

# Communication Networks

Prof. Laurent Vanbever

## Communication Networks

Spring 2018



Tobias Bühler, TA  
Thomas Holterbach, TA  
Maximilian Schütte, TA  
Alexander Dietmüller, TA  
<http://comm-net.ethz.ch/>

ETH Zürich  
May 7 2018

Last week on  
Communication Networks

## TCP Congestion Control



Congestion control aims at  
solving three problems

- #1 **bandwidth estimation** How to adjust the bandwidth of a single flow to the bottleneck bandwidth?  
could be 1 Mbps or 1 Gbps...
- #2 **bandwidth adaptation** How to adjust the bandwidth of a single flow to variation of the bottleneck bandwidth?
- #3 **fairness** How to share bandwidth "fairly" among flows, without overloading the network

Congestion control differs from flow control  
both are provided by TCP though

- Flow control prevents **one fast sender** from overloading **a slow receiver**
- Congestion control prevents **a set of senders** from overloading **the network**

The sender adapts its sending rate  
based on these two windows

- Receiving Window **RWND** How many bytes can be sent without overflowing the receiver buffer?  
based on the receiver input
- Congestion Window **CWND** How many bytes can be sent without overflowing the routers?  
based on network conditions
- Sender Window **minimum(CWND, RWND)**

## The 2 key mechanisms of Congestion Control

detecting  
congestion

reacting to  
congestion



## The 2 key mechanisms of Congestion Control



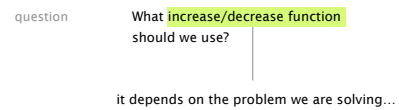
Detecting losses can be done using ACKs or timeouts, the two signal differ in their degree of severity



## The 2 key mechanisms of Congestion Control



TCP approach is to **gently increase** when not congested and to **rapidly decrease** when congested

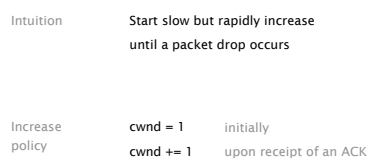


## Congestion control aims at solving three problems

- #1 **bandwidth estimation** How to adjust the bandwidth of a single flow to the bottleneck bandwidth?  
could be 1 Mbps or 1 Gbps...
- #2 **bandwidth adaptation** How to adjust the bandwidth of a single flow to variation of the bottleneck bandwidth?
- #3 **fairness** How to share bandwidth "fairly" among flows, without overloading the network

- #1 **bandwidth estimation** How to adjust the bandwidth of a single flow to the bottleneck bandwidth?  
could be 1 Mbps or 1 Gbps...

## Initially, you want to quickly get a first-order estimate of the available bandwidth



- #2 **bandwidth adaptation** How to adjust the bandwidth of a single flow to variation of the bottleneck bandwidth?

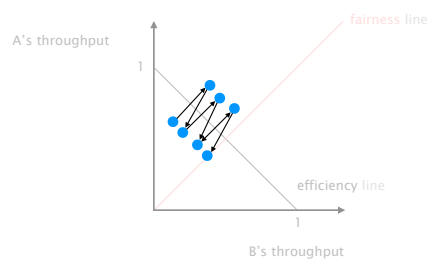


	increase behavior	decrease behavior
AIAD	gentle	gentle
AIMD	gentle	aggressive
MIAD	aggressive	gentle
MIMD	aggressive	aggressive

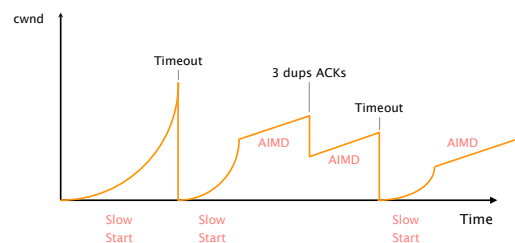
#3 **fairness**

How to share bandwidth "fairly" among flows,  
without overloading the network

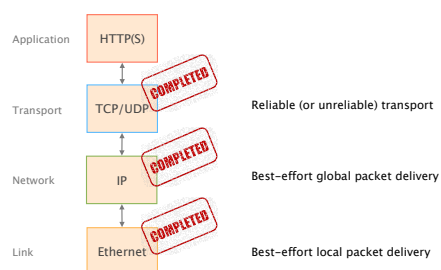
AIMD converge to fairness and efficiency,  
it then fluctuates around the optimum (in a stable way)



Congestion control makes TCP throughput  
look like a "sawtooth"



We now have completed **the transport layer (!)**



**This week on**  
**Communication Networks**

Routing Project

Reliable Transport Project

Recap, demo  
and final results

Introduction  
and demo

Python and Git  
tutorial

Routing Project

Reliable Transport Project

Recap, demo  
and final results

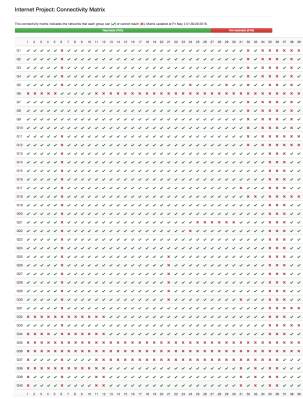
Introduction  
and demo

Python and Git  
tutorial



## Communication Networks 2018 Routing Project

Recap



73%

Proportion of valid BGP paths  
in your mini-Internet

\*From 15 traceroutes launched  
between random pairs of ASes

Your mini-Internet works!

and common services  
can run on top of it

For this project, you basically did what an  
actual network operator has to do

Including debugging and monitoring  
your configuration and connectivity

For this project, you basically did what an  
actual network operator has to do

Including debugging and monitoring  
your configuration and connectivity

Looking glass

Measurement platform

BGP is operating in STANDALONE mode.

Process	As/To/From	ASes/REB	Label/Ver	Import/Ver	Send/To/Ver	Standby/Ver
Speaker	37/200401	37/200401	37/200401	37/200401	37/200401	37/200401
Router	504	40	10000	10000	10000	10000
62.46.124.21	0	30900	1200127	452000	37/200401	0
98.249.288.30	0	4000	200000	100000	37/200401	0
98.249.288.36	0	8074	100070	100000	37/200401	0
98.249.288.60	0	6000	100000	100000	37/200401	0
98.249.288.63	0	9000	122004	100000	37/200401	0
98.249.288.65	0	5040	100000	100000	37/200401	0
98.249.288.91	0	6000	100000	100000	37/200401	0



There was often multiple ways  
to answer the questions

There was often multiple ways  
to answer the questions

and we found some interesting answers



### Enabling authentication in OSPF

```
ip address 192.168.56.1/24
ip ospf authentication message-digest
ip ospf cost 100
ip ospf message-digest-key 1 md5 08202987c5e62d101e7526459f26d5eb
ipv6 nd suppress-ra
```

### Group BGP neighbors to simplify configuration

```
neighbor internal peer-group
neighbor internal remote-as 65000
neighbor internal password 08202987c5e62d101e7526459f26d5eb
neighbor internal update-source lo
neighbor internal next-hop-self
neighbor peer-group internal
neighbor peer-group internal
neighbor peer-group internal
neighbor peer-group internal
neighbor peer-group internal
neighbor peer-group internal
neighbor peer-group internal
```

### Multiple valid answers for question 3.3

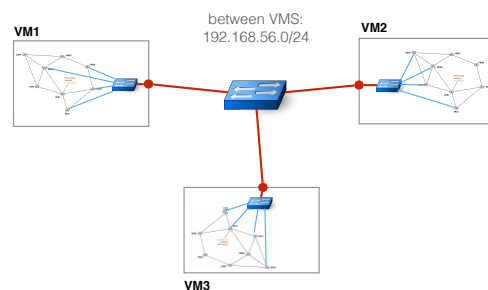
#### More specific advertisements

```
router bgp
bgp router-id 1.0.1.2
network 192.0.0.0/9
network 192.0.0.0/9
```

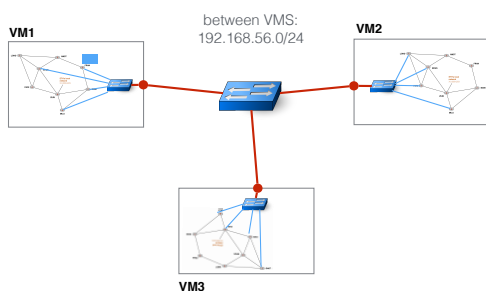
#### AS path prepending

```
route-map AMST-out permit 1
match community 2
set as-path prepend 10 10 10
```

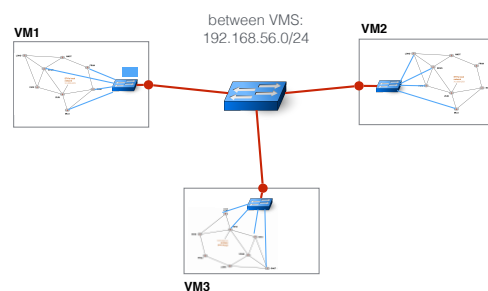
### How we have built the mini-Internet



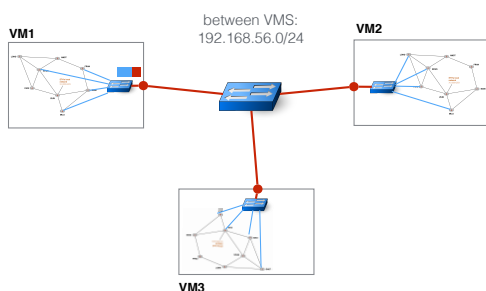
### How we have built the mini-Internet



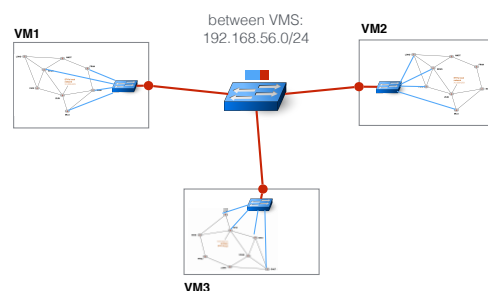
### How we have built the mini-Internet



### How we have built the mini-Internet

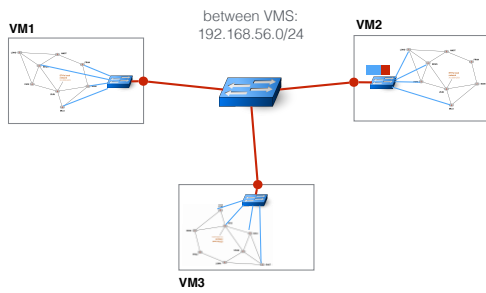


### How we have built the mini-Internet

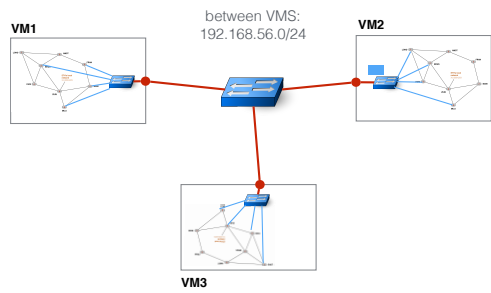




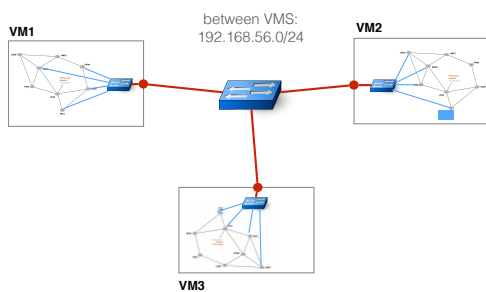
### How we have built the mini-Internet



### How we have built the mini-Internet



### How we have built the mini-Internet



### Communication Networks 2018 Routing Project

Except the grades within ~2weeks from now

#### Routing Project

Recap, demo  
and final results

#### Reliable Transport Project

Introduction  
and demo

Python and Git  
tutorial

### Implement your own **Reliable** Transport Protocol

recover from packet loss  
and reordering

### Implement your own **Reliable** Transport Protocol

recover from packet loss  
and reordering

#### Part 1

Simple Go-Back-N implementation  
Retransmit all packets after a timeout

#### Part 2

Support for Selective Repeat  
Fast retransmission after repeated ACKs

#### Part 3

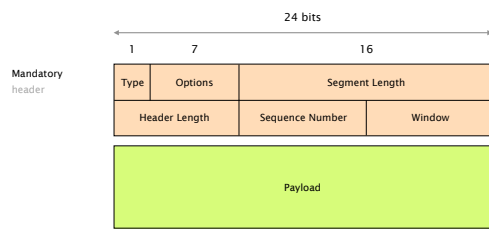
Support for Selective Acknowledgements (SACK)  
SACK contains blocks of correctly received segments

Let's see how the **final** sender and receiver  
should look like

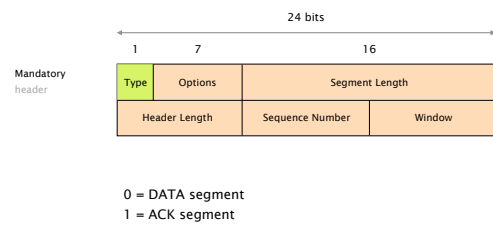




The header of our Go-Back-N protocol  
is 6 bytes long

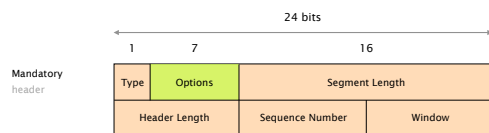


The header of our Go-Back-N protocol  
is 6 bytes long



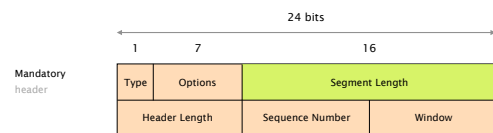
0 = DATA segment  
1 = ACK segment

The header of our Go-Back-N protocol  
is 6 bytes long



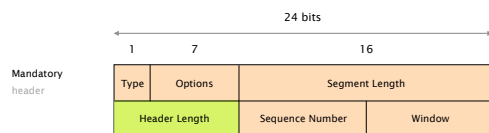
0 = no SACK support  
1 = SACK support

The header of our Go-Back-N protocol  
is 6 bytes long



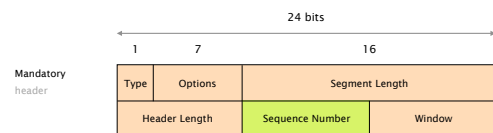
Length of the payload. Normally, 64 bytes.  
Only last segment could be smaller

The header of our Go-Back-N protocol  
is 6 bytes long



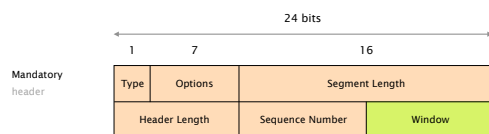
Total length of the header.  
In bytes

The header of our Go-Back-N protocol  
is 6 bytes long



In DATA: segment sequence number. Starts at 0  
In ACK: next expected in-sequence segment

The header of our Go-Back-N protocol  
is 6 bytes long



Sender respectively receiver window size.  
In number of segments

Sequence number overflow

NBITS  
maximum  
overflow  
application  
examples



NBITS	controls the maximum sequence number
maximum	assuming NBITS=3: $2^{\text{NBITS}} - 1 = 7$
overflow	... 5, 6, 7, 0, 1, 2, ...
application examples	ACK number, SACK header blocks, retransmission, ...

Sent segments:	0	1	2	3	4	5	
Receiver behavior:	0	-	2	3	4	5	Out-of-order segments are dropped
Sent ACKs:	1	-	1	1	1	1	
Retransmission:							

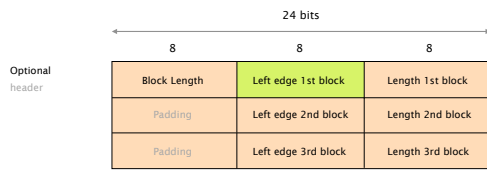
Sent segments:	0	1	2	3	4	5	
Receiver behavior:	0	-	2	3	4	5	Out-of-order segments are buffered
Sent ACKs:	1	-	1	1	1	1	
Retransmission:							

	24 bits		
	8	8	8
Optional header	Block Length	Left edge 1st block	Length 1st block
	Padding	Left edge 2nd block	Length 2nd block
	Padding	Left edge 3rd block	Length 3rd block

8 of 16

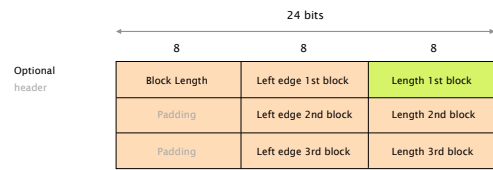


For SACK we need an **optional** header



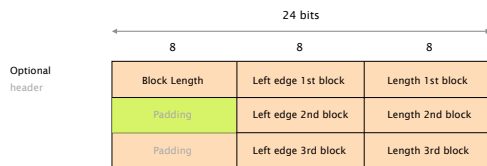
Start of the first block

For SACK we need an **optional** header



Length of the first block. In number of segments  
A block with one segment has size 1

For SACK we need an **optional** header



Padding for better alignment

SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header:

SACK header:

SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header: ACK number: 3

SACK header:

SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header:

ACK number: 3

SACK header:

#blocks	start b1	size b1
Padding	start b2	size b2
Padding	start b3	size b3

SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header: ACK number: 3

SACK header:

#blocks	4	2
Padding	start b2	size b2
Padding	start b3	size b3

SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header:

ACK number: 3

SACK header:

#blocks	4	2
Padding	8	1
Padding	start b3	size b3



### SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header: ACK number: 3

SACK header:

#blocks	4	2
Padding	8	1
Padding	10	4

### SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, ~~15~~, ~~16~~, ~~17~~ no space

Mandatory header:

ACK number: 3

SACK header:

#blocks	4	2
Padding	8	1
Padding	10	4

### SACK example - Receiver

Correctly received segments: 0, 1, 2

Buffered out-of-order segments: 4, 5, 8, 10, 11, 12, 13, 15, 16, 17

Mandatory header: ACK number: 3

SACK header:

3	4	2
Padding	8	1
Padding	10	4

### SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1:

block 1 - block 2:

block 2 - block 3:

after block 3:

### SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1:

3

block 1 - block 2:

block 2 - block 3:

after block 3:

### SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1:

3

block 1 - block 2:

6, 7

block 2 - block 3:

after block 3:

### SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1:

3

block 1 - block 2:

6, 7

block 2 - block 3:

9

after block 3:

### SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1:

3

block 1 - block 2:

6, 7

block 2 - block 3:

9

after block 3:

no retransmission



## SACK example - Sender

Receiver SACK header:

3	4	2
Padding	8	1
Padding	10	4

ACK number: 3

ACK - block 1: 3  
block 1 - block 2: 6, 7  
block 2 - block 3: 9  
after block 3: no retransmission  
important: sender window is not moved

## To test your implementation...

- ... run your sender against your receiver
- ... test with the implementation of another group
- ... **optionally**, use our test framework

Ask your questions on **Slack (#transport\_project)**  
or visit an **exercise session**

Tobias Bühler (@buehlert)  
Maximilian Schütte (@Maximilian (TA))  
Alexander Dietmüller (@Alexander (TA))  
Rüdiger Birkner (@rbirkner)  
Roland Meier (@roland)  
Thomas Holterbach (@thomas\_holterbach)

**Next week** on  
Communication Networks

**This Thursday:** Ascension Day  
**Monday:** Applications: DNS and HTTP

Routing Project

Reliable Transport Project

Recap, demo  
and final results

Introduction  
and demo

Python and Git  
tutorial

## The Hitchhiker's Guide to Efficient Python Development

Communication Networks  
Spring 2018  
ETH Zürich

## Contents

#Why we use Python

#Stop wasting time: Editors, Linters, File Sync

#Get to know the framework

#Avoiding Catastrophe: Version Control

#Git made easy: GitLab and SourceTree



## Python

Slither along with your friendly neighbourhood snake!



## Reasons to choose Python

#Interpreted Language

#Many packages available

#Simple yet powerful Syntax / Beginner Friendly

#Often used in academia and science

## Learn the Basics BEFORE You Start!

We promise the basics will pay off...

## Learn the Basics BEFORE You Start!

#One afternoon on [learnpython.org](http://learnpython.org) should suffice

#If you skip the preparation, bugs may go unnoticed and cost you points

#Also you will spend much more time on debugging than you would have to learn the python basics

## Learning Python for Pros

<https://learnxinyminutes.com/docs/python3/>

## Learning Python for Everyone

#Interactive Getting Started Guide

# <http://www.learnpython.org/>

#Short Intro

# <https://developers.google.com/edu/python/>

#Not So Short Intro

# <http://thepythonuru.com/>

# <https://docs.python.org/3/tutorial/index.html>

#Detailed Intro

# <https://learnpythonthehardway.org/python3/>

#Free Video Series for Beginners

# <https://mya.microsoft.com/en-US/training-courses/introduction-to-programming-with-python-8360>

#Udemy Lecture for Beginners

# <https://www.udemy.com/complete-python-bootcamp/>

## Learning Python for Beginners

<http://www.learnpython.org/>

## Python 2.7 or 3.x?

#Python 2.7 is slowly dying

#Python 3.x is cleaner, better, faster, stronger...

#Details

# <https://wiki.python.org/moin/Python2orPython3>

# <https://www.dataquest.io/blog/python-2-or-3/>

# <https://www.digitalocean.com/community/tutorials/python-2-vs-python-3-practical-considerations-2>

# [http://sebastianraschka.com/Articles/2014\\_python\\_2\\_3\\_key\\_diff.html](http://sebastianraschka.com/Articles/2014_python_2_3_key_diff.html)

## Which Python Shall It Be?

Two major distributions to consider...





## Which Python Shall It Be?

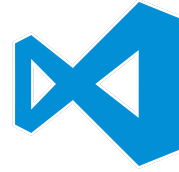
### CPython from python.org

- The „default“ distribution
- Is installed on the VMs
- Comes only with standard library
- Pip packet manager



### Anaconda by Continuum Analytics

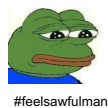
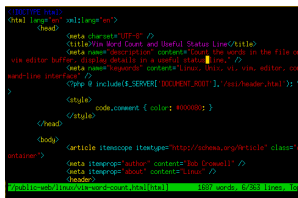
- Optimized for data science and large scale science apps
- Derived from CPython
- Ships with a big library of science related packages
- Uses Conda packet manager
- But also supports pip



## VSCode & PyLint

It's 2018, get your development workflow together!

## Editing on the console is cumbersome...



... but sometimes useful for quick fixes!

## Many good Python IDEs available!



Sublime Text



Visual Studio Code



Atom.io



JetBrains PyCharm



Eclipse PyDev

## Many good Python IDEs available!

#Any of the above will do, you the one you know and adapt it to the project!

#Top three are basic and can be used for many programming languages

#PyCharm is the most powerful Python IDE and even free for ETH students (professional edition)

## Integrated Development Environment Benefits

#Easy to set up and getting started

#Come with many supporting tools out of the box  
#IntelliSense, Syntax Checker / Linter, Auto completion...

#GUI based debugging is much faster and easier

## Linter

#A Linter performs static code analysis

#It points out...

#... errors in your code

#... redundant code

#... code that can be optimized

#... changes that improve the readability of your code

#Use it so you don't have to spend hours chasing typos!

## Secure File Transfer Protocol (SFTP)

#Available via extension for Visual Studio Code

#Makes transferring files from / to the VM super easy

#Extension shows you differences between local and vm code



## Demo Time!

#Install Python

#Install Visual Studio Code & Python / PyLint + sftp extension

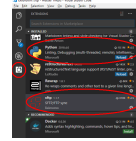
#Configure sftp & Download Project Files

#IntelliSense Demo

#CHECK SLACK FOR VIDEO DEMO! (to be released...)

## Step-by-Step Installation Reference

- # Install CPython 3.x or Anaconda / Miniconda 3.x
  - # <https://www.python.org/downloads/>
  - # <https://www.anaconda.com/download/>
- # Install Visual Studio Code
  - # <https://code.visualstudio.com/>
- # Start Visual Studio Code and click on the extensions icon on the left
- # Search for and install Python (ms-python.python) and sftp (liximomo.sftp)
- # Reload after BOTH installations have finished



## Configure Python and PyLint in VSCode

- # Press F1 and enter "Python: Select Interpreter"
- # Choose the python version that you just installed
  - # On Mac use the one in /usr/local, NOT the system installation!
- # Press F1 again and enter "Python: Select Linter" and choose "PyLint"
- # The first time you open a python file, you will receive a message box in the bottom right corner saying that PyLint is not installed. Press "install" to do so.
  - # On Mac, gcc will be installed if not installed already

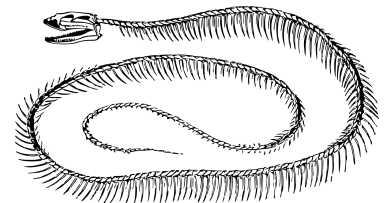
## Configure sftp and Download Code Reference

- # In VSCode, open a folder where you want your project files to be located.
- # Press F1 and enter "SFTP: Config"
- # A config file will pop up. Enter the details to your VM, as shown on the next slide. Providing a password is optional.
- # The config will be stored a subfolder **.vscode** and can be edited anytime.
- # Right click in the VSCode file browser and use the SFTP features like "download", "upload", or "sync".
- # In general, the plugin is conservative when it comes to «destructive» operations. See Extension Info page for more details.

## SFTP Example Config

```
{
  "protocol": "sftp",
  "host": "samichlaus.ethz.ch",
  "username": "root",
  "port": 3000+YOUR-GROUP-NUMBER,
  "remotePath": "./",
  "ignore": ["/.*"]
}
```

Don't forget this! It makes sure that you just copy the project related files!



## The Project Skeleton

You don't need to start from scratch...

## Sending and Receiving Packets in Python



## Sending and Receiving Packets in Python

```
from scapy.all import send, IP, TCP

Payload = b"This is some binary test data."

packet = IP(src="192.168.0.1", dst="8.8.8.8") / TCP() / payload

send(packet)
```

Combine headers with the division operator



## Sending and **Receiving** Packets in Python

Show summary and details

```
print(packet.summary())
```

```
print(packet.show())
```

Access headers and data

```
from scapy.all import IP
```

```
ip_header = packet.getlayer(IP)
```

```
source_address = ip_header.src
```

```
payload = ip_header.payload
```

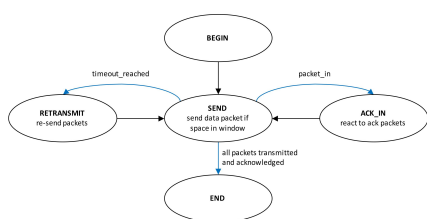
## Define **Your Own Header**

```
from scapy.all import Packet, bind_layers, BitEnumField, BitField
```

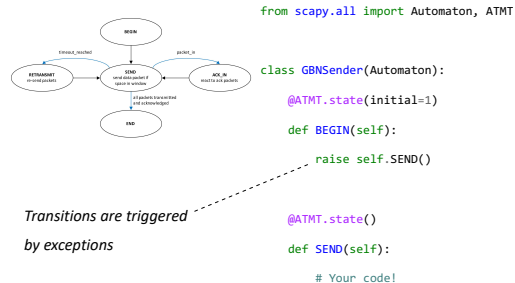
```
class GBN(Packet):
    name = 'GBN'
    fields_desc = [
        BitEnumField("type", 0, 1, {0: "data", 1: "ack"}),
        BitField("options", 0, 7),
        # other fields ...
    ]
```

```
# Tell Scapy where to look for the header when receiving a packet
bind_layers(IP, GBN, frag=0, proto=222)
```

## Our GBN Automaton is powered by Scapy



## Our GBN Automaton is powered by Scapy



## Where to start?

The GBN header is already defined...

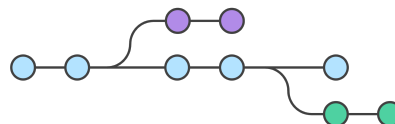
...you'll need to extend it in later questions

The automaton skeleton is fully implemented...

...no new states or transitions needed

The receiver already works for the first question...

...complete the sender, check receiver for inspiration



## Version Control

If two people are working on one problem, you get two problems...

## git Tracks Changes in Source Code



## Without git

Everyone works on the same file and uploads it to the server.

The version uploaded last **overwrites all other changes.**



## With git

Everyone works on the same file and pushes the changes to the git repository.

*All changes are combined, nothing is lost.*

## git Workflow



1. Create Repository
2. Invite Group Members



## git Workflow



`git clone <repository>`

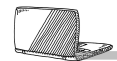


## git Workflow



`codecodecodecode...`

`git commit`



## git Workflow



1. `git pull`

2. `git push`



## git Workflow

`commit` → `pull` → `push`

Try it yourself and learn more:

<http://try.github.io/>

<https://backlog.com/git-tutorial/>



## SourceTree & Gitlab

... because no matter what they say, GUI matters.

## SourceTree

See Slack for Video Demo! (To be released)

