new Router.RouterBuilder().build();
124         Router r2 = new Router.RouterBuilder().build();
125         Router r3 = new Router.RouterBuilder().build();
126         Router r4 = new Router.RouterBuilder().build();
127
128         new Channel.ChannelBuilder().build(client, r1);
129         new Channel.ChannelBuilder().build(r1, r2);
130         new Channel.ChannelBuilder().build(r2, r3);
131         new Channel.ChannelBuilder().build(r3, r4);
132         new Channel.ChannelBuilder().build(r4, server);
133
134         client.updateRoutingTable();
135         r1.updateRoutingTable();
136         r2.updateRoutingTable();
137         r3.updateRoutingTable();
138         r4.updateRoutingTable();
139         server.updateRoutingTable();
140
141         EventHandler eventHandler = new EventHandler();
142         eventHandler.addEvent(new TCPConnectEvent(client, server));
143         eventHandler.run();
144
145         Message msg = new Message("Test");
146
147         client.send(msg);
148         eventHandler.addEvent(new RunTCPEvent(client));
149         eventHandler.run();
150
151         Assert.assertEquals(msg, server.receive().getPayload());
152     }
153
154     @Test
155     public void unorderedPacketsAreNotReceivedTest () {
156         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
157             withReceivingWindowCapacity(7).build();
158         Routable router = new Router.RouterBuilder().build();
159         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
160             withReceivingWindowCapacity(7).build();
161
162         new Channel.ChannelBuilder().build(client, router);
163         new Channel.ChannelBuilder().build(router, server);
164
165         client.updateRoutingTable();
166         router.updateRoutingTable();
167         server.updateRoutingTable();
168
169         EventHandler eventHandler = new EventHandler();
170         eventHandler.addEvent(new TCPConnectEvent(client, server));
171         eventHandler.run();
172
173         System.out.println("connected");
174
175         Message msg = new Message("test1");

```

```

173     client.send(msg);
174     eventHandler.addEvent(new RunTCPEvent(client));
175
176     Message msg2 = new Message("test2");
177     Packet packet = new PacketBuilder()
178         .withPayload(msg2)
179         .withOrigin(client)
180         .withDestination(server)
181         .withSequenceNumber(client.getConnection().getNextSequenceNumber() +
182             100)
183         .build();
184
185     eventHandler.addEvent(new RouteEvent(client, packet));
186
187     eventHandler.run();
188
189     Assert.assertEquals(msg, server.receive().getPayload());
190     Assert.assertNull(server.receive());
191 }
192
193 @Test
194 public void
195     unorderedPacketsAreDroppedAndOrderedPacketsAreReceivedWithoutBlockTest() {
196     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
197         withReceivingWindowCapacity(7).build();
198     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
199         withReceivingWindowCapacity(7).build();
200     Router r1 = new Router.RouterBuilder().build();
201     Router r2 = new Router.RouterBuilder().build();
202     Router r3 = new Router.RouterBuilder().build();
203     Router r4 = new Router.RouterBuilder().build();
204
205     new Channel.ChannelBuilder().build(client, r1);
206     new Channel.ChannelBuilder().build(r1, r2);
207     new Channel.ChannelBuilder().build(r2, r3);
208     new Channel.ChannelBuilder().build(r3, r4);
209     new Channel.ChannelBuilder().build(r4, server);
210
211     client.updateRoutingTable();
212     r1.updateRoutingTable();
213     r2.updateRoutingTable();
214     r3.updateRoutingTable();
215     r4.updateRoutingTable();
216     server.updateRoutingTable();
217
218     EventHandler eventHandler = new EventHandler();
219     eventHandler.addEvent(new TCPConnectEvent(client, server));
220     eventHandler.run();
221
222     for (int i = 0; i < server.getThisReceivingWindowCapacity(); i++) {
223         Packet packet1 = new PacketBuilder()
224             .withSequenceNumber(client.getConnection().getNextSequenceNumber()
225                 + 5000 + i)
226             .withAcknowledgmentNumber(client.getConnection().
227                 getNextAcknowledgmentNumber() + 20000 + i)
228             .withOrigin(client)
229             .withDestination(server)
230             .withPayload(new Message("should not be received 1"))
231             .withIndex(i + 50000)
232             .build();
233         eventHandler.addEvent(new RouteEvent(client, packet1));

```



```

230
231     Message msg = new Message("test " + i);
232     client.send(msg);
233     eventHandler.addEvent(new RunTCPEvent(client));
234
235     Packet packet2 = new PacketBuilder()
236         .withSequenceNumber(client.getConnection().getNextSequenceNumber
237             () + 100 + i)
238         .withAcknowledgmentNumber(client.getConnection().
239             getNextAcknowledgmentNumber() + 20 + i)
240         .withOrigin(client)
241         .withDestination(server)
242         .withPayload(new Message("should not be received 2"))
243         .withIndex(i + 100)
244         .build();
245     eventHandler.addEvent(new RouteEvent(client, packet2));
246 }
247 eventHandler.run();
248
249 for (int i = 0; i < server.getThisReceivingWindowCapacity(); i++) {
250     Packet received = server.receive();
251     Assert.assertNotNull(received);
252     Assert.assertEquals("test " + i, received.getPayload().toString());
253 }
254
255 Packet received = server.receive();
256 Assert.assertNull(received);
257 }
258
259 @Test
260 public void routedMessagesUnorderedReceiveOrderedTest() throws
261     IllegalAccessException {
262     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
263         withReceivingWindowCapacity(7).build();
264     Routable router = new Router.RouterBuilder().build();
265     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
266         withReceivingWindowCapacity(7).build();
267
268     new Channel.ChannelBuilder().build(client, router);
269     new Channel.ChannelBuilder().build(router, server);
270
271     client.updateRoutingTable();
272     router.updateRoutingTable();
273     server.updateRoutingTable();
274
275     EventHandler eventHandler = new EventHandler();
276     eventHandler.addEvent(new TCPConnectEvent(client, server));
277     eventHandler.run();
278
279     System.out.println("connected");
280
281     int seqNum = client.getConnection().getNextSequenceNumber();
282     int ackNum = client.getConnection().getNextAcknowledgmentNumber();
283
284     //increase window capacity to max
285     for (int i = 0; i < client.getThisReceivingWindowCapacity(); i++) {
286         client.getSendingWindow().increase();
287     }
288
289     for (int i = client.getThisReceivingWindowCapacity() - 1; i >= 0; i--) {
290         Packet packet = new PacketBuilder()
291             .withOrigin(client)

```

```

288         .withDestination(server)
289         .withSequenceNumber(seqNum + i)
290         .withAcknowledgmentNumber(ackNum + i)
291         .withPayload(new Message(i + ""))
292         .withIndex(i)
293         .build();
294     eventHandler.addEvent(new RouteEvent(client, packet));
295 }
296 eventHandler.run();
297
298 for (int i = 0; i < client.getThisReceivingWindowCapacity(); i++) {
299     Packet received = server.receive();
300     Assert.assertNotNull("iteration: " + i, received);
301     Assert.assertEquals(i + "", received.getPayload().toString());
302 }
303 Packet received = server.receive();
304 Assert.assertNull(received);
305 }
306
307
308 @Test
309 public void routeToManyMessagesUnorderedReceiveOrderedAndDropCorrectTest() {
310     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
311         withReceivingWindowCapacity(7).build();
312     Routable router = new Router.RouterBuilder().build();
313     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
314         withReceivingWindowCapacity(7).build();
315
316     new Channel.ChannelBuilder().build(client, router);
317     new Channel.ChannelBuilder().build(router, server);
318
319     client.updateRoutingTable();
320     router.updateRoutingTable();
321     server.updateRoutingTable();
322
323     EventHandler eventHandler = new EventHandler();
324     eventHandler.addEvent(new TCPConnectEvent(client, server));
325     eventHandler.run();
326
327     System.out.println("connected");
328
329     int seqNum = client.getConnection().getNextSequenceNumber();
330     int ackNum = client.getConnection().getNextAcknowledgementNumber();
331
332     for (int i = client.getThisReceivingWindowCapacity() * 2; i >= 0; i--) {
333         eventHandler.addEvent(new RouteEvent(client, new PacketBuilder()
334             .withOrigin(client)
335             .withDestination(server)
336             .withSequenceNumber(seqNum + i)
337             .withAcknowledgmentNumber(ackNum + i)
338             .withPayload(new Message(i + ""))
339             .withIndex(i)
340             .build()
341         ));
342         eventHandler.run();
343     }
344
345     for (int i = 0; i < client.getThisReceivingWindowCapacity(); i++) {
346         Packet received = server.receive();
347         Assert.assertNotNull(received);
348         Assert.assertEquals(i + "", received.getPayload().toString());
349     }
350     Packet received = server.receive();

```

```

349     Assert.assertNull(received);
350 }
351
352 @Test
353 public void packetIndexShouldUpdateAfterReceivingPacketInOrderTest() throws
354     IllegalAccessException {
355     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
356         withReceivingWindowCapacity(7).build();
357     Routable router = new Router.RouterBuilder().build();
358     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
359         withReceivingWindowCapacity(7).build();
360
361     new Channel.ChannelBuilder().build(client, router);
362     new Channel.ChannelBuilder().build(router, server);
363
364     client.updateRoutingTable();
365     router.updateRoutingTable();
366     server.updateRoutingTable();
367
368     EventHandler eventHandler = new EventHandler();
369     eventHandler.addEvent(new TCPConnectEvent(client, server));
370     eventHandler.run();
371
372     System.out.println("connected");
373
374     Packet packet = new PacketBuilder()
375         .withConnection(client.getConnection())
376         .withSequenceNumber(client.getConnection().getNextSequenceNumber())
377         .build();
378
379     int indexBeforeSending = server.getReceivingWindow().receivingPacketIndex(
380         packet, server.getConnection());
381     Assert.assertEquals(0, indexBeforeSending);
382
383     client.send(packet.getPayload());
384     eventHandler.addEvent(new RunTCPEvent(client));
385     eventHandler.run();
386     Assert.assertNotNull(server.receive());
387
388     int indexAfterReceived = server.getReceivingWindow().receivingPacketIndex(
389         packet, server.getConnection());
390     Assert.assertEquals(-1, indexAfterReceived);
391 }
392
393 @Test
394 public void
395     packetIndexShouldNotUpdateAfterReceivingPacketOutOfOrderButInWindowTest()
396     throws IllegalAccessException {
397     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
398         withReceivingWindowCapacity(7).build();
399     Routable router = new Router.RouterBuilder().build();
400     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
401         withReceivingWindowCapacity(7).build();
402
403     new Channel.ChannelBuilder().build(client, router);
404     new Channel.ChannelBuilder().build(router, server);
405
406     client.updateRoutingTable();
407     router.updateRoutingTable();
408     server.updateRoutingTable();

```

```

403
404     EventHandler eventHandler = new EventHandler();
405     eventHandler.addEvent(new TCPConnectEvent(client, server));
406     eventHandler.run();
407
408     System.out.println("connected");
409
410     //this test does not make sens for window size = 1
411     if (client.getThisReceivingWindowCapacity() <= 1) return;
412
413     int seqNum = client.getConnection().getNextSequenceNumber();
414     int ackNum = client.getConnection().getNextAcknowledgementNumber();
415
416
417     Packet packet = new PacketBuilder()
418         .withConnection(client.getConnection())
419         .withSequenceNumber(seqNum + client.getThisReceivingWindowCapacity()
420             - 1)
421         .withAcknowledgmentNumber(ackNum + client.
422             getThisReceivingWindowCapacity() - 1)
423         .build();
424
425     int indexBeforeSending = server.getReceivingWindow().receivingPacketIndex(
426         packet, server.getConnection());
427     Assert.assertEquals(client.getThisReceivingWindowCapacity() - 1,
428         indexBeforeSending);
429
430     eventHandler.addEvent(new RouteEvent(client, packet));
431     eventHandler.run();
432
433     int indexAfterReceived = server.getReceivingWindow().receivingPacketIndex(
434         packet, server.getConnection());
435     Assert.assertEquals(client.getThisReceivingWindowCapacity() - 1,
436         indexAfterReceived);
437 }
438
439 @Test
440 public void inWindowShouldWorkOnPacketsThatShouldBeInWindowTest() {
441     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
442         withReceivingWindowCapacity(7).build();
443     Routable router = new Router.RouterBuilder().build();
444     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
445         withReceivingWindowCapacity(7).build();
446
447     new Channel.ChannelBuilder().build(client, router);
448     new Channel.ChannelBuilder().build(router, server);
449
450     client.updateRoutingTable();
451     router.updateRoutingTable();
452     server.updateRoutingTable();
453
454     EventHandler eventHandler = new EventHandler();
455     eventHandler.addEvent(new TCPConnectEvent(client, server));
456     eventHandler.run();
457
458     System.out.println("connected");
459
460     Assert.assertNotNull(client.getConnection());
461     Assert.assertNotNull(server.getConnection());
462
463     int seqNum = client.getConnection().getNextSequenceNumber();

```

```

458     int ackNum = client.getConnection().getNextAcknowledgementNumber();
459
460     for (int i = 0; i < client.getThisReceivingWindowCapacity(); i++) {
461         Packet packet = new PacketBuilder()
462             .withOrigin(client)
463             .withDestination(server)
464             .withSequenceNumber(seqNum + i)
465             .withAcknowledgmentNumber(ackNum + i)
466             .build();
467
468
469         try {
470             ReceivingWindow receivingWindow = server.getReceivingWindow();
471             Assert.assertTrue(receivingWindow.inReceivingWindow(packet, server.
472                 getConnection()));
473         } catch (IllegalAccessException e) {
474             Assert.fail();
475         }
476     }
477
478     @Test
479     public void inWindowShouldNotWorkOnPacketsThatShouldNotBeInWindowTest() {
480         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
481             withReceivingWindowCapacity(7).build();
482         Routable router = new Router.RouterBuilder().build();
483         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
484             withReceivingWindowCapacity(7).build();
485
486         new Channel.ChannelBuilder().build(client, router);
487         new Channel.ChannelBuilder().build(router, server);
488
489         client.updateRoutingTable();
490         router.updateRoutingTable();
491         server.updateRoutingTable();
492
493         EventHandler eventHandler = new EventHandler();
494         eventHandler.addEvent(new TCPConnectEvent(client, server));
495         eventHandler.run();
496
497         System.out.println("connected");
498
499         int seqNum = client.getConnection().getNextSequenceNumber();
500         int ackNum = client.getConnection().getNextAcknowledgementNumber();
501
502         for (int i = 0; i < client.getThisReceivingWindowCapacity(); i++) {
503             Packet packet = new PacketBuilder()
504                 .withOrigin(client)
505                 .withDestination(server)
506                 .withSequenceNumber(seqNum + i + 1000)
507                 .withAcknowledgmentNumber(ackNum + i + 1000)
508                 .build();
509
510             try {
511                 ReceivingWindow receivingWindow = server.getReceivingWindow();
512                 Assert.assertFalse(receivingWindow.inReceivingWindow(packet, server.
513                     getConnection()));
514             } catch (IllegalAccessException e) {
515                 Assert.fail();
516             }
517         }
518     }

```

```

517
518 @Test
519 public void floodWithPacketsInOrderShouldWorkTest () {
520     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().setReno().
521         withReceivingWindowCapacity(7).build();
522     Rtable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.86).
523         build();
524     Rtable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.86).
525         build();
526     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().setReno().
527         withReceivingWindowCapacity(20).build();
528
529     new Channel.ChannelBuilder().build(client, r1);
530     new Channel.ChannelBuilder().build(r1, r2);
531     new Channel.ChannelBuilder().build(r2, server);
532
533     client.updateRoutingTable();
534     r1.updateRoutingTable();
535     r2.updateRoutingTable();
536     server.updateRoutingTable();
537
538     EventHandler eventHandler = new EventHandler();
539     eventHandler.addEvent(new TCPConnectEvent(client, server));
540     eventHandler.run();
541
542     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 100;
543
544     for (int i = 1; i <= numPacketsToSend; i++) {
545         Message msg = new Message("test " + i);
546         client.send(msg);
547     }
548     eventHandler.addEvent(new RunTCPEvent(client));
549     eventHandler.run();
550
551     Assert.assertTrue(client.inputBufferIsEmpty());
552     Assert.assertTrue(server.inputBufferIsEmpty());
553     Assert.assertTrue(client.outputBufferIsEmpty());
554     Assert.assertTrue(server.outputBufferIsEmpty());
555     Assert.assertTrue(r1.inputBufferIsEmpty());
556
557     for (int i = 1; i <= numPacketsToSend; i++) {
558         Message msg = new Message("test " + i);
559         Packet received = server.receive();
560         Assert.assertNotNull(received);
561         Assert.assertEquals(received.getPayload(), msg);
562     }
563 }
564
565 @Test
566 public void floodWithPacketsInBigCongestedNetworkShouldWorkTest () {
567     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
568         withReceivingWindowCapacity(10).build();
569     Rtable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
570         build();
571     Rtable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
572         build();
573     Rtable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
574         build();
575     Rtable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
576         build();
577     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
578         withReceivingWindowCapacity(30).build();

```

```

570
571     new Channel.ChannelBuilder().build(client, r1);
572     new Channel.ChannelBuilder().build(r1, r2);
573     new Channel.ChannelBuilder().build(r2, r3);
574     new Channel.ChannelBuilder().build(r3, r4);
575     new Channel.ChannelBuilder().build(r4, server);
576
577     client.updateRoutingTable();
578     r1.updateRoutingTable();
579     r2.updateRoutingTable();
580     r3.updateRoutingTable();
581     r4.updateRoutingTable();
582     server.updateRoutingTable();
583
584     EventHandler eventHandler = new EventHandler();
585     eventHandler.addEvent(new TCPConnectEvent(client, server));
586     eventHandler.run();
587
588     Assert.assertTrue(client.isConnected());
589     Assert.assertTrue(server.isConnected());
590
591     System.out.println("connected");
592
593     Assert.assertTrue(client.inputBufferIsEmpty());
594     Assert.assertTrue(server.inputBufferIsEmpty());
595     Assert.assertTrue(client.outputBufferIsEmpty());
596     Assert.assertTrue(server.outputBufferIsEmpty());
597
598     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 1000;
599     for (int i = 1; i <= numPacketsToSend; i++) {
600         Message msg = new Message("test " + i);
601         client.send(msg);
602     }
603
604     eventHandler.addEvent(new RunTCPEvent(client));
605     eventHandler.run();
606
607     Assert.assertTrue(client.inputBufferIsEmpty());
608     Assert.assertTrue(server.inputBufferIsEmpty());
609     Assert.assertTrue(client.outputBufferIsEmpty());
610     Assert.assertTrue(server.outputBufferIsEmpty());
611
612     for (int i = 1; i <= numPacketsToSend; i++) {
613         Message msg = new Message("test " + i);
614         Packet received = server.receive();
615         Assert.assertNotNull(received);
616         Assert.assertEquals(msg, received.getPayload());
617     }
618 }
619
620
621 @Test
622 public void floodWithPacketsInOrderButInLossyChannelShouldWorkTest() {
623     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
624         withReceivingWindowCapacity(30).build();
625     Routable router = new Router.RouterBuilder().build();
626     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
627         withReceivingWindowCapacity(30).build();
628
629     new Channel.ChannelBuilder().withLoss(0.001).build(client, router);
630     new Channel.ChannelBuilder().build(router, server);
631
632     client.updateRoutingTable();

```

```

631     router.updateRoutingTable();
632     server.updateRoutingTable();
633
634     EventHandler eventHandler = new EventHandler();
635     eventHandler.addEvent(new TCPConnectEvent(client, server));
636     eventHandler.run();
637
638     Assert.assertTrue(client.isConnected());
639     Assert.assertTrue(server.isConnected());
640
641     System.out.println("connected");
642
643     Assert.assertTrue(client.inputBufferIsEmpty());
644     Assert.assertTrue(server.inputBufferIsEmpty());
645     Assert.assertTrue(client.outputBufferIsEmpty());
646     Assert.assertTrue(server.outputBufferIsEmpty());
647     Assert.assertTrue(router.inputBufferIsEmpty());
648
649     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 100;
650     for (int i = 1; i <= numPacketsToSend; i++) {
651         Message msg = new Message("test " + i);
652         client.send(msg);
653     }
654
655     eventHandler.addEvent(new RunTCPEvent(client));
656     eventHandler.run();
657
658     Assert.assertTrue(client.inputBufferIsEmpty());
659     Assert.assertTrue(server.inputBufferIsEmpty());
660     Assert.assertTrue(client.outputBufferIsEmpty());
661     Assert.assertTrue(server.outputBufferIsEmpty());
662     Assert.assertTrue(router.inputBufferIsEmpty());
663
664     for (int i = 1; i <= numPacketsToSend; i++) {
665         Message msg = new Message("test " + i);
666         Packet received = server.receive();
667         Assert.assertNotNull("iteration: " + i, received);
668         Assert.assertEquals(msg, received.getPayload());
669     }
670 }
671
672 @Test
673 public void increasingWindowCapacityWithFloodWithPacketsInOrderButShouldWorkTest
674     () {
675     for (int windowCapacity = 1; windowCapacity < 100; windowCapacity++) {
676         Util.setSeed(1337);
677         Util.resetTime();
678         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
679             withReceivingWindowCapacity(windowCapacity).build();
680         Rutable router = new Router.RouterBuilder().build();
681         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
682             withReceivingWindowCapacity(windowCapacity).build();
683
684         new Channel.ChannelBuilder().build(client, router);
685         new Channel.ChannelBuilder().build(router, server);
686
687         client.updateRoutingTable();
688         router.updateRoutingTable();
689         server.updateRoutingTable();
690
691         EventHandler eventHandler = new EventHandler();
692         eventHandler.addEvent(new TCPConnectEvent(client, server));
693         eventHandler.run();

```



```

691
692     Assert.assertTrue(client.isConnected());
693     Assert.assertTrue(server.isConnected());
694
695     System.out.println("connected");
696
697     Assert.assertTrue(client.inputBufferIsEmpty());
698     Assert.assertTrue(server.inputBufferIsEmpty());
699     Assert.assertTrue(client.outputBufferIsEmpty());
700     Assert.assertTrue(server.outputBufferIsEmpty());
701     Assert.assertTrue(router.inputBufferIsEmpty());
702
703     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 10;
704     for (int i = 1; i <= numPacketsToSend; i++) {
705         Message msg = new Message("test " + i);
706         client.send(msg);
707     }
708
709     eventHandler.addEvent(new RunTCPEvent(client));
710     eventHandler.run();
711
712     Assert.assertTrue(client.inputBufferIsEmpty());
713     Assert.assertTrue(server.inputBufferIsEmpty());
714     Assert.assertTrue(client.outputBufferIsEmpty());
715     Assert.assertTrue(server.outputBufferIsEmpty());
716     Assert.assertTrue(router.inputBufferIsEmpty());
717
718     for (int i = 1; i <= numPacketsToSend; i++) {
719         Message msg = new Message("test " + i);
720         Packet received = server.receive();
721         Assert.assertNotNull(received);
722         Assert.assertEquals(msg, received.getPayload());
723     }
724 }
725 }
726
727 @Test
728 public void
increasingWindowCapacityWithFloodWithPacketsInOrderButInLossyChannelShouldWorkTest
() {
729     for (int windowCapacity = 1; windowCapacity < 50; windowCapacity++) {
730         Util.setSeed(1337);
731         Util.resetTime();
732         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
withReceivingWindowCapacity(windowCapacity).build();
733         Routable router = new Router.RouterBuilder().build();
734         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
withReceivingWindowCapacity(windowCapacity).build();
735
736         new Channel.ChannelBuilder().withLoss(0.001).build(client, router);
737         new Channel.ChannelBuilder().build(router, server);
738
739         client.updateRoutingTable();
740         router.updateRoutingTable();
741         server.updateRoutingTable();
742
743         EventHandler eventHandler = new EventHandler();
744         eventHandler.addEvent(new TCPConnectEvent(client, server));
745         eventHandler.run();
746
747         Assert.assertTrue(client.isConnected());
748         Assert.assertTrue(server.isConnected());
749

```

```

750         System.out.println("connected");
751
752         Assert.assertTrue(client.inputBufferIsEmpty());
753         Assert.assertTrue(server.inputBufferIsEmpty());
754         Assert.assertTrue(client.outputBufferIsEmpty());
755         Assert.assertTrue(server.outputBufferIsEmpty());
756         Assert.assertTrue(router.inputBufferIsEmpty());
757
758         int numPacketsToSend = server.getThisReceivingWindowCapacity() * 10;
759         for (int i = 1; i <= numPacketsToSend; i++) {
760             Message msg = new Message("test " + i);
761             client.send(msg);
762         }
763
764         eventHandler.addEvent(new RunTCPEvent(client));
765         eventHandler.run();
766
767         Assert.assertTrue(client.inputBufferIsEmpty());
768         Assert.assertTrue(server.inputBufferIsEmpty());
769         Assert.assertTrue(client.outputBufferIsEmpty());
770         Assert.assertTrue(server.outputBufferIsEmpty());
771         Assert.assertTrue(router.inputBufferIsEmpty());
772
773         for (int i = 1; i <= numPacketsToSend; i++) {
774             Message msg = new Message("test " + i);
775             Packet received = server.receive();
776             Assert.assertNotNull(received);
777             Assert.assertEquals(msg, received.getPayload());
778         }
779     }
780 }
781
782
783 @Test
784 public void floodWithPacketsBeforeConnectingShouldWorkTest() {
785     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
786         withReceivingWindowCapacity(7).build();
787     Routable router = new Router.RouterBuilder().withAverageQueueUtilization
788         (0.7).build();
789     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
790         withReceivingWindowCapacity(7).build();
791
792     new Channel.ChannelBuilder().withLoss(0.0001).build(client, router);
793     new Channel.ChannelBuilder().build(router, server);
794
795     client.updateRoutingTable();
796     router.updateRoutingTable();
797     server.updateRoutingTable();
798
799     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 1000;
800     for (int i = 1; i <= numPacketsToSend; i++) {
801         Message msg = new Message("test " + i);
802         client.send(msg);
803     }
804
805     Assert.assertFalse(client.isConnected());
806     Assert.assertFalse(server.isConnected());
807
808     EventHandler eventHandler = new EventHandler();
809     eventHandler.addEvent(new TCPConnectEvent(client, server));
810     eventHandler.run();
811
812     Assert.assertTrue(client.isConnected());

```

```

810     Assert.assertTrue(server.isConnected());
811
812     System.out.println("connected");
813
814     Assert.assertTrue(client.inputBufferIsEmpty());
815     Assert.assertTrue(server.inputBufferIsEmpty());
816     Assert.assertTrue(client.outputBufferIsEmpty());
817     Assert.assertTrue(server.outputBufferIsEmpty());
818     Assert.assertTrue(router.inputBufferIsEmpty());
819
820     for (int i = 1; i <= numPacketsToSend; i++) {
821         Message msg = new Message("test " + i);
822         Packet received = server.receive();
823         Assert.assertNotNull(received);
824         Assert.assertEquals(msg, received.getPayload());
825     }
826 }
827
828
829 @Test
830 public void serverFloodWithPacketsInOrderShouldWorkTest() {
831     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
832         withReceivingWindowCapacity(7).build();
833     Rtable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.86).
834         build();
835     Rtable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.86).
836         build();
837     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
838         withReceivingWindowCapacity(20).build();
839
840     new Channel.ChannelBuilder().build(client, r1);
841     new Channel.ChannelBuilder().build(r1, r2);
842     new Channel.ChannelBuilder().build(r2, server);
843
844     client.updateRoutingTable();
845     r1.updateRoutingTable();
846     r2.updateRoutingTable();
847     server.updateRoutingTable();
848
849     EventHandler eventHandler = new EventHandler();
850     eventHandler.addEvent(new TCPConnectEvent(client, server));
851     eventHandler.run();
852
853     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 100;
854
855     for (int i = 1; i <= numPacketsToSend; i++) {
856         Message msg = new Message("test " + i);
857         server.send(msg);
858     }
859     eventHandler.addEvent(new RunTCPEvent(server));
860     eventHandler.run();
861
862     Assert.assertTrue(client.inputBufferIsEmpty());
863     Assert.assertTrue(server.inputBufferIsEmpty());
864     Assert.assertTrue(client.outputBufferIsEmpty());
865     Assert.assertTrue(server.outputBufferIsEmpty());
866     Assert.assertTrue(r1.inputBufferIsEmpty());
867
868     for (int i = 1; i <= numPacketsToSend; i++) {
869         Message msg = new Message("test " + i);
870         Packet received = client.receive();
871         Assert.assertNotNull(received);
872         Assert.assertEquals(received.getPayload(), msg);

```

```

869     }
870 }
871
872
873 @Test
874 public void serverFloodWithPacketsBeforeConnectingShouldWorkTest() {
875     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
876         withReceivingWindowCapacity(7).build();
877     Routable router = new Router.RouterBuilder().build();
878     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
879         withReceivingWindowCapacity(7).build();
880
881     new Channel.ChannelBuilder().build(client, router);
882     new Channel.ChannelBuilder().build(router, server);
883
884     client.updateRoutingTable();
885     router.updateRoutingTable();
886     server.updateRoutingTable();
887
888     int numPacketsToSend = server.getThisReceivingWindowCapacity() * 1000;
889     for (int i = 1; i <= numPacketsToSend; i++) {
890         Message msg = new Message("test " + i);
891         server.send(msg);
892     }
893
894     Assert.assertFalse(client.isConnected());
895     Assert.assertFalse(server.isConnected());
896
897     EventHandler eventHandler = new EventHandler();
898     eventHandler.addEvent(new TCPConnectEvent(client, server));
899     eventHandler.run();
900
901     Assert.assertTrue(client.isConnected());
902     Assert.assertTrue(server.isConnected());
903
904     System.out.println("connected");
905
906     Assert.assertTrue(client.inputBufferIsEmpty());
907     Assert.assertTrue(server.inputBufferIsEmpty());
908     Assert.assertTrue(client.outputBufferIsEmpty());
909     Assert.assertTrue(server.outputBufferIsEmpty());
910     Assert.assertTrue(router.inputBufferIsEmpty());
911
912     for (int i = 1; i <= numPacketsToSend; i++) {
913         Message msg = new Message("test " + i);
914         Packet received = client.receive();
915         Assert.assertNotNull(received);
916         Assert.assertEquals(msg, received.getPayload());
917     }
918 }

```

```

1 package org.example.protocol;
2
3 import org.example.data.Message;
4 import org.example.data.Packet;
5 import org.example.data.Payload;
6 import org.example.network.Channel;
7 import org.example.network.Routable;
8 import org.example.network.Router;
9 import org.example.network.address.SimpleAddress;
10 import org.example.simulator.EventHandler;
11 import org.example.simulator.events.tcp.RunTCPEvent;
12 import org.example.simulator.events.tcp.TCPConnectEvent;
13 import org.example.util.Util;
14 import org.junit.Assert;
15 import org.junit.Before;
16 import org.junit.Rule;
17 import org.junit.Test;
18 import org.junit.rules.Timeout;
19
20 import java.util.concurrent.TimeUnit;
21
22 public class MPTCPTest {
23
24     @Rule
25     public Timeout globalTimeout = new Timeout(120, TimeUnit.SECONDS);
26
27     @Before
28     public void setup() {
29         Util.setSeed(1337);
30         Util.resetTime();
31     }
32
33     @Test
34     public void MPTCPWithTwoSubFlowsAndNonDistinctPathConnectAndSendCorrectTest () {
35         MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
36             withReceivingWindowCapacity(14).withAddress(new SimpleAddress("MPTCP-
37                 Client")).build();
38
39         Router r11 = new Router.RouterBuilder().build();
40         Router r12 = new Router.RouterBuilder().build();
41
42         Router r3 = new Router.RouterBuilder().build();
43         Router r4 = new Router.RouterBuilder().build();
44
45         MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
46             withReceivingWindowCapacity(14).withAddress(new SimpleAddress("MPTCP-
47                 Server")).build();
48
49         new Channel.ChannelBuilder().build(client, r11);
50         new Channel.ChannelBuilder().build(client, r12);
51         new Channel.ChannelBuilder().build(r11, r3);
52         new Channel.ChannelBuilder().build(r12, r3);
53         new Channel.ChannelBuilder().build(r3, r4);
54         new Channel.ChannelBuilder().build(r4, server);
55         new Channel.ChannelBuilder().build(r4, server);
56
57         client.updateRoutingTable();
58         r11.updateRoutingTable();
59         r12.updateRoutingTable();
60         r3.updateRoutingTable();
61         r4.updateRoutingTable();
62         server.updateRoutingTable();

```

```

59     EventHandler eventHandler = new EventHandler();
60     eventHandler.addEvent(new TCPConnectEvent(client, server));
61     eventHandler.run();
62     System.out.println("connected");
63
64
65
66     Message msg = new Message("Test");
67     client.send(msg);
68     eventHandler.addEvent(new RunTCPEvent(client));
69     eventHandler.run();
70
71     Packet receivedPacket = server.receive();
72     Assert.assertNotNull(receivedPacket);
73
74     Payload receivedPayload = receivedPacket.getPayload();
75     Assert.assertEquals(receivedPayload, msg);
76
77     Assert.assertNull(server.receive());
78 }
79
80 @Test
81 public void MPTCPConnectToTCPTest() {
82     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
83         withReceivingWindowCapacity(14).withAddress(new SimpleAddress("MPTCP-
84             Client")).build();
85     Routable router = new Router.RouterBuilder().withAddress(new SimpleAddress("
86         router")).build();
87     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().withAddress(new
88         SimpleAddress("server")).withReceivingWindowCapacity(7).build();
89
90     new Channel.ChannelBuilder().build(client, router);
91     new Channel.ChannelBuilder().build(client, router);
92     new Channel.ChannelBuilder().build(router, server);
93     new Channel.ChannelBuilder().build(router, server);
94
95     client.updateRoutingTable();
96     router.updateRoutingTable();
97     server.updateRoutingTable();
98
99     System.out.println(client.getChannels());
100    System.out.println(router.getChannels());
101    System.out.println(server.getChannels());
102
103    EventHandler eventHandler = new EventHandler();
104    eventHandler.addEvent(new TCPConnectEvent(client, server));
105    eventHandler.run();
106
107    Assert.assertFalse(client.isConnected());
108    Assert.assertTrue(client.getSubflows()[0].isConnected());
109    Assert.assertFalse(client.getSubflows()[1].isConnected());
110 }
111
112 @Test
113 public void MPTCPConnectThenSendMsgOverTwoSubflowsTest() {
114     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
115         withReceivingWindowCapacity(14).withAddress(new SimpleAddress("MPTCP-
116             Client")).build();
117     Routable r1 = new Router.RouterBuilder().build();
118     Routable r2 = new Router.RouterBuilder().build();
119     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().

```

```

        withReceivingWindowCapacity(7).build();
116
117 //path one
118 new Channel.ChannelBuilder().build(client, r1);
119 new Channel.ChannelBuilder().build(r1, server);
120
121 //path two
122 new Channel.ChannelBuilder().build(client, r2);
123 new Channel.ChannelBuilder().build(r2, server);
124
125
126 client.updateRoutingTable();
127 r1.updateRoutingTable();
128 r2.updateRoutingTable();
129 server.updateRoutingTable();
130
131 EventHandler eventHandler = new EventHandler();
132 eventHandler.addEvent(new TCPConnectEvent(client, server));
133 eventHandler.run();
134
135 Assert.assertTrue(client.getSubflows()[0].isConnected());
136 Assert.assertFalse(client.getSubflows()[1].isConnected());
137
138 Message msg1 = new Message("Test 1!");
139 Message msg2 = new Message("Test 2!");
140
141 client.send(msg1);
142 client.send(msg2);
143 eventHandler.addEvent(new RunTCPEvent(client));
144 eventHandler.run();
145
146 Packet received1 = server.receive();
147 Assert.assertNotNull(received1);
148 Assert.assertEquals(msg1, received1.getPayload());
149
150 Packet received2 = server.receive();
151 Assert.assertNotNull(received2);
152 Assert.assertEquals(msg2, received2.getPayload());
153 }
154
155
156 @Test
157 public void MPTCPConnectToMPTCPNonDistinctPathTest() {
158     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
        withReceivingWindowCapacity(20).withAddress(new SimpleAddress("MPTCP-
        Client")).build();
159     Routable r1 = new Router.RouterBuilder().withAddress(new SimpleAddress("A"))
        .build();
160     Routable r2 = new Router.RouterBuilder().withAddress(new SimpleAddress("B"))
        .build();
161     Routable r3 = new Router.RouterBuilder().withAddress(new SimpleAddress("C"))
        .build();
162     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(2).
        withReceivingWindowCapacity(20).withAddress(new SimpleAddress("MPTCP-
        Server")).build();
163
164     new Channel.ChannelBuilder().build(client, r1);
165     new Channel.ChannelBuilder().build(client, r2);
166
167     new Channel.ChannelBuilder().build(r3, r1);
168     new Channel.ChannelBuilder().build(r3, r2);
169
170     new Channel.ChannelBuilder().build(server, r3);

```

```

171     new Channel.ChannelBuilder().build(server, r3);
172
173
174     client.updateRoutingTable();
175     r1.updateRoutingTable();
176     r2.updateRoutingTable();
177     r3.updateRoutingTable();
178     server.updateRoutingTable();
179
180     EventHandler eventHandler = new EventHandler();
181     eventHandler.addEvent(new TCPConnectEvent(client, server));
182     eventHandler.run();
183     System.out.println("connected");
184
185     Assert.assertTrue(client.isConnected());
186
187     Message msg1 = new Message("Test 1!");
188     Message msg2 = new Message("Test 2!");
189     Message msg3 = new Message("Test 3!");
190
191     client.send(msg1);
192     client.send(msg2);
193     client.send(msg3);
194     eventHandler.addEvent(new RunTCPEvent(client));
195     eventHandler.run();
196
197     Packet received1 = server.receive();
198     Assert.assertNotNull(received1);
199     Assert.assertEquals(msg1, received1.getPayload());
200
201     Packet received2 = server.receive();
202     Assert.assertNotNull(received2);
203     Assert.assertEquals(msg2, received2.getPayload());
204
205     Packet received3 = server.receive();
206     Assert.assertNotNull(received3);
207     Assert.assertEquals(msg3, received3.getPayload());
208
209     Assert.assertNull(server.receive());
210 }
211
212 @Test
213 public void MPTCPConnectToMPTCPThenSendMsgOverThreeSubflowsTest() {
214     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
215         withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
216             Client")).build();
217     Routable r1 = new Router.RouterBuilder().withAddress(new SimpleAddress("A"))
218         .build();
219     Routable r2 = new Router.RouterBuilder().withAddress(new SimpleAddress("B"))
220         .build();
221     Routable r3 = new Router.RouterBuilder().withAddress(new SimpleAddress("C"))
222         .build();
223     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
224         withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
225             Server")).build();
226
227     //path one
228     new Channel.ChannelBuilder().build(client, r1);
229     new Channel.ChannelBuilder().build(r1, server);
230
231     //path two
232     new Channel.ChannelBuilder().build(client, r2);
233     new Channel.ChannelBuilder().build(r2, server);

```



```

227
228 //path three
229 new Channel.ChannelBuilder().build(client, r3);
230 new Channel.ChannelBuilder().build(r3, server);
231
232
233 client.updateRoutingTable();
234 r1.updateRoutingTable();
235 r2.updateRoutingTable();
236 r3.updateRoutingTable();
237 server.updateRoutingTable();
238
239 EventHandler eventHandler = new EventHandler();
240 eventHandler.addEvent(new TCPConnectEvent(client, server));
241 eventHandler.run();
242 System.out.println("connected");
243
244 Assert.assertTrue(client.getSubflows()[0].isConnected());
245 Assert.assertTrue(client.getSubflows()[1].isConnected());
246 Assert.assertTrue(client.getSubflows()[2].isConnected());
247
248 Message msg1 = new Message("Test 1!");
249 Message msg2 = new Message("Test 2!");
250 Message msg3 = new Message("Test 3!");
251
252 client.send(msg1);
253 client.send(msg2);
254 client.send(msg3);
255 eventHandler.addEvent(new RunTCPEvent(client));
256 eventHandler.run();
257
258 Packet received1 = server.receive();
259 Assert.assertNotNull(received1);
260 Assert.assertEquals(msg1, received1.getPayload());
261
262 Packet received2 = server.receive();
263 Assert.assertNotNull(received2);
264 Assert.assertEquals(msg2, received2.getPayload());
265
266 Packet received3 = server.receive();
267 Assert.assertNotNull(received3);
268 Assert.assertEquals(msg3, received3.getPayload());
269
270 Assert.assertNull(server.receive());
271 }
272
273
274 @Test
275 public void MPTCPFloodWithPacketsInOrderShouldWorkTest() {
276     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
277         withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
278             Client")).build();
279     Routable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
280         withAddress(new SimpleAddress("A")).build();
281     Routable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
282         withAddress(new SimpleAddress("B")).build();
283     Routable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.8).
284         withAddress(new SimpleAddress("C")).build();
285     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
286         withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
287             Server")).build();
288
289 //path one

```

```

283     new Channel.ChannelBuilder().build(client, r1);
284     new Channel.ChannelBuilder().build(r1, server);
285
286     //path two
287     new Channel.ChannelBuilder().build(client, r2);
288     new Channel.ChannelBuilder().build(r2, server);
289
290     //path three
291     new Channel.ChannelBuilder().build(client, r3);
292     new Channel.ChannelBuilder().build(r3, server);
293
294
295     client.updateRoutingTable();
296     r1.updateRoutingTable();
297     r2.updateRoutingTable();
298     r3.updateRoutingTable();
299     server.updateRoutingTable();
300
301     EventHandler eventHandler = new EventHandler();
302     eventHandler.addEvent(new TCPConnectEvent(client, server));
303     eventHandler.run();
304     System.out.println("connected");
305
306     Assert.assertTrue(client.getSubflows()[0].isConnected());
307     Assert.assertTrue(client.getSubflows()[1].isConnected());
308     Assert.assertTrue(client.getSubflows()[2].isConnected());
309
310     int numPacketsToSend = 10000;
311
312     for (int i = 1; i <= numPacketsToSend; i++) {
313         Message msg = new Message("test " + i);
314         client.send(msg);
315     }
316     eventHandler.addEvent(new RunTCPEvent(client));
317     eventHandler.run();
318
319     Assert.assertTrue("client still has packets to send", client.
320         outputBufferIsEmpty());
321     Assert.assertTrue(server.outputBufferIsEmpty());
322     Assert.assertTrue(client.inputBufferIsEmpty());
323     Assert.assertTrue(server.inputBufferIsEmpty());
324     Assert.assertTrue(r1.inputBufferIsEmpty());
325     Assert.assertTrue(r2.inputBufferIsEmpty());
326     Assert.assertTrue(r3.inputBufferIsEmpty());
327
328     for (int i = 1; i <= numPacketsToSend; i++) {
329         Message msg = new Message("test " + i);
330         Packet received = server.receive();
331         Assert.assertNotNull(received);
332         Assert.assertEquals("iteration " + i, received.getPayload(), msg);
333     }
334     Assert.assertNull(server.receive());
335
336 }
337
338 @Test
339 public void MPTCPFloodWithPacketsInLossyChannelsShouldWorkTest() {
340     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
341         withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
342             Client")).build();
343     Routable r1 = new Router.RouterBuilder().withAddress(new SimpleAddress("A"))
344         .build();

```

```

342     Routable r2 = new Router.RouterBuilder().withAddress(new SimpleAddress("B"))
        .build();
343     Routable r3 = new Router.RouterBuilder().withAddress(new SimpleAddress("C"))
        .build();
344     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
        withReceivingWindowCapacity(21).withAddress(new SimpleAddress("MPTCP-
        Server")).build();
345
346     //path one
347     new Channel.ChannelBuilder().withLoss(0.001).build(client, r1);
348     new Channel.ChannelBuilder().withLoss(0.01).build(r1, server);
349
350     //path two
351     new Channel.ChannelBuilder().withLoss(0.001).build(client, r2);
352     new Channel.ChannelBuilder().withLoss(0.001).build(r2, server);
353
354     //path three
355     new Channel.ChannelBuilder().withLoss(0.01).build(client, r3);
356     new Channel.ChannelBuilder().withLoss(0.001).build(r3, server);
357
358     client.updateRoutingTable();
359     r1.updateRoutingTable();
360     r2.updateRoutingTable();
361     r3.updateRoutingTable();
362     server.updateRoutingTable();
363
364     EventHandler eventHandler = new EventHandler();
365     eventHandler.addEvent(new TCPConnectEvent(client, server));
366     eventHandler.run();
367     System.out.println("connected");
368
369     Assert.assertTrue(client.getSubflows()[0].isConnected());
370     Assert.assertTrue(client.getSubflows()[1].isConnected());
371     Assert.assertTrue(client.getSubflows()[2].isConnected());
372
373     int numPacketsToSend = 10000;
374
375     for (int i = 1; i <= numPacketsToSend; i++) {
376         Message msg = new Message("test " + i);
377         client.send(msg);
378     }
379     eventHandler.addEvent(new RunTCPEvent(client));
380     eventHandler.run();
381
382     Assert.assertTrue("client still has packets to send", client.
        outputBufferIsEmpty());
383     Assert.assertTrue(server.outputBufferIsEmpty());
384
385     Assert.assertTrue(client.inputBufferIsEmpty());
386     Assert.assertTrue(server.inputBufferIsEmpty());
387     Assert.assertTrue(r1.inputBufferIsEmpty());
388     Assert.assertTrue(r2.inputBufferIsEmpty());
389     Assert.assertTrue(r3.inputBufferIsEmpty());
390
391     for (int i = 1; i <= numPacketsToSend; i++) {
392         Message msg = new Message("test " + i);
393         Packet received = server.receive();
394         Assert.assertNotNull(received);
395         Assert.assertEquals("iteration " + i, received.getPayload(), msg);
396     }
397     Assert.assertNull(server.receive());
398
399 }

```

```

400
401
402 @Test
403 public void
404     MPTCPFloodWithPacketsInOrderWithVariableNumberOfSubflowsShouldWorkTest() {
405     int maxSubflows = 5;
406     for (int numSubflows = 1; numSubflows <= maxSubflows; numSubflows++) {
407         Util.setSeed(1337);
408         Util.resetTime();
409         MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(numSubflows)
410             .withReceivingWindowCapacity(21).withAddress(new SimpleAddress("
411                 MPTCP-Client")).build();
412         MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(numSubflows)
413             .withReceivingWindowCapacity(21).withAddress(new SimpleAddress("
414                 MPTCP-Server")).build();
415
416         for (int i = 1; i <= numSubflows; i++) {
417             Routable router = new Router.RouterBuilder().withAddress(new
418                 SimpleAddress("Router " + i)).build();
419
420             new Channel.ChannelBuilder().build(client, router);
421             new Channel.ChannelBuilder().build(router, server);
422             router.updateRoutingTable();
423         }
424         client.updateRoutingTable();
425         server.updateRoutingTable();
426
427         EventHandler eventHandler = new EventHandler();
428         eventHandler.addEvent(new TCPCConnectEvent(client, server));
429         eventHandler.run();
430         System.out.println("connected");
431
432         for (TCP subflow : client.getSubflows()) {
433             Assert.assertTrue(subflow.isConnected());
434         }
435
436         int numPacketsToSend = 10000;
437         for (int i = 1; i <= numPacketsToSend; i++) {
438             Message msg = new Message("test " + i);
439             client.send(msg);
440         }
441         eventHandler.addEvent(new RunTCPEvent(client));
442         eventHandler.run();
443
444         Assert.assertTrue("client still has packets to send", client.
445             outputBufferIsEmpty());
446         Assert.assertTrue(server.outputBufferIsEmpty());
447         Assert.assertTrue(client.inputBufferIsEmpty());
448         Assert.assertTrue(server.inputBufferIsEmpty());
449         for (TCP subflow : client.getSubflows()) {
450             Assert.assertTrue(subflow.inputBufferIsEmpty());
451         }
452         for (int i = 1; i <= numPacketsToSend; i++) {
453             Message msg = new Message("test " + i);
454             Packet received = server.receive();
455             Assert.assertNotNull(received);
456             Assert.assertEquals("iteration " + i, received.getPayload(), msg);
457         }
458         Assert.assertNull(server.receive());
459     }
460 }

```

456 |
457 | }

```

1 package org.example.protocol.window;
2
3 import org.example.data.*;
4 import org.example.network.Channel;
5 import org.example.protocol.ClassicTCP;
6 import org.example.protocol.window.receiving.ReceivingWindow;
7 import org.example.protocol.window.receiving.SelectiveRepeat;
8 import org.example.protocol.window.sending.SendingWindow;
9 import org.example.protocol.window.sending.SlidingWindow;
10 import org.example.simulator.EventHandler;
11 import org.example.simulator.events.tcp.TCPConnectEvent;
12 import org.example.simulator.statistics.TCPStats;
13 import org.javatuples.Pair;
14 import org.junit.Assert;
15 import org.junit.Before;
16 import org.junit.Rule;
17 import org.junit.Test;
18 import org.junit.rules.Timeout;
19
20 import java.util.ArrayList;
21 import java.util.Comparator;
22 import java.util.List;
23 import java.util.concurrent.TimeUnit;
24
25
26 public class ReceivingWindowTest {
27
28     private static final Comparator<Packet> PACKET_COMPARATOR = Comparator.
29         comparingInt (Packet::getSequenceNumber);
30
31     @Rule
32     public Timeout globalTimeout = new Timeout (30, TimeUnit.SECONDS);
33
34     private ClassicTCP client;
35     private ClassicTCP server;
36     private ReceivingWindow receivingWindow;
37     private List<Packet> receivedPackets;
38     private List<Pair<Integer, Payload>> payloadsToSend;
39
40     @Before
41     public void setup() {
42         this.receivedPackets = new ArrayList<> ();
43         this.payloadsToSend = new ArrayList<> ();
44         this.client = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
45             (7).build();
46         this.server = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
47             (7).build();
48         this.connect (client, server);
49
50         this.receivingWindow = new SelectiveRepeat (10, PACKET_COMPARATOR, this.
51             receivedPackets);
52     }
53
54     private void connect (ClassicTCP client, ClassicTCP server) {
55         new Channel.ChannelBuilder().build (client, server);
56         client.updateRoutingTable ();
57         server.updateRoutingTable ();
58
59         EventHandler eventHandler = new EventHandler ();
60         eventHandler.addEvent (new TCPConnectEvent (client, server));
61         eventHandler.run ();
62     }
63
64 }

```

```

59     @Test
60     public void ackThisReturnsPacketTest() {
61         Packet toBeAcked = this.receivingWindow.ackThis(this.client);
62         Assert.assertNotNull(toBeAcked);
63     }
64
65
66     @Test
67     public void shouldAckReturnsTrueTest() {
68         Assert.assertTrue(this.receivingWindow.shouldAck());
69     }
70
71     @Test
72     public void getReceivedPacketsShouldReturnEmptyQueueWhenNoPacketsAreReceived() {
73         Assert.assertTrue(this.receivedPackets.isEmpty());
74     }
75
76     @Test
77     public void receiveShouldReturnFalseIfNoPacketsAreReceived() {
78         Assert.assertFalse(this.receivingWindow.receive(null));
79     }
80
81
82     @Test(expected = NullPointerException.class)
83     public void receiveAckWithNullSendingWindowArgumentCausesNullPointerException()
84     {
85         Packet ackFromServer = new PacketBuilder()
86             .withSequenceNumber(this.client.getConnection().
87                 getNextAcknowledgementNumber()) //hack to find clients expected
88                 next packet sequence number
89             .withFlags(Flag.ACK)
90             .withConnection(this.server.getConnection())
91             .build();
92         Assert.assertTrue(this.receivingWindow.inReceivingWindow(ackFromServer,
93             client.getConnection()));
94         this.receivingWindow.offer(ackFromServer);
95         Assert.assertFalse(this.receivingWindow.receive(null));
96     }
97
98     @Test
99     public void receiveReturnsFalseIfAckIsReceived() {
100         Packet ackFromServer = new PacketBuilder()
101             .withSequenceNumber(this.client.getConnection().
102                 getNextAcknowledgementNumber())
103             .withFlags(Flag.ACK)
104             .withConnection(this.server.getConnection())
105             .build();
106
107         Assert.assertTrue(this.receivingWindow.inReceivingWindow(ackFromServer,
108             client.getConnection()));
109         this.receivingWindow.offer(ackFromServer);
110
111         SendingWindow sendingWindow = new SlidingWindow(10, true, client.
112             getConnection(), PACKET_COMPARATOR, this.payloadsToSend, new TCPStats(
113                 this.client.getAddress()));
114         Assert.assertFalse(this.receivingWindow.receive(sendingWindow));
115     }
116
117     @Test
118     public void receiveReturnsTrueIfPacketIsReceived() {
119         Packet packetFromServer = new PacketBuilder()
120             .withSequenceNumber(this.client.getConnection().
121                 getNextAcknowledgementNumber())

```

```
113         .withPayload(new Message("test"))
114         .withConnection(this.server.getConnection())
115         .build();
116
117     Assert.assertTrue(this.receivingWindow.inReceivingWindow(packetFromServer,
118         client.getConnection()));
119     this.receivingWindow.offer(packetFromServer);
120
121     SendingWindow sendingWindow = new SlidingWindow(10, true, client.
122         getConnection(), PACKET_COMPARATOR, this.payloadsToSend, new TCPStats(
123         this.client.getAddress()));
124     Assert.assertTrue(this.receivingWindow.receive(sendingWindow));
125
126     Assert.assertEquals(packetFromServer, this.receivedPackets.remove(0));
127 }
```



```

1 package org.example.protocol.window;
2
3 import org.example.data.Message;
4 import org.example.data.Packet;
5 import org.example.data.PacketBuilder;
6 import org.example.network.Channel;
7 import org.example.protocol.ClassicTCP;
8 import org.example.protocol.window.sending.SendingWindow;
9 import org.example.simulator.EventHandler;
10 import org.example.simulator.events.tcp.RunTCPEvent;
11 import org.example.simulator.events.tcp.TCPConnectEvent;
12 import org.junit.Assert;
13 import org.junit.Before;
14 import org.junit.Rule;
15 import org.junit.Test;
16 import org.junit.rules.Timeout;
17
18 import java.util.concurrent.TimeUnit;
19
20
21 public class SendingWindowTest {
22
23     @Rule
24     public Timeout globalTimeout = new Timeout(30, TimeUnit.SECONDS);
25     private ClassicTCP client;
26     private ClassicTCP server;
27     private SendingWindow sendingWindow;
28     private EventHandler eventHandler;
29
30     @Before
31     public void setup() throws IllegalAccessException {
32         this.client = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
33             (20).setTahoe().build();
34         this.server = new ClassicTCP.ClassicTCPBuilder().withReceivingWindowCapacity
35             (20).setTahoe().build();
36         this.connect(client, server);
37
38         this.sendingWindow = this.client.getSendingWindow();
39
40     }
41
42     private void connect(ClassicTCP client, ClassicTCP server) {
43         new Channel.ChannelBuilder().build(client, server);
44         client.updateRoutingTable();
45         server.updateRoutingTable();
46
47         this.eventHandler = new EventHandler();
48         eventHandler.addEvent(new TCPConnectEvent(client, server));
49         eventHandler.run();
50     }
51
52     @Test
53     public void initCorrectWindowSizeTest() {
54         Assert.assertEquals(1, this.sendingWindow.getWindowCapacity());
55     }
56
57     @Test
58     public void initIsWaitingForAckIsFalseIfNoPacketsAreSentTest() {
59         Assert.assertFalse(this.sendingWindow.isWaitingForAck());
60     }
61
62     @Test

```

```

61 public void initIsWaitingForAckIsTrueIfSendingWindowIsFullTest () {
62     for (int i = 0; i < this.client.getThisReceivingWindowCapacity(); i++) {
63         this.client.send(new Message("test " + i));
64     }
65     for (int i = 0; i < this.client.getThisReceivingWindowCapacity(); i++) {
66         this.client.trySend();
67     }
68     Assert.assertTrue(this.sendingWindow.isWaitingForAck());
69 }
70
71 @Test
72 public void isWaitingForAckIsFalseIfSendingWindowIsAlmostFullTest () throws
73     IllegalAccessException {
74     for (int i = 0; i < Math.pow(this.server.getThisReceivingWindowCapacity(),
75         2); i++) {
76         this.sendingWindow.increase();
77     }
78     for (int i = 0; i < this.client.getOtherReceivingWindowCapacity() - 1; i++)
79     {
80         this.client.getSendingWindow().add(new PacketBuilder().build());
81     }
82     Assert.assertEquals(this.client.getOtherReceivingWindowCapacity(), this.
83         sendingWindow.getWindowCapacity());
84     Assert.assertFalse(this.sendingWindow.isWaitingForAck());
85 }
86
87 @Test
88 public void windowCantIncreaseToMoreThanServersReceivingWindow () {
89     for (int i = 0; i < Math.pow(this.server.getThisReceivingWindowCapacity(),
90         2) * 10; i++) {
91         this.sendingWindow.increase();
92     }
93     Assert.assertEquals(this.server.getThisReceivingWindowCapacity(), this.
94         sendingWindow.getWindowCapacity());
95 }
96
97 @Test
98 public void windowCanDecreaseWhenInMaxCapacity () {
99     for (int i = 0; i < Math.pow(this.server.getThisReceivingWindowCapacity(),
100         2); i++) {
101         this.sendingWindow.increase();
102     }
103     this.sendingWindow.decrease();
104     Assert.assertEquals(1, this.sendingWindow.getWindowCapacity());
105 }
106
107 @Test
108 public void initWindowCanNotDecrease () {
109     this.sendingWindow.decrease();
110     Assert.assertEquals(1, this.sendingWindow.getWindowCapacity());
111 }
112
113 @Test
114 public void windowWillDecreaseToDefaultValueIfDecreasedEnough () {
115     for (int i = 0; i < this.server.getThisReceivingWindowCapacity(); i++) {
116         this.sendingWindow.increase();
117     }
118     for (int i = 0; i < this.server.getThisReceivingWindowCapacity(); i++) {
119         this.sendingWindow.decrease();
120     }
121     Assert.assertEquals(1, this.sendingWindow.getWindowCapacity());
122 }

```

```

117
118     @Test
119     public void floodWithPacketsInLossyChannelShouldResultInVariableWindowCapacity()
120         throws IllegalAccessException {
121         Assert.assertTrue(server.isConnected());
122         Assert.assertTrue(client.isConnected());
123
124         System.out.println("connected");
125
126         Assert.assertTrue(client.inputBufferIsEmpty());
127         Assert.assertTrue(server.inputBufferIsEmpty());
128         Assert.assertTrue(client.outputBufferIsEmpty());
129         Assert.assertTrue(server.outputBufferIsEmpty());
130
131         int numPacketsToSend = server.getThisReceivingWindowCapacity() * 100;
132         for (int i = 1; i <= numPacketsToSend; i++) {
133             Message msg = new Message("test " + i);
134             client.send(msg);
135         }
136
137         int prevWindowCapacity = client.getSendingWindow().getWindowCapacity();
138         eventHandler.addEvent(new RunTCPEvent(client));
139         while (eventHandler.singleRun()) {
140             int curWindowCapacity = client.getSendingWindow().getWindowCapacity();
141
142             boolean loss = curWindowCapacity < prevWindowCapacity;
143             boolean packetAked = curWindowCapacity > prevWindowCapacity;
144
145             if (loss) {
146                 Assert.assertEquals(1, curWindowCapacity);
147             } else if (packetAked) {
148                 Assert.assertTrue(this.client.getOtherReceivingWindowCapacity() >=
149                     curWindowCapacity);
150             } else {
151                 Assert.assertEquals(prevWindowCapacity, curWindowCapacity);
152             }
153             prevWindowCapacity = curWindowCapacity;
154         }
155
156         Assert.assertTrue(client.inputBufferIsEmpty());
157         Assert.assertTrue(server.inputBufferIsEmpty());
158         Assert.assertTrue(client.outputBufferIsEmpty());
159         Assert.assertTrue(server.outputBufferIsEmpty());
160
161         for (int i = 1; i <= numPacketsToSend; i++) {
162             Message msg = new Message("test " + i);
163             Packet received = server.receive();
164             Assert.assertNotNull(received);
165             Assert.assertEquals(msg, received.getPayload());
166         }
167     }
168 }
169

```

```

1 package org.example.simulator;
2
3 import org.example.data.Message;
4 import org.example.data.Packet;
5 import org.example.network.Channel;
6 import org.example.network.Router;
7 import org.example.protocol.ClassicTCP;
8 import org.example.simulator.events.Event;
9 import org.example.simulator.events.tcp.RunTCPEvent;
10 import org.example.simulator.events.tcp.TCPConnectEvent;
11 import org.example.util.Util;
12 import org.junit.Assert;
13 import org.junit.Before;
14 import org.junit.Rule;
15 import org.junit.Test;
16 import org.junit.rules.Timeout;
17
18 import java.util.ArrayList;
19 import java.util.Queue;
20 import java.util.concurrent.TimeUnit;
21
22
23 public class EventHandlerTest {
24
25     @Rule
26     public Timeout globalTimeout = new Timeout(30, TimeUnit.SECONDS);
27
28     @Before
29     public void setup() {
30         Util.setSeed(1337);
31         Util.resetTime();
32     }
33
34     @Test
35     public void runRunsWithoutErrorTest() {
36         EventHandler eventHandler = new EventHandler();
37         Event eventOne = new Event(Util.getTime()) {
38             @Override
39             public void run() {
40                 System.out.println(this.getInstant());
41             }
42
43             @Override
44             public void generateNextEvent(Queue<Event> events) {
45                 events.add(new Event(Util.getTime()) {
46                     @Override
47                     public void run() {
48                         System.out.println(this.getInstant());
49                     }
50
51                     @Override
52                     public void generateNextEvent(Queue<Event> events) {
53
54                 }
55             });
56         }
57     };
58
59     Event eventTwo = new Event(Util.getTime()) {
60         @Override
61         public void run() {
62             System.out.println(this.getInstant());

```

```

63     }
64
65     @Override
66     public void generateNextEvent (Queue<Event> events) {
67         events.add(new Event (Util.getTime()) {
68             @Override
69             public void run() {
70                 System.out.println(this.getInstant());
71             }
72
73             @Override
74             public void generateNextEvent (Queue<Event> events) {
75
76                 }
77         });
78     }
79 };
80
81 //start condition
82 eventHandler.addEvent (eventOne);
83 eventHandler.addEvent (eventTwo);
84
85 eventHandler.run();
86
87 Assert.assertEquals(0, eventHandler.getNumberOfEvents());
88 }
89
90
91 @Test
92 public void runTest() {
93     EventHandler eventHandler = new EventHandler();
94
95     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
96         withReceivingWindowCapacity(7).build();
97     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
98         withReceivingWindowCapacity(7).build();
99     Router r1 = new Router.RouterBuilder().build();
100
101     new Channel.ChannelBuilder().build(client, r1);
102     new Channel.ChannelBuilder().build(r1, server);
103
104     client.updateRoutingTable();
105     r1.updateRoutingTable();
106     server.updateRoutingTable();
107
108     eventHandler.addEvent (new TCPConnectEvent (client, server));
109     eventHandler.run();
110
111     client.send(new Message ("test"));
112     eventHandler.addEvent (new RunTCPEvent (client));
113     eventHandler.run();
114
115     Assert.assertEquals(0, eventHandler.getNumberOfEvents());
116 }
117
118 @Test
119 public void runFloodWithPacketsInOrderButInLossyChannelShouldWorkTest() {
120     EventHandler eventHandler = new EventHandler();
121
122     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
123         withReceivingWindowCapacity(7).build();
124     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().

```

```

123         withReceivingWindowCapacity(7).build();
124 Router r1 = new Router.RouterBuilder().build();
125
126 new Channel.ChannelBuilder().build(client, r1);
127 new Channel.ChannelBuilder().build(r1, server);
128
129 client.updateRoutingTable();
130 r1.updateRoutingTable();
131 server.updateRoutingTable();
132
133 eventHandler.addEvent(new TCPConnectEvent(client, server));
134 eventHandler.run();
135
136 int multiplier = 100;
137 int numPacketsToSend = server.getThisReceivingWindowCapacity() * multiplier;
138
139 for (int i = 1; i <= numPacketsToSend; i++) {
140     Message msg = new Message("test " + i);
141     client.send(msg);
142 }
143 eventHandler.addEvent(new RunTCPEvent(client));
144 eventHandler.run();
145
146 for (int i = 1; i <= numPacketsToSend; i++) {
147     Message msg = new Message("test " + i);
148     Packet received = server.receive();
149     Assert.assertNotNull(received);
150     Assert.assertEquals(received.getPayload(), msg);
151 }
152
153 private ArrayList<Event> allEventsList(int numPacketsToSend, double
154     noiseTolerance) {
155     Util.setSeed(1337);
156     Util.resetTime();
157     EventHandler eventHandler = new EventHandler();
158
159     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
160         withReceivingWindowCapacity(7).build();
161     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
162         withReceivingWindowCapacity(7).build();
163     Router r1 = new Router.RouterBuilder().build();
164
165     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(client, r1);
166     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(r1, server);
167
168     client.updateRoutingTable();
169     r1.updateRoutingTable();
170     server.updateRoutingTable();
171
172     eventHandler.addEvent(new TCPConnectEvent(client, server));
173     eventHandler.run();
174
175     Assert.assertNull(eventHandler.peekEvent());
176
177     for (int i = 1; i <= numPacketsToSend; i++) {
178         Message msg = new Message("test " + i);
179         client.send(msg);
180     }
181     eventHandler.addEvent(new RunTCPEvent(client));
182
183     ArrayList<Event> eventList = new ArrayList<>();

```

```

182     while (eventHandler.peekEvent() != null) {
183         eventList.add(eventHandler.peekEvent());
184         eventHandler.singleRun();
185     }
186
187     Assert.assertNull(server.dequeueInputBuffer());
188     Assert.assertNull(client.dequeueInputBuffer());
189     Assert.assertNull(r1.dequeueInputBuffer());
190     Assert.assertNull(eventHandler.peekEvent());
191
192     return eventList;
193 }
194
195 @Test
196 public void eventArrangementsAreConsistent() {
197     double noiseTolerance = 0.01;
198     int numPacketsToSend = 1001;
199     ArrayList<Event> eventList1 = this.allEventsList(numPacketsToSend,
200         noiseTolerance);
201     ArrayList<Event> eventList2 = this.allEventsList(numPacketsToSend,
202         noiseTolerance);
203
204     for (Event event : eventList1) {
205         Assert.assertEquals(event.getClass(), eventList2.remove(0).getClass());
206     }
207 }
208
209 @Test
210 public void eventAreRunningInCorrectOrderWithRespectToTime() {
211     double noiseTolerance = 0.01;
212     EventHandler eventHandler = new EventHandler();
213
214     ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
215         withReceivingWindowCapacity(7).build();
216     ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
217         withReceivingWindowCapacity(7).build();
218     Router r1 = new Router.RouterBuilder().build();
219
220     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(client, r1);
221     new Channel.ChannelBuilder().withLoss(noiseTolerance).build(r1, server);
222
223     client.updateRoutingTable();
224     r1.updateRoutingTable();
225     server.updateRoutingTable();
226
227     eventHandler.addEvent(new TCPConnectEvent(client, server));
228     eventHandler.run();
229
230     Assert.assertNull(eventHandler.peekEvent());
231
232     Util.setSeed(1337);
233
234     int numPacketsToSend = 1000;
235
236     for (int i = 1; i <= numPacketsToSend; i++) {
237         Message msg = new Message("test " + i);
238         client.send(msg);
239     }
240     eventHandler.addEvent(new RunTCPEvent(client));
241
242     Event prevEvent;
243     while (true) {

```

```
241     prevEvent = eventHandler.peekEvent();
242     eventHandler.singleRun();
243
244     if (eventHandler.peekEvent() == null) break;
245
246     Event curEvent = eventHandler.peekEvent();
247     Assert.assertTrue(prevEvent.getInstant() <= curEvent.getInstant());
248 }
249 Assert.assertNull(server.dequeueInputBuffer());
250 Assert.assertNull(client.dequeueInputBuffer());
251 Assert.assertNull(r1.dequeueInputBuffer());
252 Assert.assertNull(eventHandler.peekEvent());
253 }
254 }
```



```

1 package org.example.simulator;
2
3 import org.example.data.Message;
4 import org.example.data.PacketBuilder;
5 import org.example.network.Channel;
6 import org.example.protocol.ClassicTCP;
7 import org.example.protocol.TCP;
8 import org.example.simulator.events.ChannelEvent;
9 import org.example.simulator.events.Event;
10 import org.example.simulator.events.RouteEvent;
11 import org.example.simulator.events.tcp.RunTCPEvent;
12 import org.example.simulator.events.tcp.TCPConnectEvent;
13 import org.example.simulator.events.tcp.TCPRetransmitEventGenerator;
14 import org.example.util.Util;
15 import org.junit.Assert;
16 import org.junit.Test;
17
18 import java.util.PriorityQueue;
19 import java.util.Queue;
20
21 public class EventTest {
22
23     private Queue<Event> events;
24     private ClassicTCP tcp;
25     private ClassicTCP host;
26
27     public void connect(TCP linkedClient, TCP linkedServer) {
28         EventHandler eventHandler = new EventHandler();
29         eventHandler.addEvent(new TCPConnectEvent(linkedClient, linkedServer));
30         eventHandler.run();
31         Assert.assertTrue(linkedClient.isConnected());
32         Assert.assertTrue(linkedServer.isConnected());
33     }
34
35     @Test(expected = IllegalArgumentException.class)
36     public void expectIllegalArgumentExceptionIfNullChannelGiven() {
37         new ChannelEvent(null);
38     }
39
40     @Test(expected = IllegalArgumentException.class)
41     public void expectIllegalArgumentExceptionIfNullNodeGiven() {
42         new RunTCPEvent(null);
43     }
44
45     @Test
46     public void routeEventEqualsTest() {
47         Event event1 = new RouteEvent(new ClassicTCP.ClassicTCPBuilder().build(),
48             new PacketBuilder().build());
49         Event event2 = new RouteEvent(new ClassicTCP.ClassicTCPBuilder().build(),
50             new PacketBuilder().build());
51         Assert.assertNotEquals(event1, event2);
52     }
53
54     @Test
55     public void EventsOccurInAPreMatchedSequence() {
56         int numberOfRuns = 100;
57         for (int i = 0; i < numberOfRuns; i++) {
58             Util.resetTime();
59             Util.setSeed(1337);
60             this.events = new PriorityQueue<>();
61             this.tcp = new ClassicTCP.ClassicTCPBuilder().

```

```

61         withReceivingWindowCapacity(7).build();
        this.host = new ClassicTCP.ClassicTCPBuilder().
            withReceivingWindowCapacity(7).build();
62
63     new Channel.ChannelBuilder().build(this.tcp, this.host);
64
65     this.tcp.updateRoutingTable();
66     this.host.updateRoutingTable();
67     this.connect(this.tcp, this.host);
68
69     this.tcp.send(new Message("hello"));
70     this.events.add(new RunTCPEvent(this.tcp));
71
72     String debugString = "Iteration: " + (i + 1);
73
74     Event curEvent = this.events.poll();
75     Assert.assertNotNull(curEvent);
76     Assert.assertEquals(debugString, RunTCPEvent.class, curEvent.getClass())
77         ;
78     curEvent.run();
79     curEvent.generateNextEvent(this.events);
80
81     curEvent = this.events.poll();
82     Assert.assertNotNull(curEvent);
83     Assert.assertEquals(debugString, ChannelEvent.class, curEvent.getClass())
84         );
85     curEvent.run();
86     curEvent.generateNextEvent(this.events);
87
88     curEvent = this.events.poll();
89     Assert.assertNotNull(curEvent);
90     Assert.assertEquals(debugString, RunTCPEvent.class, curEvent.getClass())
91         ;
92     curEvent.run();
93     curEvent.generateNextEvent(this.events);
94
95     curEvent = this.events.poll();
96     Assert.assertNotNull(curEvent);
97     Assert.assertEquals(debugString, ChannelEvent.class, curEvent.getClass())
98         );
99     curEvent.run();
100    curEvent.generateNextEvent(this.events);
101
102    curEvent = this.events.poll();
103    Assert.assertNotNull(curEvent);
104    Assert.assertEquals(debugString, ChannelEvent.class, curEvent.getClass())
105        );
106    curEvent.run();
107    curEvent.generateNextEvent(this.events);
108
109    curEvent = this.events.poll();
110    Assert.assertNotNull(curEvent);
111    Assert.assertEquals(debugString, RunTCPEvent.class, curEvent.getClass())
112        ;
113    curEvent.run();
114    curEvent.generateNextEvent(this.events);
115
116    Assert.assertTrue(debugString, this.tcp.outputBufferIsEmpty());
117    Assert.assertTrue(debugString, this.host.outputBufferIsEmpty());
118    Assert.assertTrue(debugString, this.tcp.inputBufferIsEmpty());
119    Assert.assertTrue(debugString, this.host.inputBufferIsEmpty());
120    Assert.assertEquals(debugString, 0, this.events.size());
121 }

```

```
116     }  
117  
118  
119 }
```

```

1 package org.example.simulator.statistics;
2
3 import org.example.data.Message;
4 import org.example.data.Packet;
5 import org.example.network.Channel;
6 import org.example.network.Routable;
7 import org.example.network.Router;
8 import org.example.network.address.SimpleAddress;
9 import org.example.protocol.ClassicTCP;
10 import org.example.protocol.MPTCP;
11 import org.example.simulator.EventHandler;
12 import org.example.simulator.events.tcp.RunTCPEvent;
13 import org.example.simulator.events.tcp.TCPConnectEvent;
14 import org.example.util.Util;
15 import org.junit.Assert;
16 import org.junit.Before;
17 import org.junit.Test;
18
19 public class TCPStatsTest {
20
21     @Before
22     public void setup() {
23         Util.resetTime();
24         Util.setSeed(1);
25     }
26
27     @Test
28     public void floodWithPacketsInBigCongestedNetworkShouldWorkTest() {
29         ClassicTCP client = new ClassicTCP.ClassicTCPBuilder().
30             withReceivingWindowCapacity(10).setReno().withAddress(new SimpleAddress(
31                 "Client")).build();
32         Routable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.91).
33             withAddress(new SimpleAddress("Router 1")).build();
34         Routable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.91).
35             withAddress(new SimpleAddress("Router 2")).build();
36         Routable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.91).
37             withAddress(new SimpleAddress("Router 3")).build();
38         Routable r4 = new Router.RouterBuilder().withAverageQueueUtilization(0.91).
39             withAddress(new SimpleAddress("Router 4")).build();
40         ClassicTCP server = new ClassicTCP.ClassicTCPBuilder().
41             withReceivingWindowCapacity(30).setReno().withAddress(new SimpleAddress(
42                 "Server")).build();
43
44         new Channel.ChannelBuilder().build(client, r1);
45         new Channel.ChannelBuilder().build(r1, r2);
46         new Channel.ChannelBuilder().withLoss(0.01).build(r2, r3);
47         new Channel.ChannelBuilder().build(r3, r4);
48         new Channel.ChannelBuilder().build(r4, server);
49
50         client.updateRoutingTable();
51         r1.updateRoutingTable();
52         r2.updateRoutingTable();
53         r3.updateRoutingTable();
54         r4.updateRoutingTable();
55         server.updateRoutingTable();
56
57         EventHandler eventHandler = new EventHandler();
58         eventHandler.addEvent(new TCPConnectEvent(client, server));
59         eventHandler.run();
60
61         Assert.assertTrue(client.isConnected());
62     }
63 }

```

```

55     Assert.assertTrue(server.isConnected());
56
57     System.out.println("connected");
58
59     Assert.assertTrue(client.inputBufferIsEmpty());
60     Assert.assertTrue(server.inputBufferIsEmpty());
61     Assert.assertTrue(client.outputBufferIsEmpty());
62     Assert.assertTrue(server.outputBufferIsEmpty());
63
64     int numPacketsToSend = 10000;
65     for (int i = 1; i <= numPacketsToSend; i++) {
66         Message msg = new Message("test " + i);
67         client.send(msg);
68     }
69
70     eventHandler.addEvent(new RunTCPEvent(client));
71     eventHandler.run();
72
73     Assert.assertTrue(client.inputBufferIsEmpty());
74     Assert.assertTrue(server.inputBufferIsEmpty());
75     Assert.assertTrue(client.outputBufferIsEmpty());
76     Assert.assertTrue(server.outputBufferIsEmpty());
77
78     for (int i = 1; i <= numPacketsToSend; i++) {
79         Message msg = new Message("test " + i);
80         Packet received = server.receive();
81         Assert.assertNotNull(received);
82         Assert.assertEquals(msg, received.getPayload());
83     }
84     TCPStats stat = server.getStats();
85     System.out.println(stat);
86     stat.createArrivalChart();
87     stat.createDepartureChart();
88     stat.createInterArrivalChart();
89     stat.createTimeInSystemChart();
90     stat.createNumberOfPacketsInSystemChart();
91
92 }
93
94 @Test
95 public void MPTCPFloodWithPacketsInOrderShouldWorkTest() {
96     MPTCP client = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
97         withReceivingWindowCapacity(60).withAddress(new SimpleAddress("MPTCP-
98         Client")).build();
99     Routable r1 = new Router.RouterBuilder().withAverageQueueUtilization(0.94).
100         withAddress(new SimpleAddress("A")).build();
101     Routable r2 = new Router.RouterBuilder().withAverageQueueUtilization(0.94).
102         withAddress(new SimpleAddress("B")).build();
103     Routable r3 = new Router.RouterBuilder().withAverageQueueUtilization(0.9).
104         withAddress(new SimpleAddress("C")).build();
105     MPTCP server = new MPTCP.MPTCPBuilder().withNumberOfSubflows(3).
106         withReceivingWindowCapacity(60).withAddress(new SimpleAddress("MPTCP-
107         Server")).build();
108
109     //path one
110     new Channel.ChannelBuilder().build(client, r1);
111     new Channel.ChannelBuilder().withLoss(0.015).build(r1, server);
112
113     //path two
114     new Channel.ChannelBuilder().build(client, r2);
115     new Channel.ChannelBuilder().withLoss(0.015).build(r2, server);
116
117     //path three

```

```

111     new Channel.ChannelBuilder().build(client, r3);
112     new Channel.ChannelBuilder().withLoss(0.005).build(r3, server);
113
114     client.updateRoutingTable();
115     r1.updateRoutingTable();
116     r2.updateRoutingTable();
117     r3.updateRoutingTable();
118     server.updateRoutingTable();
119
120     EventHandler eventHandler = new EventHandler();
121     eventHandler.addEvent(new TCPConnectEvent(client, server));
122     eventHandler.run();
123     System.out.println("connected");
124
125     Assert.assertTrue(client.getSubflows()[0].isConnected());
126     Assert.assertTrue(client.getSubflows()[1].isConnected());
127     Assert.assertTrue(client.getSubflows()[2].isConnected());
128
129     int numPacketsToSend = 10000;
130
131     for (int i = 1; i <= numPacketsToSend; i++) {
132         Message msg = new Message("test " + i);
133         client.send(msg);
134     }
135     eventHandler.addEvent(new RunTCPEvent(client));
136     eventHandler.run();
137
138     Assert.assertTrue("client still has packets to send", client.
139         outputBufferIsEmpty());
140     Assert.assertTrue(server.outputBufferIsEmpty());
141     Assert.assertTrue(client.inputBufferIsEmpty());
142     Assert.assertTrue(server.inputBufferIsEmpty());
143     Assert.assertTrue(r1.inputBufferIsEmpty());
144     Assert.assertTrue(r2.inputBufferIsEmpty());
145     Assert.assertTrue(r3.inputBufferIsEmpty());
146
147     for (int i = 1; i <= numPacketsToSend; i++) {
148         Message msg = new Message("test " + i);
149         Packet received = server.receive();
150         Assert.assertNotNull(received);
151         Assert.assertEquals("iteration " + i, received.getPayload(), msg);
152     }
153     Assert.assertNull(server.receive());
154
155     //receiver
156     for (TCPStats stat : server.getTcpStats()) {
157         System.out.println(stat.toString());
158         stat.createArrivalChart();
159         stat.createDepartureChart();
160         stat.createInterArrivalChart();
161         stat.createTimeInSystemChart();
162         stat.createNumberOfPacketsInSystemChart();
163     }
164
165     //sender
166     for (TCPStats stat : client.getTcpStats()) {
167         System.out.println(stat.toString());
168         stat.createCWNDChart();
169     }
170
171 }
172

```



```

1 package org.example.util;
2
3 import org.example.data.Packet;
4 import org.example.data.PacketBuilder;
5 import org.junit.Assert;
6 import org.junit.Test;
7
8 public class BoundedPriorityBlockingQueueTest {
9
10
11     @Test
12     public void bothConstructorsWorkTest() {
13         BoundedQueue bq1 = new BoundedPriorityBlockingQueue(10);
14         BoundedQueue<Packet> bq2 = new BoundedPriorityBlockingQueue<>(10, (packet,
15             t1) -> packet.getSequenceNumber() - t1.getSequenceNumber());
16         Assert.assertTrue(bq1 instanceof BoundedPriorityBlockingQueue);
17         Assert.assertTrue(bq2 instanceof BoundedPriorityBlockingQueue);
18     }
19
20     @Test
21     public void offerInsertsOrdered() {
22         BoundedPriorityBlockingQueue<Packet> boundedPriorityBlockingQueue =
23             new BoundedPriorityBlockingQueue<>(5, (packet, t1) -> packet.
24                 getSequenceNumber() - t1.getSequenceNumber());
25
26         Packet one = new PacketBuilder()
27             .withSequenceNumber(1)
28             .build();
29         Packet two = new PacketBuilder()
30             .withSequenceNumber(2)
31             .build();
32         Packet three = new PacketBuilder()
33             .withSequenceNumber(3)
34             .build();
35         Packet four = new PacketBuilder()
36             .withSequenceNumber(4)
37             .build();
38         Packet five = new PacketBuilder()
39             .withSequenceNumber(5)
40             .build();
41
42         boundedPriorityBlockingQueue.offer(three);
43         boundedPriorityBlockingQueue.offer(four);
44         boundedPriorityBlockingQueue.offer(one);
45         boundedPriorityBlockingQueue.offer(five);
46         boundedPriorityBlockingQueue.offer(two);
47
48         Assert.assertEquals(one, boundedPriorityBlockingQueue.poll());
49         Assert.assertEquals(two, boundedPriorityBlockingQueue.poll());
50         Assert.assertEquals(three, boundedPriorityBlockingQueue.poll());
51         Assert.assertEquals(four, boundedPriorityBlockingQueue.poll());
52         Assert.assertEquals(five, boundedPriorityBlockingQueue.poll());
53         Assert.assertEquals(null, boundedPriorityBlockingQueue.poll());
54     }
55
56     @Test
57     public void isFullWorksReturnsTrueIfFullTest() {
58         BoundedPriorityBlockingQueue<Packet> boundedPriorityBlockingQueue =
59             new BoundedPriorityBlockingQueue<>(5, (packet, t1) -> packet.
60                 getSequenceNumber() - t1.getSequenceNumber());

```



```

60     Packet one = new PacketBuilder()
61         .withSequenceNumber(1)
62         .build();
63     Packet two = new PacketBuilder()
64         .withSequenceNumber(2)
65         .build();
66     Packet three = new PacketBuilder()
67         .withSequenceNumber(3)
68         .build();
69     Packet four = new PacketBuilder()
70         .withSequenceNumber(4)
71         .build();
72     Packet five = new PacketBuilder()
73         .withSequenceNumber(5)
74         .build();
75
76     boundedPriorityBlockingQueue.offer(one);
77     boundedPriorityBlockingQueue.offer(two);
78     boundedPriorityBlockingQueue.offer(three);
79     boundedPriorityBlockingQueue.offer(four);
80     boundedPriorityBlockingQueue.offer(five);
81
82     Assert.assertTrue(boundedPriorityBlockingQueue.isFull());
83 }
84
85
86 @Test
87 public void isFullWorksReturnsFalseIfNotFullTest() {
88     BoundedPriorityBlockingQueue<Packet> boundedPriorityBlockingQueue =
89         new BoundedPriorityBlockingQueue<>(5, (packet, t1) -> packet.
90             getSequenceNumber() - t1.getSequenceNumber());
91
92     Packet one = new PacketBuilder()
93         .withSequenceNumber(1)
94         .build();
95     Packet two = new PacketBuilder()
96         .withSequenceNumber(2)
97         .build();
98     Packet three = new PacketBuilder()
99         .withSequenceNumber(3)
100        .build();
101     Packet four = new PacketBuilder()
102         .withSequenceNumber(4)
103         .build();
104
105     boundedPriorityBlockingQueue.offer(one);
106     boundedPriorityBlockingQueue.offer(two);
107     boundedPriorityBlockingQueue.offer(three);
108     boundedPriorityBlockingQueue.offer(four);
109
110     Assert.assertFalse(boundedPriorityBlockingQueue.isFull());
111 }
112
113 @Test
114 public void resizeWorksTest() {
115     BoundedPriorityBlockingQueue<Packet> boundedPriorityBlockingQueue =
116         new BoundedPriorityBlockingQueue<>(5, (packet, t1) -> packet.
117             getSequenceNumber() - t1.getSequenceNumber());
118
119     boundedPriorityBlockingQueue.setBound(10);
120     Assert.assertEquals(10, boundedPriorityBlockingQueue.bound());
121 }

```

```

121
122     @Test
123     public void resizeAllowsMoreElementsInQueue() {
124         BoundedPriorityBlockingQueue<Packet> boundedPriorityBlockingQueue =
125             new BoundedPriorityBlockingQueue<>(4, (packet, tl) -> packet.
126                 getSequenceNumber() - tl.getSequenceNumber());
127         Packet one = new PacketBuilder()
128             .withSequenceNumber(1)
129             .build();
130         Packet two = new PacketBuilder()
131             .withSequenceNumber(2)
132             .build();
133         Packet three = new PacketBuilder()
134             .withSequenceNumber(3)
135             .build();
136         Packet four = new PacketBuilder()
137             .withSequenceNumber(4)
138             .build();
139         boundedPriorityBlockingQueue.offer(one);
140         boundedPriorityBlockingQueue.offer(two);
141         boundedPriorityBlockingQueue.offer(three);
142         boundedPriorityBlockingQueue.offer(four);
143
144         Assert.assertTrue(boundedPriorityBlockingQueue.isFull());
145
146         boundedPriorityBlockingQueue.setBound(8);
147         Assert.assertEquals(8, boundedPriorityBlockingQueue.bound());
148
149
150         Packet five = new PacketBuilder()
151             .withSequenceNumber(5)
152             .build();
153         Packet six = new PacketBuilder()
154             .withSequenceNumber(6)
155             .build();
156         Packet seven = new PacketBuilder()
157             .withSequenceNumber(7)
158             .build();
159         Packet eight = new PacketBuilder()
160             .withSequenceNumber(8)
161             .build();
162
163         boundedPriorityBlockingQueue.offer(five);
164         boundedPriorityBlockingQueue.offer(six);
165         boundedPriorityBlockingQueue.offer(seven);
166         boundedPriorityBlockingQueue.offer(eight);
167
168         Assert.assertTrue(boundedPriorityBlockingQueue.isFull());
169
170
171         Packet nine = new PacketBuilder()
172             .withSequenceNumber(9)
173             .build();
174         boundedPriorityBlockingQueue.offer(nine);
175
176         Assert.assertTrue(boundedPriorityBlockingQueue.isFull());
177         Assert.assertEquals(8, boundedPriorityBlockingQueue.size());
178     }
179
180 }

```

```

1 package org.example.util;
2
3 import org.junit.Assert;
4 import org.junit.Test;
5
6 import java.util.ArrayList;
7
8 public class UtilTest {
9
10
11     @Test
12     public void getGaussianIsStable() {
13         ArrayList<Double> doubles = new ArrayList<>();
14         Util.setSeed(1337);
15
16         int arraySize = 10000;
17         for (int i = 0; i < arraySize; i++) {
18             doubles.add(Util.getNextGaussian());
19         }
20
21         Assert.assertEquals(doubles.size(), arraySize);
22
23         Util.setSeed(1337);
24         for (double gaussian : doubles) {
25             Assert.assertEquals(gaussian, Util.getNextGaussian(), 0);
26         }
27     }
28
29
30     @Test
31     public void getNextRandomIsStable() {
32         ArrayList<Integer> ints = new ArrayList<>();
33         Util.setSeed(1337);
34
35         int arraySize = 10000;
36         for (int i = 0; i < arraySize; i++) {
37             ints.add(Util.getNextRandomInt());
38         }
39
40         Assert.assertEquals(ints.size(), arraySize);
41
42         Util.setSeed(1337);
43         for (double rand : ints) {
44             Assert.assertEquals(rand, Util.getNextRandomInt(), 0);
45         }
46     }
47
48
49     @Test
50     public void getNextRandomWithBoundIsStable() {
51         ArrayList<Integer> ints = new ArrayList<>();
52         int bound = 100;
53         Util.setSeed(1337);
54
55         int arraySize = 10000;
56         for (int i = 0; i < arraySize; i++) {
57             ints.add(Util.getNextRandomInt(bound));
58         }
59
60         Assert.assertEquals(ints.size(), arraySize);
61
62         Util.setSeed(1337);

```

```

63     for (double rand : ints) {
64         Assert.assertEquals(rand, Util.getNextRandomInt (bound), 0);
65     }
66 }
67
68
69 @Test
70 public void allIsStable() {
71     ArrayList<Integer> ints = new ArrayList<>();
72     ArrayList<Integer> boundedInts = new ArrayList<>();
73     ArrayList<Double> doubles = new ArrayList<>();
74     final int bound = 100;
75     Util.setSeed(1337);
76
77     int arraySize = 10000;
78     for (int i = 0; i < arraySize; i++) {
79         ints.add(Util.getNextRandomInt());
80         boundedInts.add(Util.getNextRandomInt (bound));
81         doubles.add(Util.getNextGaussian());
82     }
83
84     Assert.assertEquals(ints.size(), arraySize);
85     Assert.assertEquals(boundedInts.size(), arraySize);
86     Assert.assertEquals(doubles.size(), arraySize);
87
88     Util.setSeed(1337);
89     for (int i = 0; i < arraySize; i++) {
90         Assert.assertEquals(ints.get(i), Util.getNextRandomInt(), 0);
91         Assert.assertEquals(boundedInts.get(i), Util.getNextRandomInt (bound), 0)
92             ;
93         Assert.assertEquals(doubles.get(i), Util.getNextGaussian(), 0);
94     }
95 }
96
97 }

```