

Cristina Nita-Rotaru



7610: Distributed Systems

Introduction. Class Policy. Examples.



1: Intro

Welcome to cs7610

- ▶ Introducing the team
- ▶ How we will communicate
- ▶ **Take questions**
- ▶ Class policies
- ▶ **Take questions**
- ▶ Why distributed systems
- ▶ Class syllabus
- ▶ **Take questions**
- ▶ Docker and project zero
- ▶ **Take questions**

How we communicate in this class

- ▶ **ZOOM** – main class delivery channel
- ▶ **CANVAS** – main hub
- ▶ **PIAZZA** – main communication channel
- ▶ **Email in case of emergency - use cs7610 in subject**
c.nitarotaru@northeastern.edu

ZOOM (Link is in canvas and piazza)

- ▶ **Before class:**
 - ▶ Mute and no video before class start
 - ▶ Zoom opens at 1:15, join early
- ▶ **Start class:**
 - ▶ Please show your video for a few minutes to say hi
- ▶ **During class: you can close the video when I start teaching**
- ▶ **Questions: Please type them in ZOOM chat window**
 - ▶ Emergency for onsite students – you mention them to the Fengqi or somebody that can type in zoom chat
 - ▶ I will go through questions every 10-15 minutes
- ▶ **We will take a 5 minutes break after 50 minutes**

CANVAS

- ▶ It's a hub
- ▶ Have links there to class and piazza
- ▶ Add things that need to be protected
- ▶ Plan to have the grades there

PIAZZA

- ▶ Main communication environment where I will post
 - ▶ announcements
 - ▶ homework and projects
 - ▶ questions from class, etc
- ▶ You can post privately just to me and TA
- ▶ Public questions are anonymous to your colleagues
- ▶ **MAKE SURE YOU CHECK IT OFTEN**

How to ask on Piazza

- ▶ Read slides, notes, homework or project description
- ▶ Use #hashtags (#lecture2, #project3, #hw1, etc.)
- ▶ Describe the problem clearly, using the right terms
- ▶ Add code in attached files
- ▶ Add output from compiler or debugging information
- ▶ Add any other relevant information
- ▶ **Don't post solutions on piazza**
- ▶ **Anything that relates to solution post PRIVATELY**

OFFICE HOURS

- ▶ There are 2 hours of office hours every day with either me or the TA, also additional availability outside the allocated time

Schedule is in piazza post @15

<https://piazza.com/class/kdyjt8cld8c7ot?cid=15>

Exceptional situations

- ▶ Anything that impacts you and class please let me know
- ▶ We will accommodate the situation and find a solution

- ▶ I expect that deadlines will be difficult to make if you will be impacted by covid 19, so just let me know and we will work together to accommodate the situation

- ▶ **DO NOT WORRY !!!!**

Weather / Emergency

- ▶ In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control.

Academy integrity

- ▶ It is allowed to discuss homework problems before writing them down; however, **WRITING IS INDIVIDUAL**
 - ▶ if you look at another student's written or typed answers, or let another student look at your written or typed answers, that is considered cheating.
- ▶ It is allowed to discuss your project with your colleagues, but **DO NOT SHARE CODE**
- ▶ Never have a copy of someone else's homework or program in your possession and never give your homework (or password) or program to someone else.
- ▶ **NO CHEATING WILL BE TOLERATED.**

Individual meeting

- ▶ You are required to meet with me at least once per semester
- ▶ I will update piazza @15 with how to sign up to meet with me during office hours or outside office hours
- ▶ If needed you can set up additional appointments by sending me a private message on piazza

How to stay engaged during lecture and outside lecture

- ▶ Come to lecture, having a structure helps
- ▶ Take notes
- ▶ Ask questions
- ▶ Chat with colleagues
- ▶ Make appointments with colleagues to work together on homework and projects
 - ▶ They are individual but you can discuss them
- ▶ Ask/answer questions on piazza
- ▶ Meet with the TAs
- ▶ Final project is in teams (more about it next)

QUESTIONS:

Please type in zoom chat window

Why do we need distributed systems

- ▶ Distribute load
- ▶ Faster response by placing replicas closer to clients
- ▶ Increased resources, computation and storage
- ▶ Resilience to failures and attacks

What is a distributed system?

A distributed computing system is a set of computer programs executing on one or more computers and coordinating actions by exchanging messages.

A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable.

Attributed to Leslie Lamport

What do we expect from distributed systems

- ▶ **Reliability**: provide continuous service
- ▶ **Availability**: ready to use
- ▶ **Safety**: systems do what they are supposed to do, avoiding catastrophic consequences
- ▶ **Security**: withstands passive/active attacks from outsiders or insiders

...not easy to achieve because

- ▶ Computers and networks fail in many (often unpredictable) ways
- ▶ Computers get compromised
- ▶ Real-time constraints
- ▶ Performance requirements
- ▶ Complexity

Why do computer systems fail?

- ▶ *Why Do Computers Stop and What can be done about it? Jim Gray, 1985*
 - ▶ System administration (operator actions, system configuration and maintenance)
 - ▶ Software faults, environmental failures
 - ▶ Hardware failures (disks and communication controllers)
 - ▶ Power outages
- ▶ *Why do Internet services fail, and what can be done about it? D. Oppenheimer, A. Ganapathi and D.A. Patterson, 2003.*
 - ▶ Operator error (particularly configuration errors) is the leading cause of failures
 - ▶ Failures in custom-written front-end software
 - ▶ Not enough on-line testing

Why do computers get compromised?

- ▶ Software bugs
- ▶ Administration errors
- ▶ Lack of diversity, same vulnerability is exploited
- ▶ The explosion of the Internet facilitates the spread of malware
- ▶ Social engineering attacks

..how do computer systems fail...

- ▶ **Halting failures:** no way to detect except by using timeout
- ▶ **Fail-stop failures:** accurately detectable halting failures
- ▶ **Send-omission failures**
- ▶ **Receive-omission failures**
- ▶ **Network failures**
- ▶ **Network partitioning failures**
- ▶ **Timing failures:** temporal property of the system is violated
- ▶ **Byzantine failures:** arbitrary failures, include both benign and malicious failures

Example 1: Boeing 737 MAX and sensor inputs

- ▶ Disclaimer: this is not a detailed description of the MAX design, but an exemplification on how not following the fault-tolerance principle on sensor inputs can lead to severe problems.

Based on article by Gregory Travis

Redundancy in 737 design

- ▶ Boeing included the requisite redundancy in instrumentation and sensors, and flight computers – one on the pilot's side and one on the co-pilot's side
- ▶ The flight computers (among many other things)
 - ▶ act as the autopilot (i.e. fly the plane by computer) when commanded
 - ▶ make sure that the human pilots do not do anything wrong when the autopilot is not flying the plane

737MAX and MCAS review

- ▶ MCAS was put into the 737 MAX because the larger engines and their placement, make an aerodynamic stall more likely in a 737 MAX than in previous 737 models
- ▶ MCAS pushes the nose of the plane down when the MCAS system thinks the plane might exceed its angle of attack limits – in order to avoid an aerodynamic stall
- ▶ MCAS is implemented in the flight computer software.
 - ▶ When MCAS senses that the angle of attack is too high, it commands the aircraft's trim system (the system that makes the plane go up or down) to lower the nose
 - ▶ It pushes the pilot's control columns in the down direction

No redundancy on sensor input

- ▶ In the 737 MAX only one of the flight management computers is active at once
- ▶ **And that computer takes inputs ONLY from the sensors on the side of the aircraft corresponding to which flight computer is in control**
- ▶ **If one sensor has erroneous information it will incorrectly infer stall and push the nose down**
- ▶ **How was fault-tolerance achieved before?
Human in the loop – the pilot and co-pilot were able to see that the two sensors have different inputs and infer that something is wrong**


Example 2:

A network partition's impact on github


- ▶ What happened: A routine maintenance work to replace failing 100G optical equipment resulted in the loss of connectivity for **43 seconds** between POP US East Coast and DC East Coast.
- ▶ Result:
 - ▶ Degraded service for 24 hours and 11 minutes.
 - ▶ Multiple internal systems were affected which resulted in displaying of information that was out of date and inconsistent.
 - ▶ No user data was lost; but required manual reconciliation of a few seconds of database writes that took hours
 - ▶ For the majority of the incident, GitHub was also unable to serve webhook events or build and publish GitHub Pages sites.

<https://blog.github.com/2018-10-30-oct21-post-incident-analysis/>

Example 3: Amazon



At 6 AM ET error rates for the company's massive NoSQL database named [DynamoDB](#) began skyrocketing in AWS's US-East Virginia region - the oldest and largest of its nine global regions. By 7:52 AM ET, AWS determined the cause of the problems: an issue with how the database manages metadata had gone awry, impacting the service's partitions and tables.

 [RESOLVED] Increased API error rates

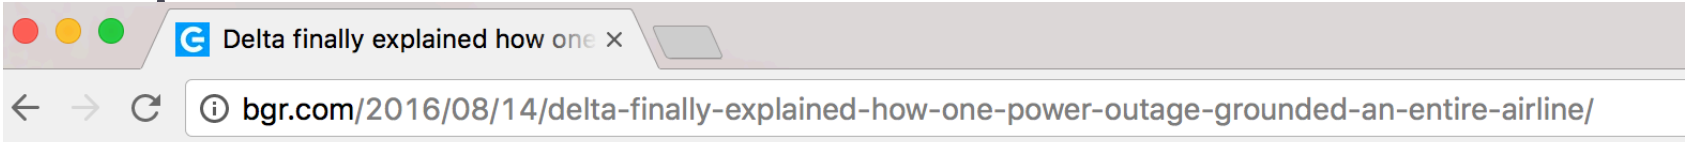
3:00 AM PDT We are investigating increased error rates for API requests in the US-EAST-1 Region.
3:26 AM PDT We are continuing to see increased error rates for all API calls in DynamoDB in US-East-1. We are actively working on resolving the issue.
4:05 AM PDT We have identified the source of the issue. We are working on the recovery.
4:41 AM PDT We continue to work towards recovery of the issue causing increased error rates for the DynamoDB APIs in the US-EAST-1 Region.
4:52 AM PDT We want to give you more information about what is happening. The root cause began with a portion of our metadata service within DynamoDB. This is an internal sub-service which manages table and partition information. Our recovery efforts are now focused on restoring metadata operations. We will be throttling APIs as we work on recovery.
5:22 AM PDT We can confirm that we have now throttled APIs as we continue to work on recovery.
5:42 AM PDT We are seeing increasing stability in the metadata service and continue to work towards a point where we can begin removing throttles.
6:19 AM PDT The metadata service is now stable and we are actively working on removing throttles.
7:12 AM PDT We continue to work on removing throttles and restoring API availability but are proceeding cautiously.
7:22 AM PDT We are continuing to remove throttles and enable traffic progressively.
7:40 AM PDT We continue to remove throttles and are starting to see recovery.
7:50 AM PDT We continue to see recovery of read and write operations and continue to work on restoring all other operations.
8:16 AM PDT We are seeing significant recovery of read and write operations and continue to work on restoring all other operations.
9:12 AM PDT Between 2:13 AM and 8:15 AM PDT we experienced high error rates for API requests in the US-EAST-1 Region. The issue has been resolved and the service is operating normally.

Amazon Web Services

Amazon Web Service's Health Dashboard shows

Because of the intricate interconnectivity of AWS's services, the issue snowballed to impact 34 total services (out of 117) that the company's [Service Health Dashboard](#) monitors. Everything from Elastic Compute Cloud (EC2) virtual machines to the Glacier storage service to its Relational Database Service were impacted. According to [media reports](#), other companies that rely on AWS experienced outages too, ranging from [Netflix to IMDB, to Tinder, Pocket and Buffer](#).

Example 4: Delta



Chris Mills  @chrisfills
August 14th, 2016 at 12:00 PM

 Share

 Tweet

[Earlier this week](#), Delta passengers worldwide were stranded as a computer failure completely screwed up operations. The ensuing chaos provided a good look at how the robots are actually going to kill us, but also raised some good questions: how does one power outage ground an airline, and how fired is the sysadmin?

The Week spoke to Delta's COO, Giles West, to try and understand what happened to take the entire network offline. It's a sad story of backups that should've worked, knock-on effects, and one seriously expensive outage.

DON'T MISS: [The iPhone 7 is going to be so much more exciting than you think](#)

"Monday morning a critical power control module at our Technology Command Center malfunctioned, causing a surge to the transformer and a loss of power," West said. "When this happened, critical systems and network equipment didn't switch over to backups. Other systems did. And now we're seeing instability in these systems," West told *The Week*.

In other words: a power surge caused by one malfunctioning piece of equipment tripped a power transformer, killing everything at Delta's command center in Atlanta. Clearly, this shouldn't have happened, and there should have been a backup power system in place (or an entire backup command system).

Examples of distributed systems

- ▶ Air Traffic Control
- ▶ Space Shuttle
- ▶ Banking Systems
- ▶ Grid Power Systems
- ▶ Cloud Computing



QUESTIONS:

Please type in zoom chat window



2: Syllabus and class policy

Looking under the hood of distributed systems

THEORY + SYSTEMS

- ▶ **Theory**
 - ▶ Fundamental problems
 - ▶ Algorithms solving this problems
 - ▶ Impossibility results
 - ▶ Trade-offs between properties provided by distributed systems
- ▶ **Systems**
 - ▶ What can go wrong when designing, implementing, testing and deploying a distributed service
 - ▶ Design of existing and popular software
 - ▶ Dependencies between different services

Course Information

- ▶ **Meetings**

- ▶ TuF 1:25-3:15 pm Sept. – Dec.
- ▶ Zoom opens at 1:15, join early

- ▶ **Office hours:**

- ▶ Office hours: Fri 3:30 – 5:30 pm see piazza post @15

- ▶ **Class webpage**

http://cnitarot.github.io/courses/ds_Fall_2020/index.html

- ▶ **Piazza for class communication**

- ▶ Use Piazza for questions and postings
- ▶ Hw and projects posted on piazza

Course overview

MODULE I – FUNDAMENTAL TOPICS

- ▶ Ordering events and distributed snapshots
 - ▶ Time in distributed systems. Clock synchronization. Global states and distributed snapshots. Detecting failures.
- ▶ Consensus
 - ▶ Synchronous systems, asynchronous systems, byzantine failures (including randomized solutions).
- ▶ Distributed commit and consistency models
 - ▶ 2PC and 3PC. Weak and strong consistency in partitioned database systems. Linearizability. CAP Theorem.

Course overview

MODULE II – ADVANCED TOPICS

- ▶ **Process Groups**
 - ▶ Leader election, membership, reliable multicast, virtual synchrony. Gossip protocols.
- ▶ **Quorums**
 - ▶ Paxos. Viewstamped replication. BFT.
- ▶ **Peer-to-peer systems**
 - ▶ File sharing, lookup services, streaming, publish-subscribe

Course overview

MODULE III – SYSTEMS

- ▶ **Files systems**
 - GFS, HDFS
- ▶ **Databases**
 - ▶ BigTable, HBase, Spanner, DynamoDB, Casandra
- ▶ **Lock services**
 - ▶ Chubby, Zookeeper, Zab
- ▶ **Computational services**
 - ▶ MapReduce, Spark

Course overview

MODULE III – SYSTEMS (cont.)

- ▶ **Distributed ledgers:**
 - ▶ Digital currency (BitCoin), smart contracts (Ethereum), credit systems (Ripple)
- ▶ **Infrastructure for ML**
 - ▶ TensorFlow, GraphLab
- ▶ **Microservices**
 - ▶ AWS Lambda

Reference Material

- ▶ **Textbooks**
 - ▶ Ken Birman: *Reliable Distributed Systems*
- ▶ **Recommended reading**
 - ▶ Research papers that will be specified for each lecture

Prerequisites

- ▶ Strong systems and networking background
- ▶ Socket programming
- ▶ Fluency in many languages
 - ▶ C/C++
 - ▶ Java
 - ▶ Go
 - ▶ Python or some other scripting language
- ▶ Linux command line proficiency
- ▶ Some computer security and cryptography fundamentals

Grading policy

- ▶ **Written assignments (3)** **30%**
 - ▶ **Programming projects (2)** **40%**
 - ▶ Do not include ungraded projects to get you started
 - ▶ **Final project** **30%**
-
- ▶ **There is no curve for grades**

Written assignments

- ▶ **Purpose of the written assignments is to make you understand the theoretical results discussed in class**
 - ▶ Read the material before solving them and solve them with closed books and notebooks
- ▶ 3 written theoretical assignments
- ▶ Homework is individual
- ▶ Homework must be typed – PDF submission format only
- ▶ For submission, see piazza post @10

Programming projects

- ▶ **Purpose of the programming projects is to help you understand practical aspects of things discussed in class**
 - ▶ Read all material in class and the description of the project in details before starting
- ▶ 2 programming projects
- ▶ Programming projects are individual
- ▶ All the code must be from scratch
- ▶ Follow the project description

Final project

- ▶ **Purpose of the final project is to help you understand some existing software or start a research project**
- ▶ You must work in teams of 2
 - ▶ Start looking for a partner at the beginning of the semester, don't wait till the project is assigned/chosen
- ▶ You can choose the final project, or I can assign one
- ▶ Project proposal presentation + 1 page description
- ▶ Final presentation in class + report of 3 pages submitted with the code and presentation

Late policy

- ▶ Each of you gets 5 LATE DAYS that can be used any way you want for homework and projects (but not the final project); you do not need to let me know if you plan to take any late day; just submit late
 - ▶ Keep track of your late days used
 - ▶ 20% off from grade obtained for that project or homework per day late
- ▶ Follow the requirements from project description to see how to submit
- ▶ **Assignments are due at 9:59:59 pm, no exceptions**
 - ▶ **1 second late = 1 hour late = 1 day late**

Regrading

- ▶ **YOU HAVE 1 WEEK to ASK for REGRADING** of a homework or project from the moment solutions were posted on piazza or discussed in class
- ▶ Make sure you read and understand the solution before asking for a regrade
- ▶ Request for a regrade will result in the regrading of the entire homework, project

Debugging distributed protocols

- ▶ They are known to be difficult to debug
- ▶ Write proactively – print all the info you send/receive over the network;
- ▶ Have state machine design before implementation and make sure you understand what your state machine is supposed to do before you implement your code
- ▶ Have message detailed description in design before implementation
- ▶ Focus on testcases to understand specific behavior
 - ▶ Delay, interleave, drop messages
 - ▶ Crash participants

One last word ...

- ▶ **No meetings will be accepted with the TA or instructor the day homework or projects are due**
- ▶ Start early, plan carefully
- ▶ Develop your solution gradually, test gradually so you always have functionality for which you can receive a grade; **YOUR CODE MUST WORK**
- ▶ Do not wait to submit your code last minute

Required Reading

- ▶ Chapter 1 and 2 from *Reliable Distributed Systems*
- ▶ Why do Internet services fail, and what can be done about it? D. Oppenheimer, A. Ganapathi and D. A. Patterson, 2003.
- ▶ Why Do Computers Stop and What can be done about it? Jim Gray, 1985.



QUESTIONS:

Please type in zoom chat window



Containers

Basic docker commands

Why containerize applications

-- build once, run anywhere --

- ▶ A container packs together an application and all its dependencies and isolates the application from the rest of the machine it runs on.
- ▶ **Running multiple instances:** Because the dependencies are isolated from each other you can run multiple containers on the same machine without them interfering with each other.
- ▶ **Automated installation on clusters:** Orchestrators (such as Kubernetes) automatically distribute containerized applications across a cluster of servers so you do not have to manually install applications.

Linux features that make containers work

- ▶ **Control groups (Cgroups):** limits the resources, such as memory, CPU, and network input/output, that a group of Linux processes can use
 - ▶ By limiting the resources a process can use, containers provide protection against attacks that consume excessive resources
- ▶ **Linux Namespaces:** restricts visibility of resources to a process
 - ▶ By putting a process in a namespace, you can restrict the resources that are visible to that process
- ▶ **Changing the Root Directory:** limits the set of files and directories that a process can see
 - ▶ By changing the root directory when the container is created, a container can not see the host's entire filesystem

Containers vs VMs

- ▶ **Hypervisors:** a pure virtual machine environment, a dedicated kernel-level VMM program runs instead of the OS kernel.
- ▶ **Hosted VM:** VMs are hosted by the host OS, a VMM runs on the host OS and the guest OS runs on the VMM – e.g. VirtualBox on your laptop
- ▶ **Containers** share the kernel with the OS on the machine they are running on
- ▶ VM – fixed resources, overhead of running a whole kernel.
- ▶ Faster to start a container than a kernel
- ▶ VMs offer better isolation



Basic docker commands

What is Docker:

- ▶ Docker is an application that allows you to create and build containers
 - ▶ Set of commands to manipulate images and containers
- ▶ **Image:** complete and executable version of an application
- ▶ **Container:** is the instantiation of an image
- ▶ Docker maintains a repository of images

Docker basic commands: push/pull

- ▶ Allow you to pull an image from repo or push; you need to have a docker account to run push

```
docker push [OPTIONS] NAME[:TAG]
```

```
docker pull [OPTIONS] NAME[:TAG|@DIGEST]
```

```
$ docker push registry-host:5000/myadmin/rhel-httpd
```

```
$ docker pull Debian
```

<https://docs.docker.com/engine/reference/commandline/push/>

<https://docs.docker.com/engine/reference/commandline/pull/>

Docker basic commands: run/stop

- ▶ Run a command in a new container

```
docker run [OPTIONS] IMAGE [COMMAND] [ARG...]
```

```
docker stop [OPTIONS] CONTAINER [CONTAINER...]
```

```
$ docker run mydockerhello
```

```
$ docker stop mydockerhello
```

<https://docs.docker.com/engine/reference/commandline/run/>

<https://docs.docker.com/engine/reference/commandline/stop/>

Docker basic commands: ps

- ▶ Allows to see containers running or finished (finished containers are saved on disk)

```
docker ps [OPTIONS]
```

Examples:

```
$ docker ps
```

```
$ docker ps -a
```

<https://docs.docker.com/engine/reference/commandline/ps/>

Docker basic commands: images

- ▶ List all images
- ▶ `docker images [OPTIONS] [REPOSITORY[:TAG]]`
- ▶ `$docker images java`
- ▶ <https://docs.docker.com/engine/reference/commandline/images/>

Docker basic commands: rm/rmi

- ▶ Allows to remove containers/images, you can not remove an image before you removed all containers that are using it
- ▶ `docker rm [OPTIONS] CONTAINER [CONTAINER...]`
- ▶ `docker rmi [OPTIONS] IMAGE [IMAGE...]`

```
$ docker rm redis
```

```
$ docker rmi test:latest
```

<https://docs.docker.com/engine/reference/commandline/rm/>

<https://docs.docker.com/engine/reference/commandline/rmi/>

Docker basic commands: volume mapping

- ▶ When the container ends, data is not persistent, it is lost.
- ▶ To have data persistent after the container ends you can mount in the container a volume (file, directory) from the host

```
$ docker run -v /opt/data:/var/lib/mysql mysql
```

`/var/lib/mysql` of the container is mapped to `/opt/data` of the host.

<https://docs.docker.com/engine/reference/commandline/run>

Docker basic commands: port mapping

- ▶ Containers are not accessible via networking from outside the host by default
- ▶ Using port mapping you can access them from outside, using host ports
- ▶ You can run multiple instances, but each host port can be mapped only once

```
$ docker run -p 8001:5000 mycontainer1
```

```
$ docker run -p 8002:5000 mycontainer2
```

When accessing via host port 8001 you will access mycontainer1

<https://docs.docker.com/engine/reference/commandline/run>

Docker basic commands: stdin, stdout,stderr

- ▶ Containers are not mapped to stdin, stdout, stderr by default, option `-i` for run command

```
$ docker run -i mycontainer |
```

When running mycontainer, if you're reading from stdin, you will be prompted to input from the keyboard

<https://docs.docker.com/engine/reference/commandline/run>

Docker basic commands: attach/detach

- ▶ run option `-d` runs the container in the background
- ▶ run option `-a` attaches stdout, stdin, stderr for a container

```
$ docker run -d mycontainer
```

```
$ docker run -a mycontainer |
```

```
$ docker run -i -a STDERR mycontainer
```

<https://docs.docker.com/engine/reference/commandline/run>

Docker basic commands: exec

- ▶ Run a command in a running container

```
docker exec [OPTIONS] CONTAINER COMMAND [ARG...]
```

```
$ docker run --name ubuntu_bash --rm -i -t ubuntu bash
```

```
$ docker exec -d ubuntu_bash touch /tmp/execWorks
```

Result is creating a new file /tmp/execWorks inside the running container ubuntu_bash, in the background.

<https://docs.docker.com/engine/reference/commandline/exec>

Docker basic commands: Networking

- ▶ When you install docker it creates three networks: *bridge*, *host* and *none*
- ▶ Bridge is an internal private network, all containers get an IP address on this internal network (172.17)
- ▶ Containers can talk to each other using this IP
- ▶ If you want to access them from outside there are several solutions, one is port mapping
- ▶ You can also create your own private network using docker network create

<https://docs.docker.com/engine/reference/commandline/network/>

Docker basic commands: Dockerfile

- ▶ It tells docker how to build a container

```
FROM ubuntu:latest
```

```
RUN apt-get update
```

```
RUN apt-get install -y gcc
```

```
ADD hello.c /app/
```

```
WORKDIR /app/ RUN gcc hello.c -o hello
```

```
ENTRYPOINT /app/hello
```

<https://github.com/asadsalman/docker-tutorial>

Docker basic commands: passing args

```
FROM ubuntu:latest
```

```
RUN apt-get update
```

```
RUN apt-get install -y gcc
```

```
ADD hello.c /app/
```

```
WORKDIR /app/ RUN gcc hello.c -o hello
```

```
ENTRYPOINT /app/hello classnumber
```

```
$docker run hello 7610
```

Project 0: Objective

- ▶ Get started with docker
- ▶ Get started with UDP sockets
- ▶ Bootstrap a set of distributed processes communicating just through messages

- ▶ It's a pass/fail to prepare you for the first project
- ▶ The main goal is to do it completely not to get a grade, it is the basis for project 1

Bootstrapping

- ▶ One problem in distributed systems: how do processes *find each other*, how do you bootstrap
- ▶ Who are the processes
 - ▶ We are going to use a file with list of processes
 - ▶ List needs to be the same
- ▶ Are they up?
 - ▶ Each process needs to know that the other processes came up
 - ▶ Simple approach, *I am alive message*
 - ▶ It is sent with UDP thus can be lost, you have to send a few times