

01: Introduction

Distributed Software Architectures

Stefan Huber <shuber.lba@fh-salzburg.ac.at>

March 8, 2019

Section 1

Course formalities

This course is about [Software Architectures](#) of [Distributed Systems](#).

Topics:

- ▶ Ways to organize distributed software systems
- ▶ Foundations of distributed systems
 - ▶ Communication
 - ▶ Coordination
 - ▶ Virtualization & operating systems
 - ▶ Name systems
 - ▶ Replication and consistency
 - ▶ Security for distributed systems

Related fields and courses:

- ▶ Software engineering
- ▶ Computer networks
- ▶ Operating systems

The course is split into lecture (VL) and tutorial (UE) units.

Grading:

- ▶ 60% lecture: Two 45 min. tests
- ▶ 40% tutorial: Exercises, presentations, coding assignments

FR 08.03.	08:15–09:45	VL	
FR 22.03.	08:15–09:45	VL	10:00–11:30 UE
TH 04.04.	08:15–09:45	VL	10:00–11:30 UE
FR 26.04.	08:15–09:45	VL	10:00–11:30 UE
TH 09.05.	08:15–09:45	VL	10:00–11:30 UE Test
FR 07.06.	08:15–09:45	VL	10:00–11:30 UE
FR 14.06.	08:15–09:45	VL	
MO 24.06.	08:15–09:45	UE	Test

Literature:

- ▶ van Steen, Tanenbaum (2017): Distributed Systems
Free digital copy: <https://www.distributed-systems.net>

Short bio:

- ▶ B&R Industrial Automation:
 - ▶ Headed an R&D team of 5 PhDs in mathematics, computer science and control theory.
 - ▶ Invented the software system operating *ACOPOStrak*.
 - ▶ Recently switched to a new project *Digital Automation*.
- ▶ PostDoc researcher at IST Austria in the field of *computational topology*.
- ▶ Senior Scientist (math), PostDoc (CS), and PhD student (CS) at the Univ. of Salzburg in the field of *computational geometry*.
- ▶ Doctorate in CS, Master in Math, Master in CS, Bachelor in Math and Bachelor in CS in 2011, 2009, 2008, 2007 and 2006.

Section 2

Introduction to Distributed Systems

What is a Distributed System?

A computing **system** that is **distributed** on a computer network.

- ▶ Collection of *autonomous* components (nodes).
- ▶ Nodes cooperate to form a single *coherent* system.
- ▶ Nodes are distributed on and communicate via a computer network.

Characterization (van Steen)

A *distributed system* is a collection of autonomous computing elements that appears to its users as a single coherent system.

WWW Master-slave. The World Wide Web is the probably the largest distributed system. Web browsers and web servers as nodes that communicate via HTTP protocol over the Internet. For the user the WWW appears as a single system.

Bittorrent A peer-to-peer file sharing protocol using so-called distributed hash tables to find peers.

WWW Master-slave. The World Wide Web is the probably the largest distributed system. Web browsers and web servers as nodes that communicate via HTTP protocol over the Internet. For the user the WWW appears as a single system.

Bittorrent A peer-to-peer file sharing protocol using so-called distributed hash tables to find peers.

GIMPS Great Internet Mersenne Prime Search is an example for *distributed computing*.

NFS The Network File System is a distributed file system based on Remote Procedure Calls.

WWW Master-slave. The World Wide Web is the probably the largest distributed system. Web browsers and web servers as nodes that communicate via HTTP protocol over the Internet. For the user the WWW appears as a single system.

Bittorrent A peer-to-peer file sharing protocol using so-called distributed hash tables to find peers.

GIMPS Great Internet Mersenne Prime Search is an example for *distributed computing*.

NFS The Network File System is a distributed file system based on Remote Procedure Calls.

Car A car comprises a distributed system of several dozens ECUs, controlling the engine, brakes, doors, or HMI. A typical automotive network is CAN bus.

Automation An industrial machine comprises controllers, drives, sensors, HMI with realtime communication over a fieldbus.

- ▶ Nodes are in principle independent from each other, but cooperate.
- ▶ Nodes run concurrently.
- ▶ The collection may be very heterogeneous.
- ▶ Each node has its own notion of time. There may not be a global clock.
- ▶ Cooperation requires communication, e.g., message passing.
- ▶ Notion of an overlay network, e.g., the communication topology:
 - ▶ Structured overlay: Tree, ring, grid, et cetera.
 - ▶ Unstructured overlay: Randomly chosen neighborhood.
- ▶ Group management for the collection of nodes:
 - ▶ Open group: Any node can join
 - ▶ Closed group: Access is restricted

Characteristics: Single coherent system

- ▶ Wiktionary on *coherent*: Unified; sticking together; making up a whole.
- ▶ Level of coherence versus the level of distribution.
- ▶ Distribution transparency:
 - ▶ User does not need to know where something is processed, stored, et cetera.
 - ▶ Half of system engineering is abstraction! Abstraction means hiding details.
- ▶ Fault tolerance

Characterization (Lamport)

[A distributed system is] one in which the failure of a computer you didn't even know to exist can render your computer unusable.

Section 3

Python

Python is popular:

- ▶ 3rd on TIOBE index 01/2019: Java, C, [Python](#), C++, VB.net, JS, C#

Python is very simple and easy to learn:

- ▶ Very popular at MIT courses
- ▶ Syntax is simple and expressive, reasonably close to pseudo code
- ▶ *Distributed Systems* uses Python for examples

Python is technologically rich:

- ▶ Originated in 1990
- ▶ Interpreted, platform independent
- ▶ Duck typing, very dynamic
- ▶ General-purpose, multi-paradigm: object-oriented, functional, aspect-oriented, ...
- ▶ A rich standard library and a huge universe of third-party modules
- ▶ IPython and friends make it a Matlab alternative

This is a Python 3 hello world example:

```
1 #!/usr/bin/env python3
2
3 print("Hello World!")
```

Listing 1: helloworld.py

Learning Python:

- ▶ Python tutorial:
<https://docs.python.org/3/tutorial/>
- ▶ Python in 10 minutes:
<https://www.stavros.io/tutorials/python/>
- ▶ Python 101:
http://www.davekuhlman.org/python_101.html
For Python 2, but mentions the differences to Python 3.

Brief tutorials, but for Python 2:

- ▶ Google for Education:
<https://developers.google.com/edu/python/introduction>
- ▶ A very short introduction by scipy:
http://scipy-lectures.org/language/python_language.html

Python can be used in an interactive interpreter:

```
1 >>> a = 2
2 >>> b = True
3 >>> print(a, b)
4 2 True
5 >>> type(a), type(b), type("hello"), type('world')
6 (<class 'int'>, <class 'bool'>, <class 'str'>, <class 'str'>)
7 >>> "hello" + 'world'
8 'helloworld'
9 >>> 2 + 2, 3 * 4, 4**2, 17 / 3, 17 // 3, 3 * "no", True and False or True
10 (4, 7, 16, 5.666666666666667, 5, 'nonono', True)
11 >>> r"c:\who\uses\windows\anyways"
12 'c:\\who\\uses\\windows\\anyways'
13 >>> """A multiline
14 string can be
15 handy"""
16 'A multiline\nstring can be\nhandy'
17 >>> c = "slicing"
18 >>> c[0], c[-1], c[1:4], c[-3:], c[:], c[::2], c[::-1], len(c)
19 ('s', 'g', 'lic', 'ing', 'slicing', 'siig', 'gnicils'), 7
20 >>> c[:2] + c[2:]
21 'slicing'
```

Listing 2: Variables and basic types

A list is a heterogeneous sequence of elements.

```
1 >>> squares = [1, 4, 9, 16, 25]
2 >>> type(squares)
3 <class 'list'>
4 >>> squares[::-1]           # Indexing and slicing works for all sequence types
5 [25, 16, 9, 4, 1]
6 >>> squares + [36, 49, 64, 81, 100]
7 [1, 4, 9, 16, 25, 36, 49, 64, 81, 100]
8 >>> bag = ["hello", 42, [0, 1]]           # Can be heterogeneous
9 >>> bag[0] = "hi!"
10 >>> bag.append(3.14)
11 >>> bag
12 ['hi!', 42, [0, 1], 3.14]
```

Listing 3: Lists

A tuple is, like a list, a sequence type, but immutable.

```
1 >>> tri = 23, 42, 100           # The same as (23, 42, 100)
2 >>> type(tri)
3 <class 'tuple'>
4 >>> tri[0], len(tri)           # Note that this expression is a tuple of two
5 (23, 3)
6 >>> a, b, c = tri              # Assign tri to the tuple (a, b, c)
7 >>> print(a, b, c)
8 23 42 100
```

Listing 4: Tuples

```
1 #!/usr/bin/env python3
2
3 a, b = 0, 1
4 while a < 10:
5     print(a)
6     a, b = b, a + b
```

Listing 5: The Fibonacci series

Some remarks:

- ▶ We use tuples for multiple assignments and variable definitions
- ▶ Indentation defines blocks. It is 4 spaces, see <https://www.python.org/dev/peps/pep-0008/>.
- ▶ The function `help()` in the interpreter gives help on anything, e.g., `print()`.

```
1 #!/usr/bin/env python3
2
3 x = int(input("Enter a whole number: "))
4
5 if x < 0:
6     x = 0
7     print("Negative changed to zero")
8 elif x == 0:
9     print("Zero")
10 else:
11     print("Positive")
12
13
14 for w in "hello!":
15     print(w)
16     if w == 'o':
17         print("exit loop")
18         break
19
20 for i in range(5):
21     print(i)
22
23 while True:
24     pass
```

Iterating over a sequence

range(5) generates the numbers 0, 1, 2, 3, 4

Endless loop. Exit by pressing ctrl+c

Listing 6: More control statements

```
1 #!/usr/bin/env python3
2
3 def fib_print(n):
4     """Print the Fibonacci up to at most n."""           # The __doc__ string
5     a, b = 0, 1
6     while a < n:
7         print(a, end=', ')
8         a, b = b, a + b
9     print()
10
11 def fib_print_ext(n, start=0):
12     """Print the Fibonacci series up to at most n, but at least start."""
13     a, b = 0, 1
14     while a < n:
15         if a >= start:
16             print(a, end=', ')
17             a, b = b, a + b
18     print()
19
20 print("Our fancy module", __name__, "loaded")
21
22 if __name__ == "__main__":                               # If called from interpreter
23     fib_print(30)
24     fib_print_ext(30, 4)
```

Listing 7: Functions

A set has unique elements and adding, removing, ownership-testing is fast.

```
1 >>> bucket = {'hi', 0, 1, 2, 2, 2}
2 >>> len(bucket)
3 4
4 >>> 1 in bucket
5 True
```

Listing 8: set

A dictionary is a key-value map.

```
1 >>> codes = {5020: "Salzburg", 6010: "Innsbruck", 4010: "Linz"}
2 >>> len(codes)
3 3
4 >>> codes[5020]
5 'Salzburg'
```

Listing 9: dict

Python: Exceptions

```
1 #!/usr/bin/env python3
2
3 if __name__ == "__main__":
4     try:
5         x = int(input("Enter whole number: "))
6         print("Entered", x)
7     except ValueError:
8         print("No whole number given")
9
10    def whiney(): # A function can be defined in any scope
11        raise Exception("A whiney function")
12
13    try:
14        whiney()
15    except Exception as err:
16        print(err)
17
18    try:
19        f = open("doesnotexist.txt", "r")
20    except OSError as err:
21        print("Error opening file:", err)
22    except:
23        print("Some unknown error")
24    else:
25        print("Close again")
26        f.close()
```

Listing 10: Exceptions

```
1 #!/usr/bin/env python3
2
3 class Vector:
4     def __init__(self, x, y):                # Constructor
5         self.x, self.y = x, y
6
7     def printit(self):                       # All methods are functions with self as 1st arg
8         print("{} {}".format(self.x, self.y)) # String formatting
9
10    def add(self, v):                          # We do not care about type(v)...
11        self.x += v.x                        # ... only that it has members x and y
12        self.y += v.y                        # It is called Duck Typing
13
14 class Object:
15     pass
16
17 if __name__ == "__main__":
18     v = Vector(1.0, 2.0)
19     v.add(Vector(5.0, 6.0))
20     v.printit()
21
22     u = Object()                             # u has no members
23     u.x, u.y = 10.0, 20.0                    # Now we dynamically add members x and y
24     v.add(Vector(5.0, 6.0))
25     v.printit()
```

Listing 11: Classes

```
1 #!/usr/bin/env python3
2
3 class A:
4     def __init__(self, x, y):
5         self.x, self.y = x, y
6
7     def f(self):
8         print("(%g, %g)" % (self.x, self.y))           # String formatting
9
10 class B(A):                                           # Inheritance
11     def __init__(self, x, y, z):
12         A.__init__(self, x, y)                       # Calling base class constructor
13         self.z = z
14
15     def f(self):                                     # Polymorphism
16         print("(%g, %g, %g)" % (self.x, self.y, self.z))
17
18 if __name__ == "__main__":
19     b = B(1, 2, 3)
20     b.f()
```

Listing 12: Inheritance

Python: Beyond Hello World

```
1 #!/usr/bin/env python3
2
3 # A function
4 def print_square(n):
5     print("{} squared is {}".format(n, n**2))
6     if n % 2 == 0:
7         print("    even", n)
8
9 # Will also be called when this module is imported somewhere
10 print("module loaded")
11
12 # The module name is "__main__" if called by the interpreter rather than
13 # being imported
14 if __name__ == "__main__":
15     # Loop over the list [0, 1, 2, 3, 4]
16     for i in range(5):
17         print_square(i)
18
19     li = []
20     i = 0
21     while i < 5:
22         li.append(i*i)
23         i += 1
24
25     # List comprehension: Functional programming rocks!
26     assert(li == [x*x for x in range(5)])
```

Listing 13: morehelloworld.py

Python: Further beyond Hello World

```
1 #!/usr/bin/env python3
2
3 # A class
4 class PowerPrinter:
5     # A constructor
6     def __init__(self, exp=2):
7         self.exp = exp
8
9     # A method
10    def print(self, n):
11        fmt = "{} to the power of {} is {}"
12        print(fmt.format(n, self.exp, n**self.exp))
13
14 if __name__ == "__main__":
15     # Loop over a tuple with two elements
16     for p in PowerPrinter(), PowerPrinter(3):
17         # Loop over the list of whole numbers [0..5)
18         for i in range(5):
19             p.print(i)
```

Listing 14: moremorehelloworld.py

Section 4

Exercise formalities

Hard facts:

- ▶ Exercise sheets on the course wiki page.
- ▶ Deadline given on the sheet.
- ▶ Exercises submitted via git repo on `gitlab.mediacube.at`.
- ▶ Please respect the dictionary name formatting.

Soft facts:

- ▶ Use static code analyzers, like `flake8` or `pylint`, to improve your results and avoid trivial mistakes.
- ▶ A sample git repository is on `gitlab.mediacube.at` to ease the getting started phase.

- ▶ Work load primarily on getting familiar with Python at home.
- ▶ Therefore, only two small warm-up exercises.