

Linux Kernel Programming

Introduction to the Linux Kernel

Pierre Olivier

Systems Software Research Group @ Virginia Tech

January 17, 2017

Im doing a (free) operating system (just a hobby, wont be big and professional like gnu) for 386(486) AT clones.

Linus Torvalds
08-25-1991

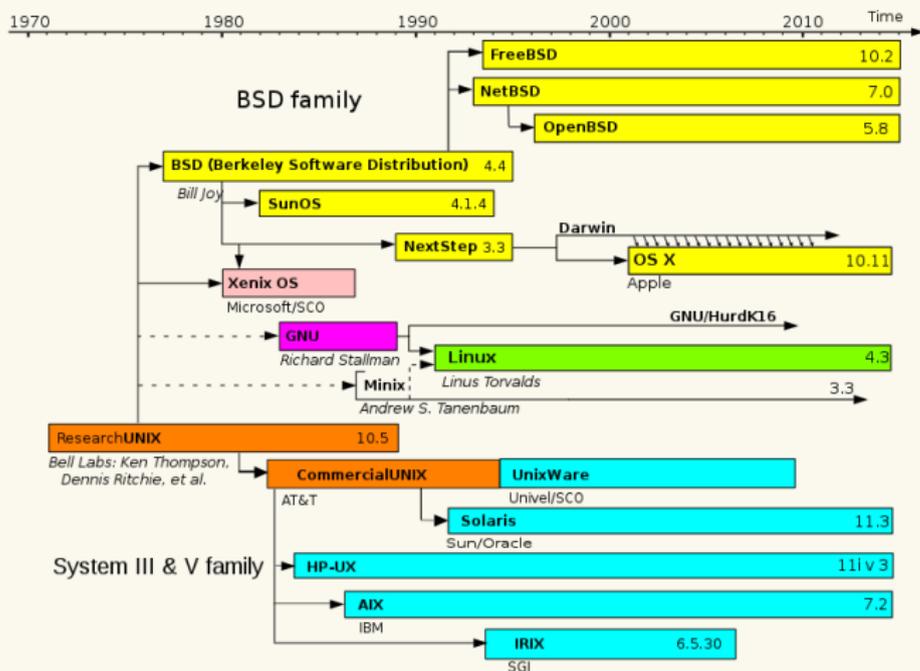
Outline

- 1 A bit of history
- 2 Linux usage today
- 3 Linux open source model & the community
- 4 Global overview of the kernel

Outline

- 1 A bit of history
- 2 Linux usage today
- 3 Linux open source model & the community
- 4 Global overview of the kernel

The Unix OS family



Source: [9]

The popularity of Unix

- ▶ Unix is very popular and ported to multiple architectures
 - ▶ This is due to its **simple design** and ease of use
- ▶ **Simplicity:**
 - ▶ Small number of system calls with clearly defined design
 - ▶ *Everything is a file*
 - ▶ Written in C for **portability**
 - ▶ Easy and fast process creation (`fork()`)
 - ▶ Simple and efficient Inter-Process Communication mechanisms (IPC)

Enters Linux

```
1 From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)
2 Newsgroups: comp.os.minix
3 Subject: What would you like to see most in minix?
4 Summary: small poll for my new operating system
5 Message-ID: <1991Aug25.205708.9541@klaava.Helsinki.FI>
6 Date: 25 Aug 91 20:57:08 GMT
7 Organization: University of Helsinki
8
9 Hello everybody out there using minix
10
11 I'm doing a (free) operating system (just a hobby, won't be big and
12 professional like gnu) for 386(486) AT clones. This has been brewing
13 since april, and is starting to get ready. I'd like any feedback on
14 things people like/dislike in minix, as my OS resembles it somewhat
15 (same physical layout of the file-system (due to practical reasons)
16 among other things).
17
18 I've currently ported bash(1.08) and gcc(1.40), and things seem to work.
19 This implies that I'll get something practical within a few months, and
20 Id like to know what features most people would want. Any suggestions
21 are welcome, but I won't promise I'll implement them
22
23 Linus (torvalds@kruuna.helsinki.fi)
24
25 PS. Yes      it's free of any minix code, and it has a multi-threaded fs.
26 It is NOT protable (uses 386 task switching etc), and it probably never
27 will support anything other than AT-harddisks, as that's all I have :-).
```

Enters Linux (2)

- ▶ 1991: First apparition, author: *Linus Torvalds*
- ▶ 1992: GPL License, first Linux distributions
- ▶ 1994: Linux v1.0 - Single CPU for i386, then quickly ported to Alpha, Sparc, Mips architectures
- ▶ 1996: v2.0 - *Symmetric multiprocessing* (SMP) support
- ▶ 1999: v2.2 - *Big Kernel Lock* removed
- ▶ 2001: v2.4 - USB, later: RAID, Bluetooth, etc.
- ▶ 2003: v2.6 - Lots of features: Physical Address Expansion (PAE), new architectures, new file systems, etc.
- ▶ 2011: v3.0 - "Nothing" [1]
- ▶ 2015: v4.0 - Livepatch, Kernel Address Space Sanitizer (KASAN)
- ▶ Today stable version: *Let's check it out* → www.kernel.org

Sources: [10, 5]

Outline

- 1 A bit of history
- 2 Linux usage today**
- 3 Linux open source model & the community
- 4 Global overview of the kernel

A Few Numbers

Linux usage in today's computer systems landscape

Embedded systems

- ▶ Smartphones, tablets, etc.: **65%** market share (Android)
 - ▶ vs iOS (24%)
- ▶ Embedded systems in general: **29%** (Android, embedded Linux solutions)



Servers

- ▶ **36%** market share (top 10M web servers)
 - ▶ vs other Unix-like OS (30%) and Windows (30%)
- ▶ **98%** (top 1M)

Source: [11]

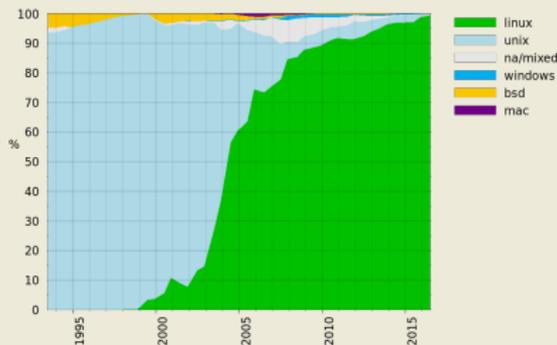
A Few Numbers (2)

Linux usage in today's computer systems landscape

Source: [11]

High-performance computing

▶ **99% market share**



Desktop Computers

▶ **2% market share**

▶ vs Microsoft Windows (90%) and Mac OSX (8%)



▶ **Indisputably one of the most popular operating systems over the whole spectrum of today's computer systems**

Jobs & Linux

- ▶ *Linux Jobs Report* released by the Linux foundation in 2014 [3] and 2015 [4]
 - ▶ 1000+ hiring managers polled as well as 3400+ Linux professionals
 - ▶ From large corporations to medium/small businesses, and government organizations
- ▶ **97% of the managers report they will hire Linux talent in the next six month**
 - ▶ Up from 46% in 2014
- ▶ **86% of Linux professionals report knowing Linux has given them more career opportunities**
 - ▶ **64% says they chose to work with Linux because it is pervasive**

Outline

- 1 A bit of history
- 2 Linux usage today
- 3 Linux open source model & the community**
- 4 Global overview of the kernel

Linux open source model

- ▶ Linux is licensed under **GPLv2**
 - ▶ Sources are freely available (<https://www.kernel.org/>)

Quick summary of the license terms:

You may copy, distribute and modify the software as long as you track changes/dates in source files. Any modifications to a software including (via compiler) GPL-licensed code must also be made available under the GPL along with build & install instructions.

<https://tldrlegal.com/license/gnu-general-public-license-v2>



- ▶ Third party can distribute modules as proprietary binaries
 - ▶ Compliance with GPLv2 is still an outstanding question, but Linus allows this

Linux open source model (2)

Benefits from the open source model

Benefits of open source vs closed development:

- ▶ **Security, stability:** *Given enough eyeballs, all bugs are shallow* [12]
 - ▶ Not a panacea though (OpenSSL Heartbleed, Bash Shellshock, etc.)
- ▶ **Code quality**
- ▶ **Freely modifiable** by anyone having an interest to do so
 - ▶ Hardware and systems vendors, distributors, end users
 - ▶ Enable innovation !
- ▶ **Education, research**
 - ▶ That's us !
 - ▶ We can `PEEK` at the code to understand how it works
 - ▶ And we can `POKE` it to experiment and build new systems

The Linux community

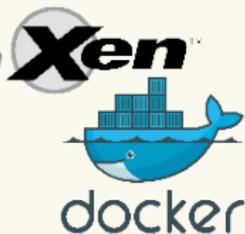
▶ **Developers:**

- ▶ Anybody can propose modifications (*patches*)
 - ▶ Patches posted on **mailing lists**
 - ▶ Reviewed and commented by the subscribers
 - ▶ Then by **kernel maintainers**
 - ▶ More info here: [7]
- ▶ Most of Linux contributors are actually employed by companies producing hardware (ex: Intel), systems (ex: Google), or doing consulting (ex: Red Hat)

▶ **Linux foundation** [2]

▶ **Conventions & Symposiums**

▶ **Software ecosystem** - very few examples:



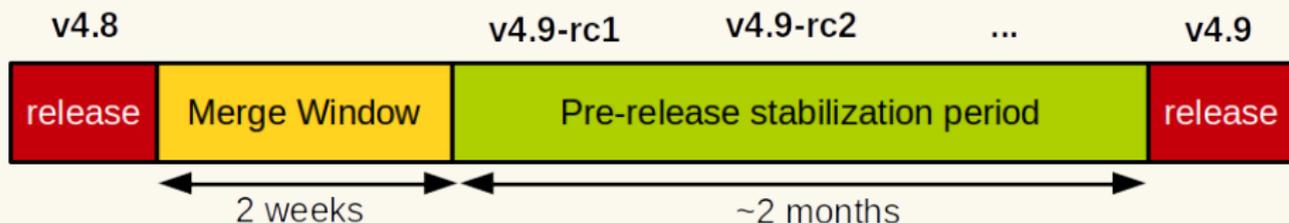
etc.



Linux open source model & the community

Kernel version numbering system

- ▶ **<major>.<minor>.<stable>**
 - ▶ ex: 4.8.17



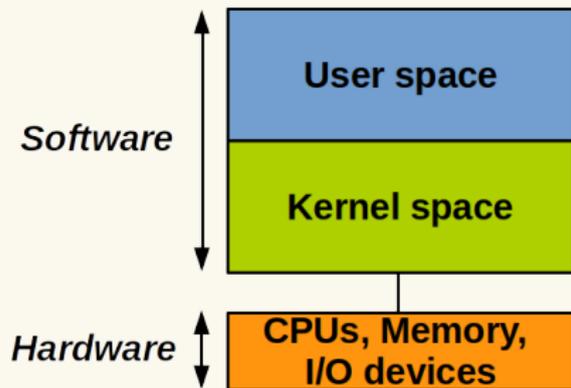
- ▶ **Stable version numbers** (bugs fixes) are added after release
- ▶ **Long Term Support** for a subset of releases:
 - ▶ Ex: 4.4.43 or 3.2.84 vs 4.8.17 [EOL]

Outline

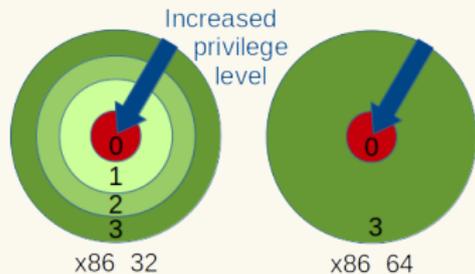
- 1 A bit of history
- 2 Linux usage today
- 3 Linux open source model & the community
- 4 Global overview of the kernel**

Global overview of the kernel

User vs Kernel space

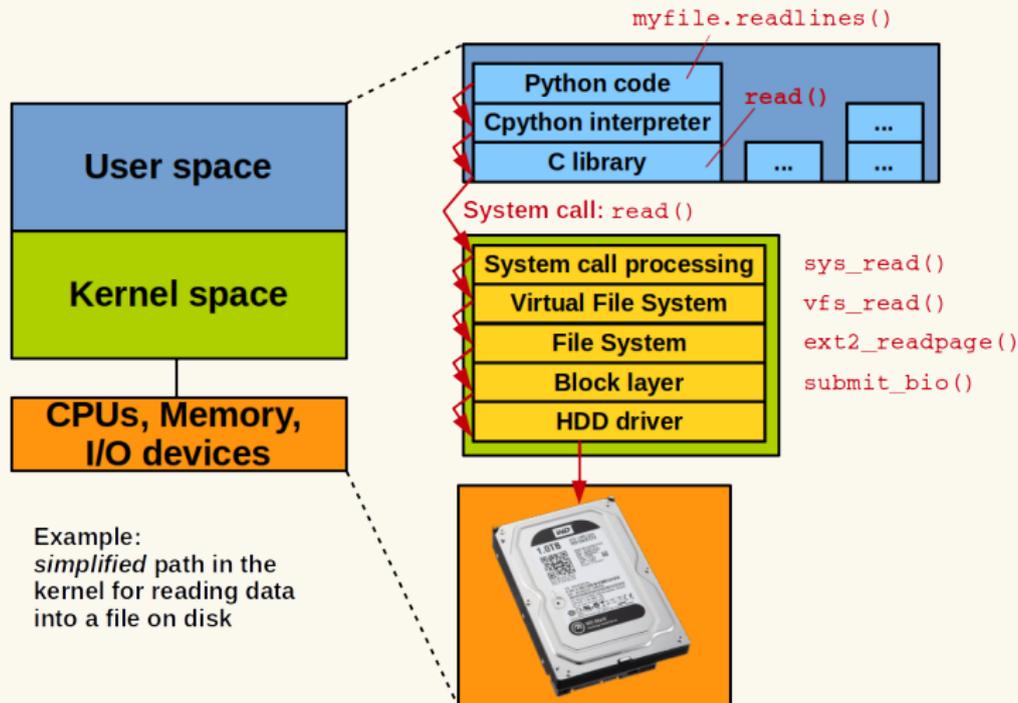


- ▶ At any given moment a CPU is executing in **user** or in **kernel** mode
- ▶ Only the kernel is allowed to perform privileged operations and access I/O devices



Global overview of the kernel

User vs Kernel space (2)



Example:
simplified path in the
kernel for reading data
into a file on disk

Global overview of the kernel

Monolithic Model

- ▶ Linux is **monolithic**
 - ▶ **All the OS services run in kernel, privileged mode**
 - ▶ In the **same address space**
- ▶ Opposed to the *microkernel* model:
 - ▶ Only the core services run in full privileged mode
 - ▶ Scheduling, interrupt handling, etc.
 - ▶ Other services run with reduced privileged in their own address space
 - ▶ Communication through message passing
 - ▶ Ex: Minix, L4, etc.
 - ▶ More secure and modular, but communication brings performance down
 - ▶ See the *Tanenbaum-Torvalds Debate* [8]
 - ▶ Other classes of OS: Exokernel, MultiKernel, etc.

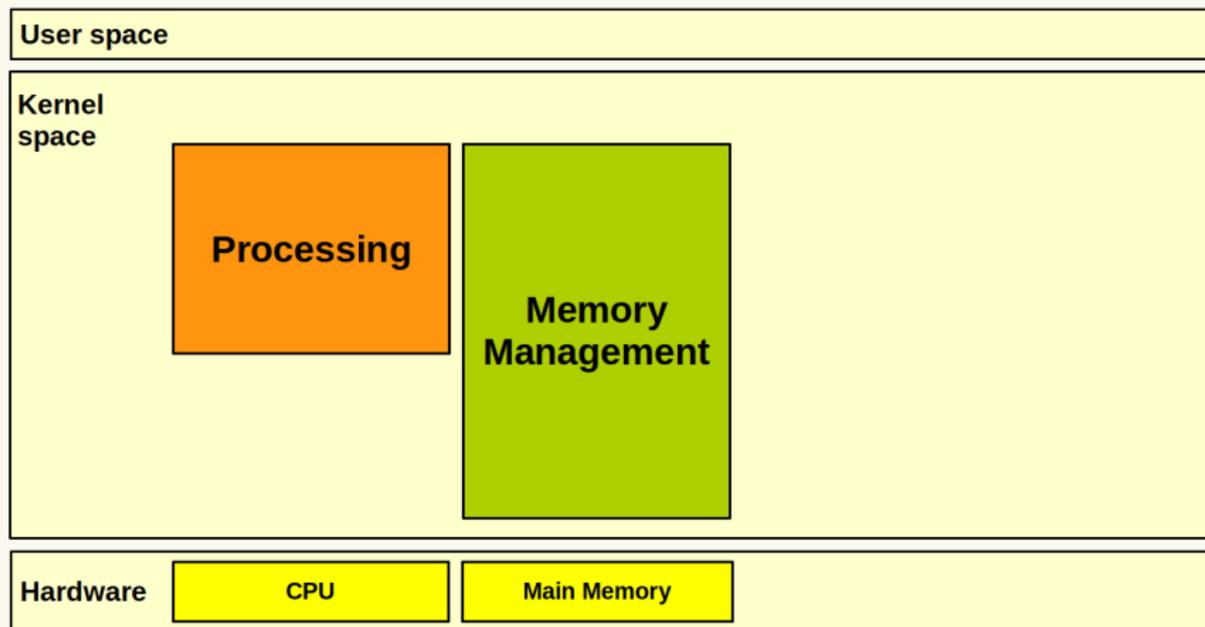
Global overview of the kernel

Kernel vs user level programming

- ▶ **No C library** mostly because of speed and size
 - ▶ However, the kernel implements lots of usefull functions from the C library
 - ▶ Ex: `strcat`, `vmalloc`, `ssleep`, etc.
 - ▶ `printf` → `printk`
- ▶ **2% assembly, the rest in GNU C**
 - ▶ GCC extensions
- ▶ **No memory protection**
- ▶ **Highly concurrent**
 - ▶ Preemption, interrupts, running on SMP
 - ▶ Race conditions without proper synchronization

Global overview of the kernel

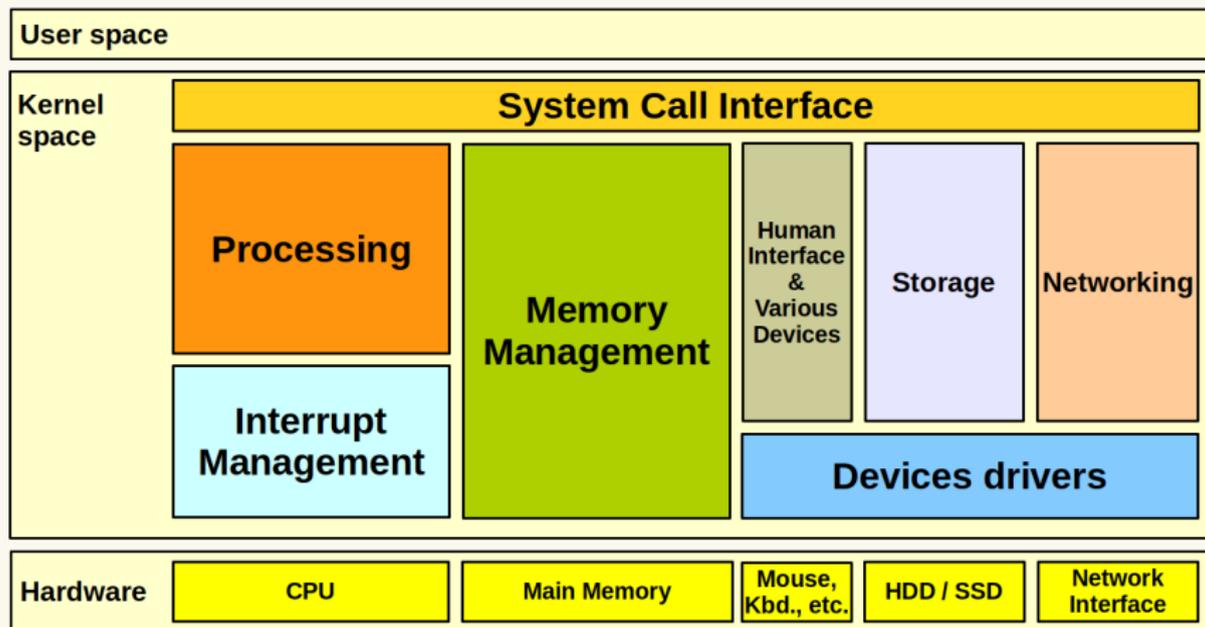
Kernel & course map



Inspired from [6].

Global overview of the kernel

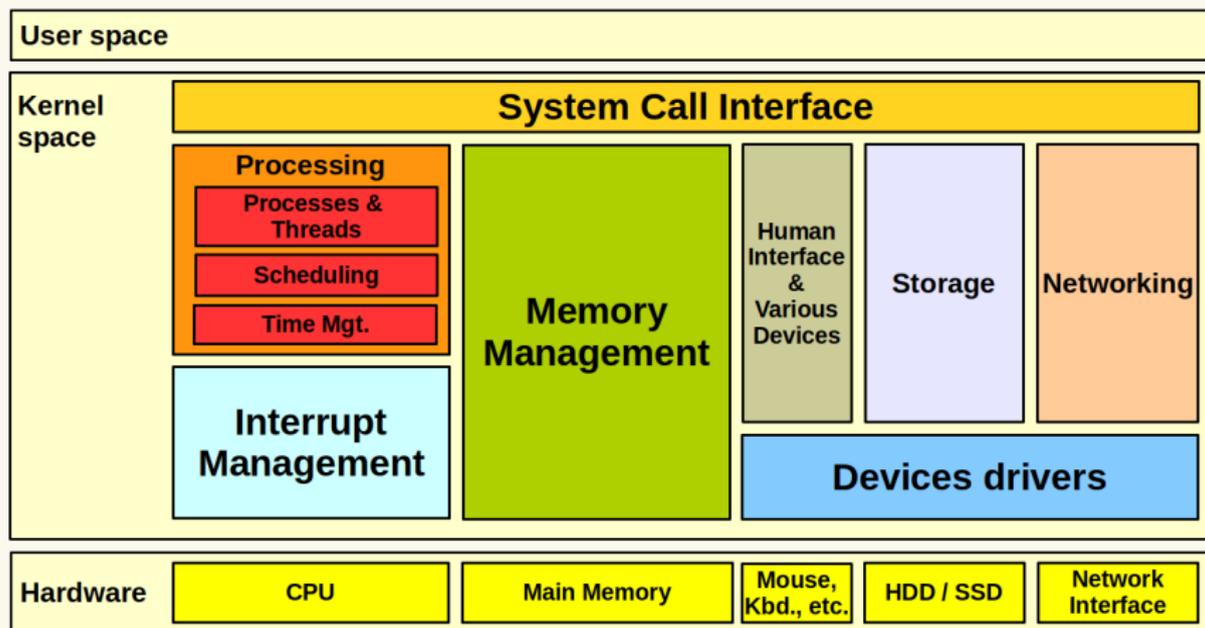
Kernel & course map



Inspired from [6].

Global overview of the kernel

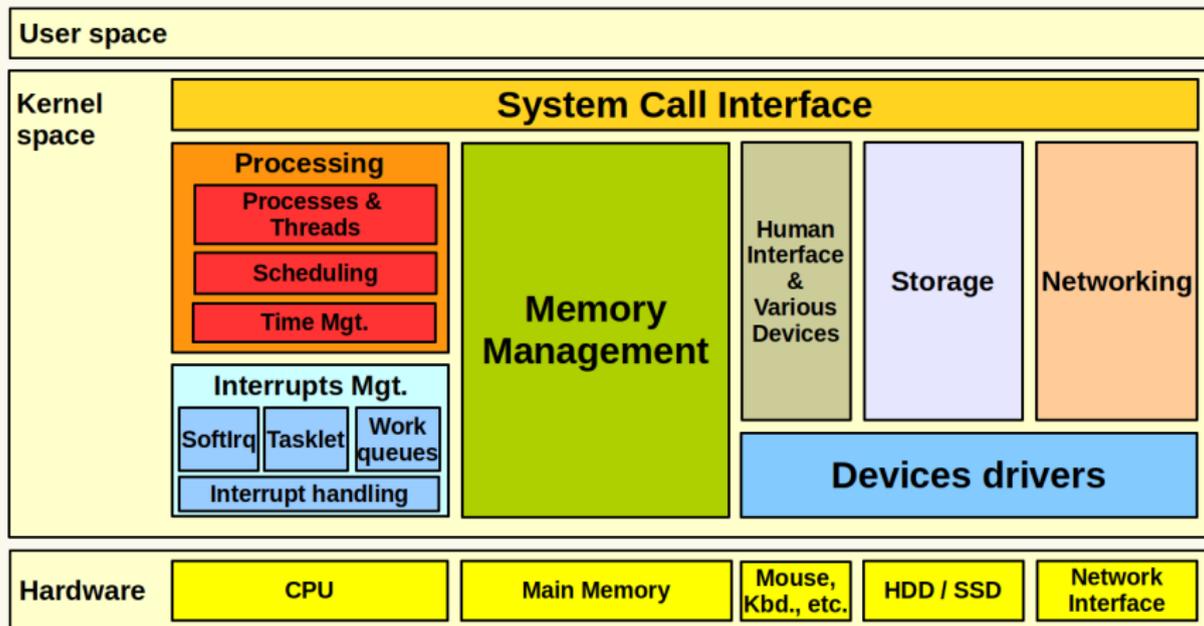
Kernel & course map



Inspired from [6].

Global overview of the kernel

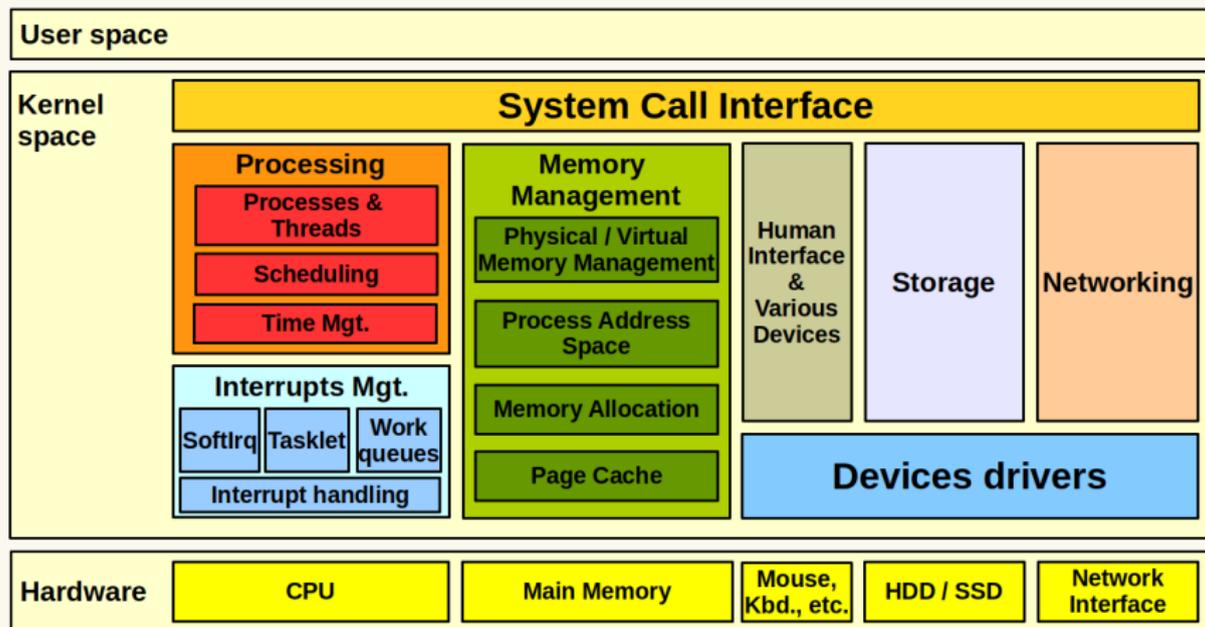
Kernel & course map



Inspired from [6].

Global overview of the kernel

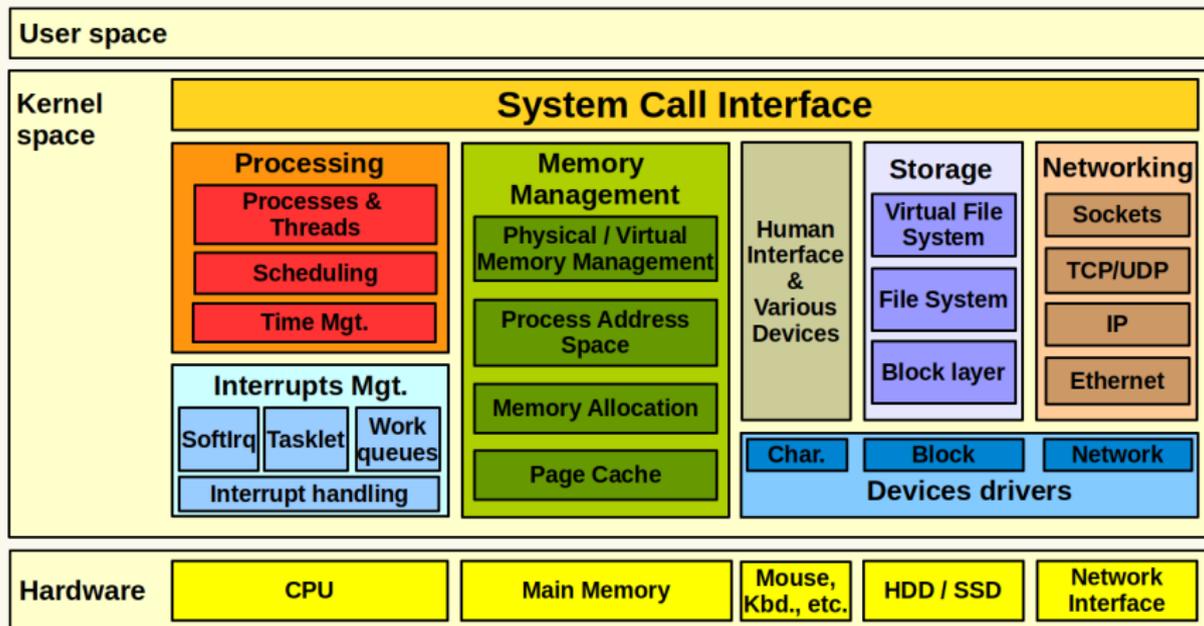
Kernel & course map



Inspired from [6].

Global overview of the kernel

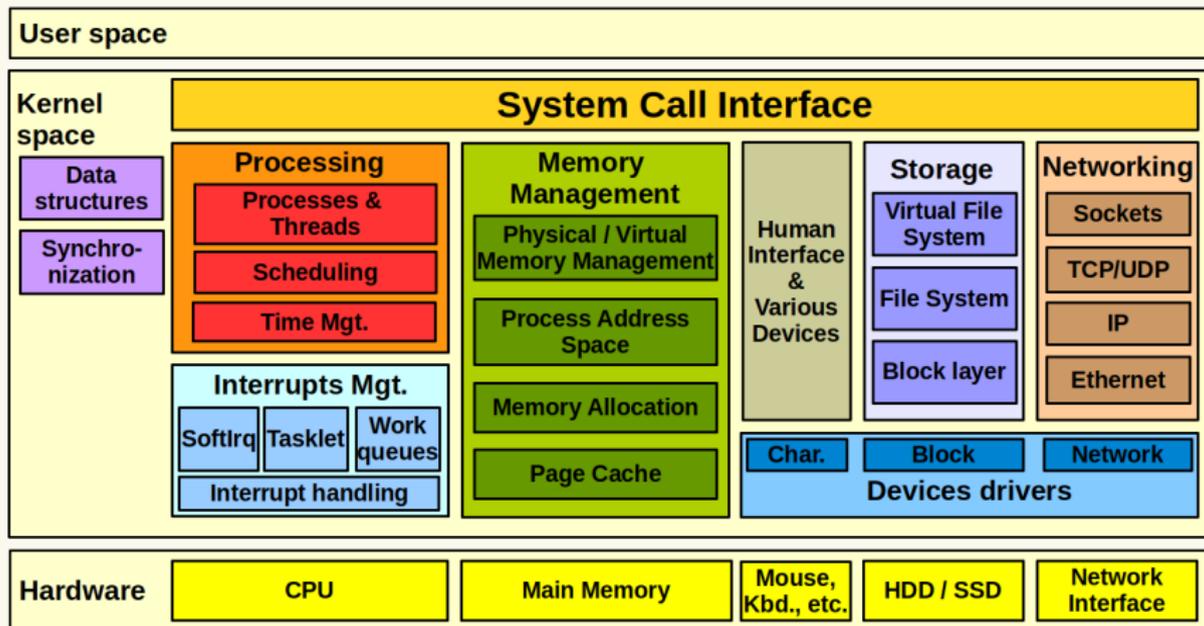
Kernel & course map



Inspired from [6].

Global overview of the kernel

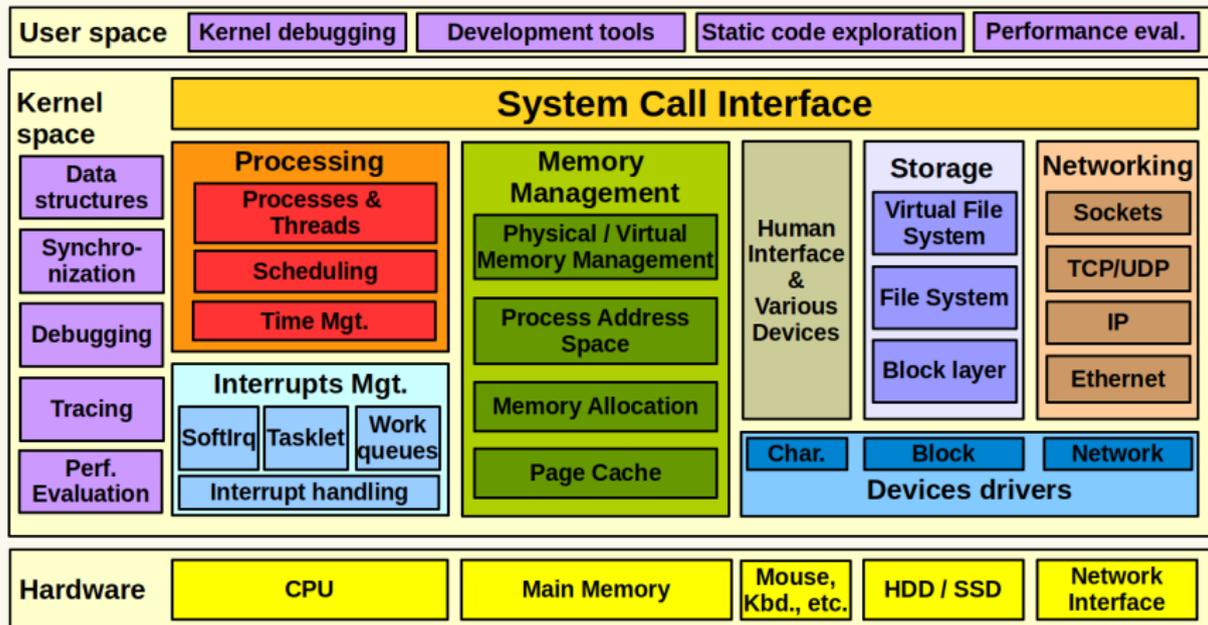
Kernel & course map



Inspired from [6].

Global overview of the kernel

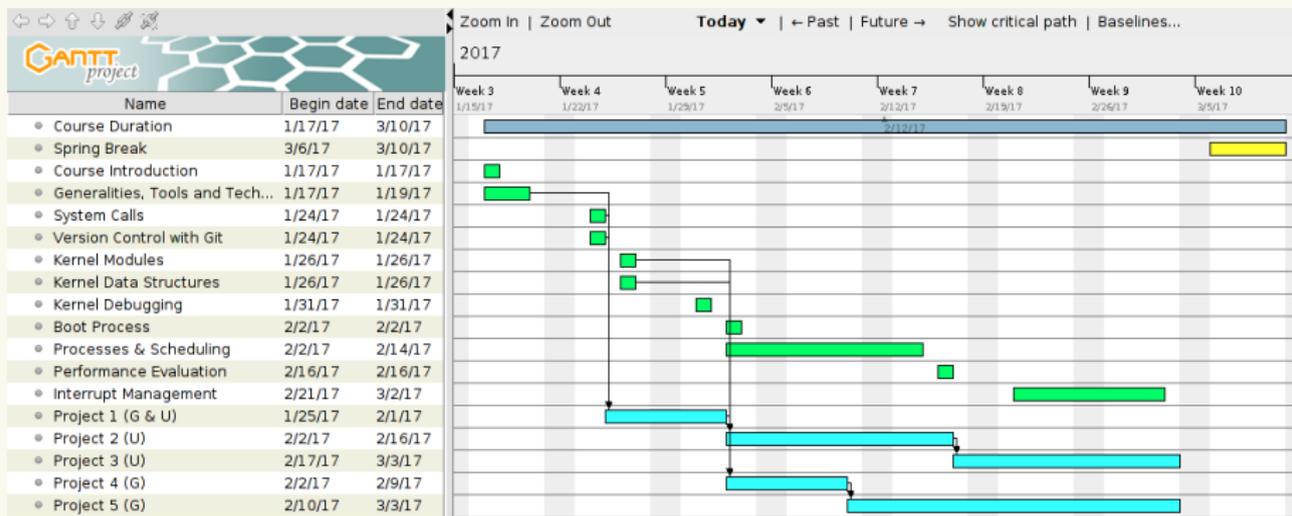
Kernel & course map



Inspired from [6].

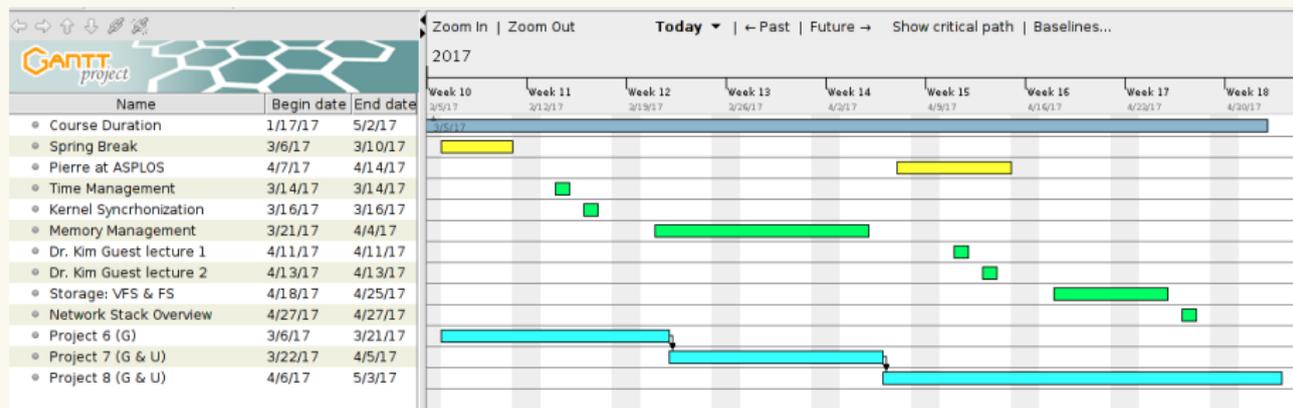
Course schedule (tentative)

- ▶ From the start of the course to the spring break:



Course schedule (tentative) (2)

► After the spring break:



Bibliography I

- [1] [Linux 3.0-rc1.](https://lkml.org/lkml/2011/5/29/204)
[https://lkml.org/lkml/2011/5/29/204.](https://lkml.org/lkml/2011/5/29/204)
Accessed: 2016-09-30.
- [2] [Linux foundation website.](https://www.linuxfoundation.org/)
[https://www.linuxfoundation.org/.](https://www.linuxfoundation.org/)
Accessed: 2016-09-30.
- [3] [Linux jobs report 2014.](https://www.linux.com/publications/linux-jobs-report-2014)
[https://www.linux.com/publications/linux-jobs-report-2014.](https://www.linux.com/publications/linux-jobs-report-2014)
Accessed: 2017-01-12.
- [4] [Linux jobs report 2015.](https://www.linux.com/publications/linux-jobs-report-2015)
[https://www.linux.com/publications/linux-jobs-report-2015.](https://www.linux.com/publications/linux-jobs-report-2015)
Accessed: 2017-01-12.
- [5] [Linux kernel newbies - kernel versions.](https://kernelnewbies.org/LinuxVersions)
[https://kernelnewbies.org/LinuxVersions.](https://kernelnewbies.org/LinuxVersions)
Accessed: 2016-09-30.
- [6] [makelinux.net - linux kernel map.](http://www.makelinux.net/kernel_map/)
[http://www.makelinux.net/kernel_map/.](http://www.makelinux.net/kernel_map/)
Accessed: 2016-12-26.
- [7] [A survivor's guide to contributing to the linux kernel.](http://events.linuxfoundation.org/sites/events/files/slides/klf2015_slides_javier_martinez_0.pdf)
[http://events.linuxfoundation.org/sites/events/files/slides/klf2015_slides_javier_martinez_0.pdf.](http://events.linuxfoundation.org/sites/events/files/slides/klf2015_slides_javier_martinez_0.pdf)
Accessed: 2017-01-16.
- [8] [The tanenbaum-torvalds debate.](http://www.oreilly.com/openbook/opensources/book/appa.html)
[http://www.oreilly.com/openbook/opensources/book/appa.html.](http://www.oreilly.com/openbook/opensources/book/appa.html)
Accessed: 2017-01-12.

Bibliography II

- [9] [Wikipedia - darwin os.](#)
[https://en.wikipedia.org/wiki/Darwin_\(operating_system\).](https://en.wikipedia.org/wiki/Darwin_(operating_system))
Accessed: 2016-09-30.
- [10] [Wikipedia - linux history.](#)
[https://en.wikipedia.org/wiki/History_of_Linux.](https://en.wikipedia.org/wiki/History_of_Linux)
Accessed: 2016-09-30.
- [11] [Wikipedia - usage share of operating systems.](#)
[https://en.wikipedia.org/wiki/Usage_share_of_operating_systems.](https://en.wikipedia.org/wiki/Usage_share_of_operating_systems)
- [12] [RAYMOND, E.](#)
The cathedral and the bazaar, vol. 12.
Springer, 1999.