

Operating Systems

Daniel Gruss

2024-03-07

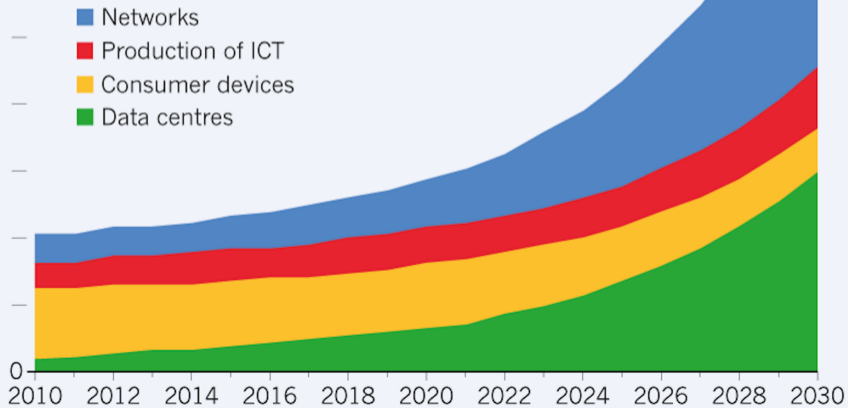
Why OS?

9,000 terawatt hours (TWh)

©nature

ENERGY FORECAST

20.9% of projected
electricity demand



It's 2030. You get to make an OS design decision:

one page table per process


vs

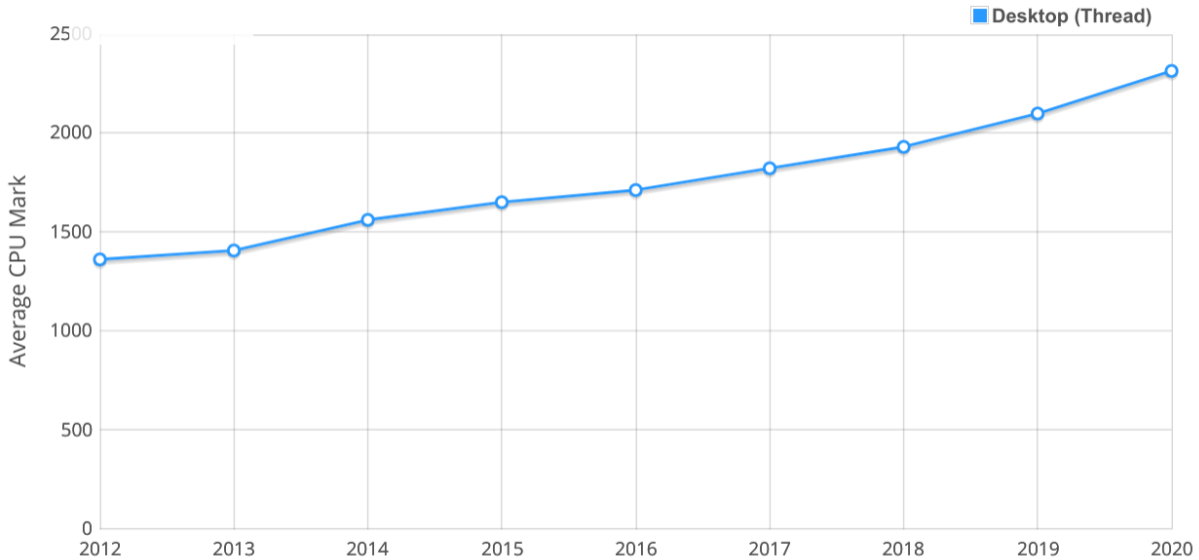
half a page table per process + 1 for the kernel

0.1%

Moore's Law – The number of transistors on integrated circuit chips (1971-2018)

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important as other aspects of technological progress – such as processing speed or the price of electronic products – are linked to Moore's law.








YESS...

MOAR SOFTWARE COMPLEXITY




SWEB is C and C++

Which programming language would you prefer?

Table 4. Normalized global results for Energy, Time, and Memory

Total					
Energy		Time		Mb	
(c) C	1.00	(c) C	1.00	(c) Pascal	1.00
(c) Rust	1.03	(c) Rust	1.04	(c) Go	1.05
(c) C++	1.34	(c) C++	1.56	(c) C	1.17
(c) Ada	1.70	(c) Ada	1.85	(c) Fortran	1.24
(v) Java	1.98	(v) Java	1.89	(c) C++	1.34
(c) Pascal	2.14	(c) Chapel	2.14	(c) Ada	1.47
(c) Chapel	2.18	(c) Go	2.83	(c) Rust	1.54
(v) Lisp	2.27	(c) Pascal	3.02	(v) Lisp	1.92
(c) Ocaml	2.40	(c) Ocaml	3.09	(c) Haskell	2.45
(c) Fortran	2.52	(v) C#	3.14	(i) PHP	2.57
(c) Swift	2.79	(v) Lisp	3.40	(c) Swift	2.71
(c) Haskell	3.10	(c) Haskell	3.55	(i) Python	2.80
(v) C#	3.14	(c) Swift	4.20	(c) Ocaml	2.82
(c) Go	3.23	(c) Fortran	4.20	(v) C#	2.85
(i) Dart	3.83	(v) F#	6.30	(i) Hack	3.34
(v) F#	4.13	(i) JavaScript	6.52	(v) Racket	3.52
(i) JavaScript	4.45	(i) Dart	6.67	(i) Ruby	3.97
(v) Racket	7.91	(v) Racket	11.27	(c) Chapel	4.00
(i) TypeScript	21.50	(i) Hack	26.99	(v) F#	4.25
(i) Hack	24.02	(i) PHP	27.64	(i) JavaScript	4.59
(i) PHP	29.30	(v) Erlang	36.71	(i) TypeScript	4.69
(v) Erlang	42.23	(i) Jruby	43.44	(v) Java	6.01
(i) Lua	45.98	(i) TypeScript	46.20	(i) Perl	6.62
(i) Jruby	46.54	(i) Ruby	59.34	(i) Lua	6.72
(i) Ruby	69.91	(i) Perl	65.79	(v) Erlang	7.20
(i) Python	75.88	(i) Python	71.90	(i) Dart	8.64
(i) Perl	79.58	(i) Lua	82.91	(i) Jruby	19.84



You and SWEB

Organizational Details

SWEB

Booting SWEB

Back to Assignment 1

You and SWEB

Organizational Details


SWEB

Booting SWEB

Back to Assignment 1



- Knowledge from earlier lectures
 - ESP, OOP1(!), SLP(!!), CON, ...
 - **Everything** taught in SLP, OOP1, and ESP will be practically applied **by you** in OS
- Reasonable C/C++ experience
- Team work + Time management



- Depends significantly on your knowledge and experience ...
- ... and on your team members.
- Short good solutions add around 4000 lines of code → 1000 lines of code per member



- No one should be afraid of OS.
- We even added a course to prepare you for OS (and other courses): SLP!

→ SLP is sort of the “booster” for OS

- Statistics confirm: if you did well in SLP, OS becomes very much doable
- Steep learning curve in the beginning? No worries, you will manage!




Note to students

While we have made an effort to simplify these projects for you, most Duke students find these projects sufficiently difficult to dominate their lives during the one-semester course. A common misconception from earlier semesters is that we are sadistic individuals who enjoy seeing students suffer. Actually, this is not the case. We enjoy seeing students who are proud of what they have accomplished and excited by the power that flows from a relatively small set of simple abstractions in an operating system, even a toy one like Nachos.


Wie beurteilen Sie Ihren für diese LV nötigen Arbeitsaufwand - gemessen an den ECTS-Punkten, die Sie erhalten? /
How would you rate the workload for this course - measured against the ECTS credit points you received? (71 x answered)

sehr niedrig / very low (-2)	0.0% (0x)
niedrig / low (-1)	2.8% (2x)
angemessen / appropriate (0)	9.9% (7x)
hoch / high (1)	28.2% (20x)
sehr hoch / very high (2)	59.2% (42x)



Wie beurteilen Sie die inhaltlichen Anforderungen (Schwierigkeit) dieser LV? /
How would you rate the content requirements (difficulty) of this course? (71 x answered)

sehr niedrig / very low (-2)	0.0% (0x)
niedrig / low (-1)	0.0% (0x)
angemessen / appropriate (0)	14.1% (10x)
hoch / high (1)	40.9% (29x)
sehr hoch / very high (2)	45.1% (32x)



Wie zufrieden sind Sie mit der LV insgesamt? /

How satisfied are you with this course in general? (71 x answered)

sehr unzufrieden / <i>highly dissatisfied</i> (1)	<div style="width: 2.8%;"><div style="width: 2.8%;"></div></div> 2.8% (2x)
(2)	<div style="width: 11.3%;"><div style="width: 11.3%;"></div></div> 11.3% (8x)
(3)	<div style="width: 8.5%;"><div style="width: 8.5%;"></div></div> 8.5% (6x)
(4)	<div style="width: 11.3%;"><div style="width: 11.3%;"></div></div> 11.3% (8x)
(5)	<div style="width: 33.8%;"><div style="width: 33.8%;"></div></div> 33.8% (24x)
sehr zufrieden / <i>highly satisfied</i> (6)	<div style="width: 32.4%;"><div style="width: 32.4%;"></div></div> 32.4% (23x)

Mean value: 4,59 Standard dev.: 1,44



Wie stark trifft die folgende Aussage Ihrer Meinung nach auf die LV zu?:


"In dieser LV werden alle Studierenden gleichermaßen fair behandelt (unabhängig von Lerntyp, Geschlecht, Ethnie)" /

How well does the following statement describe this course, in your opinion?



"On this course, all students are treated fairly and equally (regardless of their learning style, gender and ethnicity)." (71 x answered)

trifft überhaupt nicht zu / <i>not at all well</i> (1)	<div style="width: 2.8%;"><div style="width: 2.8%;"></div></div> 2.8% (2x)
(2)	<div style="width: 2.8%;"><div style="width: 2.8%;"></div></div> 2.8% (2x)
(3)	<div style="width: 1.4%;"><div style="width: 1.4%;"></div></div> 1.4% (1x)
(4)	<div style="width: 4.2%;"><div style="width: 4.2%;"></div></div> 4.2% (3x)
(5)	<div style="width: 19.7%;"><div style="width: 19.7%;"></div></div> 19.7% (14x)
trifft völlig zu / <i>very well</i> (6)	<div style="width: 69.0%;"><div style="width: 69.0%;"></div></div> 69.0% (49x)

Mean value: 5,42 Standard dev.: 1,15



I was well prepared for this course from previous courses. (70 x answered)



Comments on "trifft zu"


- Nur SLP konnte hier wirklich als Vorbereitung angesehen werden, da alle sonstigen LVen eigentlich nicht viel mit Low-Level Development zu tun hatten.
- Vor allem SLP
- SLP
- As a somewhat unrelated point I'd just like to say that I think prof. Gruss puts way too much stock into A0 as an indicator of preparedness for SLP. From what I've seen actual grades at SLP and OS haven't really gotten worse in recent times.
- SLP bereitet einen perfekt vor

Comments on "trifft eher zu"


- SLP für locking, Threads und shared Memory und C++
- I would say i was prepared, but of course a lot of new topics
- SLP ja
- Der einzige Kurs, welcher als Vorbereitung anzusehen ist, ist SLP. Die Kurse davor haben meines Erachtens sehr wenig für OS beigetragen.
- schon eher gut vorbereitet, durch SLP aber könnte besser sein
- slp war enorme Hilfe
- Durch SLP kriegt man schon eine gute Vorbereitung, doch es wart trotzdem recht schwer.
- I lacked in: basics like good git-handling in a group, good use of debugging and asserts, etc
- SLP, OOP und ESP prinzipiell sehr Hilfreich. Aber die LV war am anfang schon etwas überwältigend
- In SLP könnte mehr Perspektive gegeben werden, wie genau die Inhalte in OS dann gebraucht werden einfach ein kleiner Spoiler, dass man leichter das große Ganze versteht.

Der Stil der Abgabegespräche mit meinem Tutor war (70 x answered)

sehr angenehm (3) 5 Comm.
ok (2) 3 Comm.
unangenehm wegen (1)
eine Qual weil (0) 1 Comm.



Mean value: 2,81 Standard dev.: 0,49



Comments on "sehr angenehm"

- our tutor did not make a big deal out of the assignment interview and it did help me feel more relaxed, even though i shat my pants every times before it. in the interview after talking about our code our tutor said ok great now only the coding example and then youre done. i really liked how they made the coding example out to be very easy and definitely doable, it took the stress away a bit. of course our tutor graded us accordingly afterwards and gave deductions for mistakes, so they didnt just give away free points.
- -
- not to difficult or stressful and the tutor really wanted to see our understanding of the os, which seems like the best approach
- Da er ein sehr gemütlicher Typ war, waren dementsprechend auch seine Abgabegespräche, sofern man auch mitgearbeitet hat.
- Even as the tutor was absolutely eviscerating us in A1 going through RC after RC it felt like a great learning experience. Both tutors gave us chance to argue for more points and we actually got a few in A2. My only complaint is that I think some of the A2 AG tasks my teammates got were too difficult but I wrote about that elsewhere. Additionally, I really appreciated that we got a git patch with comments highlighting the mistakes we made in the first assignment.

Gab es Lehrveranstaltungen, die Ihnen als besonders POSITIV in Erinnerung geblieben sind? Wenn ja, benennen Sie diese bitte im Textfeld und begründen Sie Ihre Antwort. (18 x answered)

ja	15 Comm.	<div style="background-color: #008000; color: white; padding: 2px;">83.3% (15x)</div>
nein		<div style="background-color: #008000; color: white; padding: 2px;">16.7% (3x)</div>

Comments on "ja"

- OS
- Betriebssysteme. Es war zwar super schwierig, aber in keinem anderen Fach hab ich so viel gelernt und auch nichts so spannendes.
- Objektorientierte Programmierung 2, Projektmanagement, Agile Software Development
- InfoSec
- 373.336 Bürgerliches Recht und Unternehmensrecht / 373.314 Betriebssoziologie; Sehr spannende Aufbereitung des Stoffes mit vielen praxisnahen Beispielen; 716.062 Softwareentwicklung und Wissensmanagement; Interessante Lv-Organisation
- InfoSec, CompStat, Web Technology
- Operating Systems - super organisiert und man lernt extrem viel - auch wenn es sehr aufwendig ist
- Operating Systems
- Information Security
- Operating Systems, System Level Programming, Objektorientierte Programmierung 1 & 2, Computer Grafik & Vision - Mit Abstand die interessantesten Fächer meines Bachelorstudiums. Diese hatten das vermittelte Wissen, das man sich in sehr vielen (anderen) Fächern angeeignet hatte, am besten in die Praxis umgesetzt mit sehr sehr interessanten Projekten!
- Bachelorarbeit und Operating systems
- Auch wenn man sehr gefordert wurde, war Betriebssysteme doch eine sehr gute LV.
- Fundamentals of Geometry Processing
- Ja, OS, weil super präsentiert. FoCS, Prof Slany hat einen super Vortragsstil für das erste Semester.
- System Level Programming, Rechnernetze und Organisation, Software Architecture

Vielleicht hätte man hervorheben sollen, dass uSTL NICHT kompatibel zu STL ist! `ustl::list` gehört eigentlich verboten. Man kommt in Teufels Küche, wenn man von ihr ableiten möchte. Man gibt sich einer Illusion hin, die nicht stimmt. Einfügen und Löschen sind NICHT effizient und Iteratoren verlieren die Gültigkeit, wo sie es nicht sollten. Trotzdem war es hilfreich.

Check the standard!

```
iterator erase (iterator position);  
iterator erase (iterator first, iterator last);
```

erase verwenden ohne den iterator upzudaten geht mit der GNU C++ STL *meistens...*
→ gefährlich sich anzugewöhnen den iterator zu ignorieren

Conclusion

To conclude, we can get some facts about each data structure:

- `std::list` is very very slow to iterate through the collection due to its very poor spatial locality.
- `std::vector` and `std::deque` perform always faster than `std::list` with very small data
- `std::list` handles very well large elements
- `std::deque` performs better than a `std::vector` for inserting at random positions (especially at the front, which is constant time)
- `std::deque` and `std::vector` do not support very well data types with high cost of copy/assignment

This draw simple conclusions on usage of each data structure:

- Number crunching: use `std::vector` or `std::deque`
- Linear search: use `std::vector` or `std::deque`
- Random Insert/Remove:
 - Small data size: use `std::vector`
 - Large element size: use `std::list` (unless if intended principally for searching)
- Non-trivial data type: use `std::list` unless you need the container especially for searching. But for multiple modifications of the container, it will be very slow.
- Push to front: use `std::deque` or `std::list`

[http://baptiste-wicht.com/posts/2012/12/
cpp-benchmark-vector-list-deque.html](http://baptiste-wicht.com/posts/2012/12/cpp-benchmark-vector-list-deque.html)



- We try to improve our support constantly
- Feedback, Evaluations

“We will not lower the bar, but we will do what we can to help you over it.”



- Bad time management
- Pair programming (it's not efficient enough!!)
- Problems with working in a team
- No C/C++ experience
- Not willing to learn and use C/C++

You and SWEB

Organizational Details

SWEB

Booting SWEB

Back to Assignment 1



- Website: <https://iaik.tugraz.at/os>
- Discord (tag your tutor, also in your own group channels!)
- Email: os@iaik.tugraz.at
- Consultation hour: send us an email
- *During the question days: slower response times on Discord/email ;)*



- Take place before the design deadlines
- You are expected to bring a sketch/summary of your design
- You are expected to have done proof of concept implementations by then
- Prepare to defend your ideas



- Only code! No design document!
- Proof of concept implementation
- You get up to 5 points, depending on the test system score of your submission.



- You have to:
 - push regularly into the provided repository
 - use your real name and email address for commits




- Personal maximum number of points
- Unlock points:
 - For each commit you unlock 0.5 points
 - For each 10 LoC added you unlock 0.5 points
 - Max. 10 points per day can be unlocked
- Average: > 55 points unlocked
- Test system shows your current personal maximum number of points



- We even give you master access to your repository, **but:**
 - renaming repo = your group fails the course
 - changing repo path = your group fails the course
 - adding/removing group members = your group fails the course



Every team member must understand all parts and contribute points to the mandatory task categories!



- Set internal deadlines for your team members
 - No idea how to split the work and choose good internal deadlines? Ask tutor for recommendations!
- Missed deadline? → re-assign task
- Repeatedly missed deadline → tell tutor to remove team member from team

What if you as a team don't fit together?



- Team member(s) not contributing enough → contact tutor via mail!
 - There will be more students with the same problem as you have
 - Team doesn't want to split but you want to leave? → also possible
- Merges are possible
- Don't pull anyone through
- small number of bonus points as compensation (if **we** think it's appropriate)

Not happy with your tutor's resolution? → **Talk to me**, we'll figure out a solution

- Existing syscall wrapper functions are POSIX-compatible
 - Do not change signatures!
 - We have automated tests using the POSIX interface
- If you implement new functions, try to make signatures POSIX-compatible

- Writing **numerous** test programs is unavoidable
- Test programs are supposed to show whether your implementation works according to the assignment
- You will get points for the test programs
- We have our own secret test programs



**REALIZING THAT YOU
ONLY NEED ABOUT 20 MORE
HOURS OF WORK TO GET A 1**

IT'S 2 HOURS BEFORE THE DEADLINE

- Tag the commit you want to submit: `git tag SubmissionD1 [commit_hash]`
- Push to repository: `git push / git push --tags`
- Read the commit ID:

```
git show SubmissionD1
commit 196bc4a704f37d7f969d27a258b513693e3b30f4
Author: Daniel Gruss <daniel.gruss@iaik.tugraz.at>
```

Test System will acknowledge your submission!

- Small fixes → small deduction
- Do this at the same time:
 - Mail to os@iaik.tugraz.at!
 - Submit and retag your fixed version!

Reference (=100%)

- Design: 5 points
- Task 1: 20 points
- Task 2: 10 points
- Elective tasks*: 15 points (or more)

Minimum

- Design: 1 point
- Mandatory tasks: 15 points
- And in total: 25 points (=50%)

(*) See the list of elective tasks on the Website!

Reference (=100%)

- Design: 5 points
- Mandatory tasks: 40 points
- Elective tasks*: 5 points (or more)

Minimum

- Design: 1 point
- Mandatory tasks: 15 points
- And in total: 25 points (=50%)

(*) See the list of elective tasks on the Website!

Minimum requirements

- A1: 25 of 50 points
- A2: 25 of 50 points

Limits

- A1: max. 60 points
 - Unlimited if $A2 \geq 40$ points
- A2: unlimited points



- mid-term exam: 35%
- the practicals: 65%



To pass the class, you have to acquire

- overall at least 55%
- at least 50% of the possible points on written exam
- at least 50% of the possible points in the practicals



- Didn't pass any assignment or exam?
- No worries, we have a second chance for everyone.
- DD → you can continue without Second Chance but with 0 DD/PoC points
- PoC → Second Chance PoC (examined by Lecturers)
- Exercise → Second Chance Exercise (examined by Tutors + Lecturers)
- Exam → Second Chance Exam

Getting a positive grade remains as easy, getting a 1 gets harder (with every second chance)



Points

1. Take the team score
2. cap it by the personal unlocked points
3. apply DD deduction
4. apply AG deduction

→ final points for assignment



Points (Example)

1. You get 48 points as a team → 48
2. Your personal unlocked points are 43 → 43
3. Your group had 4 points on the PoC but you didn't say enough in the DD → 39
4. You get a 75% deduction because you couldn't really figure out where to do what in your SWEB in the AG → 10 points → second chance



Points (Example)

1. You get 48 points as a team → 48
2. Your personal unlocked points are 57 → 48
3. You have a 25% deduction because you didn't say enough in the DD → 36
4. You get a 20% deduction because you took too long with making changes → 29
→ barely positive



Points (Example)

1. You get 45 points as a team → 45
2. Your personal unlocked points are 70 → 45
3. You participated in the DD properly → 45
4. You get a 5% deduction because you took rather long with making changes → 43
→ good grade

Marks

- genügend: 55-66
- befriedigend: 67-78
- gut: 79-89
- sehr gut: 90+



We want you to evaluate. From 30% upwards: for every 10 percentage points, **everyone** needs 1 point less for a better positive grade!

As a group

- You explain what you implemented and how you tested it
- Together with the teaching assistant you determine the points of your group

As a group member

- You are able to read, explain and change the code
- You can implement small new features or extend existing ones
 - just quickly, no rigorous testcase writing
 - tutor will stop you as soon as it's clear that you deserve the points
- Otherwise you will get less points

- Discussions with other teams are appreciated
- But: **no collaboration!**
- We check for plagiarism
- Similarities → teams are questioned
- Both teams: 0 points in either case
- At least one team: “Ungültig/Täuschung” (with all its consequences)

- Do not provide your source code to other teams
- Make sure your source code is protected against unintended access from others
- Do not use source code from previous years
 - Code from another team → plagiarism
 - Your own code (not exactly the same team)
→ not allowed
 - Your own code (exactly the same team)
→ allowed, probably not the best idea ;)

Questions so far?

- Tutors want to help you.
- Tutors have a competition:
 - The higher their students score **on average**, the better they score!
→ they want to win, they want to make you pass, help them!
- You don't know whether you have enough points yet? Don't ask the tutor for your points, but ask them:
“should we do more? would you recommend that we still do something? what would you recommend we do?”

You and SWEB

Organizational Details

SWEB

Booting SWEB

Back to Assignment 1



- VU Amsterdam: Minix (1987), Minix3 (2005)
- Berkeley: Nachos (1992)
- Stanford: Pintos (2004)
- Graz: SWEB (2006)
 - BS KU W04
 - Philip Lawatsch, Bernhard Tittelbach, many other students and projects



- Minimalistic operating system kernel
- Runs on x86-32, x86-64, ARM
- Emulated using qemu
- Important features are missing
- Your task: Make your SWEB a beautiful, feature-rich kernel

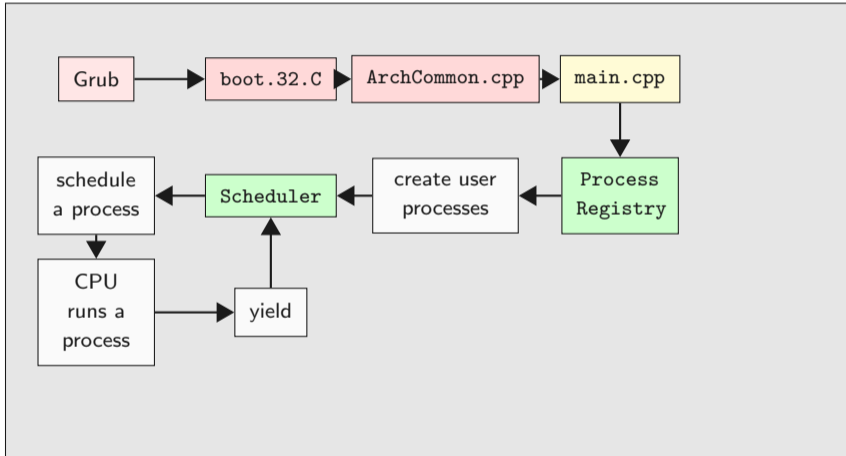
You and SWEB

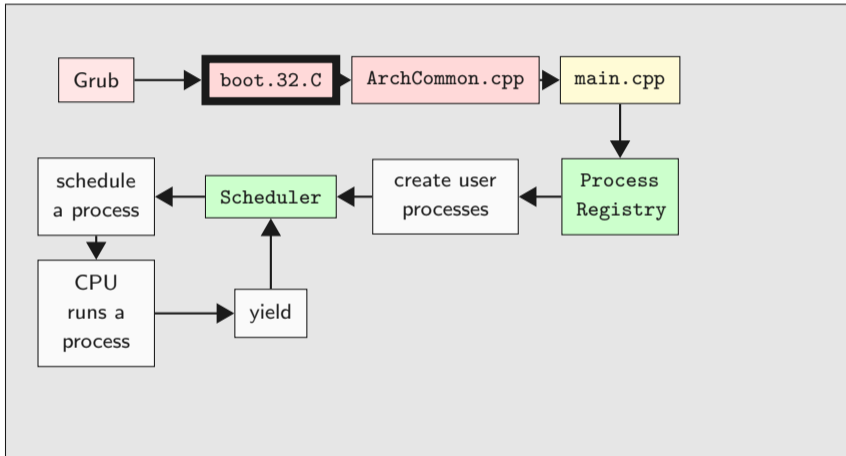
Organizational Details

SWEB

Booting SWEB

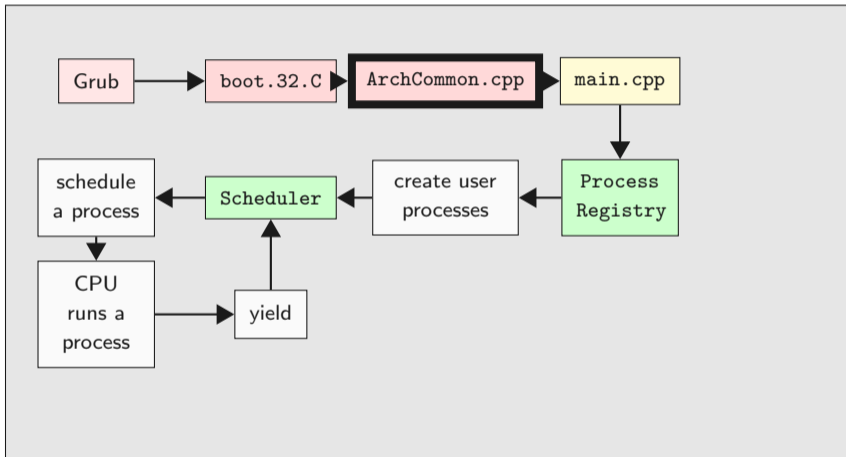
Back to Assignment 1






```
extern "C" void entry()
{
    PRINT("Booting...\n");
    // ...
    PRINT("Initialize Kernel Paging Structs\n");
    // ...
    PRINT("Enable Paging...\n");
    // ...
    PRINT("Calling entry64()...\n");
    asm("ljmp %[cs], $entry64-BASE\n" : : [cs]"i"(KERNEL_CS));
}
```

- 32 bit
- Works on physical addresses
- Setup hardware
- Setup paging
- Jump into sane, virtual C world



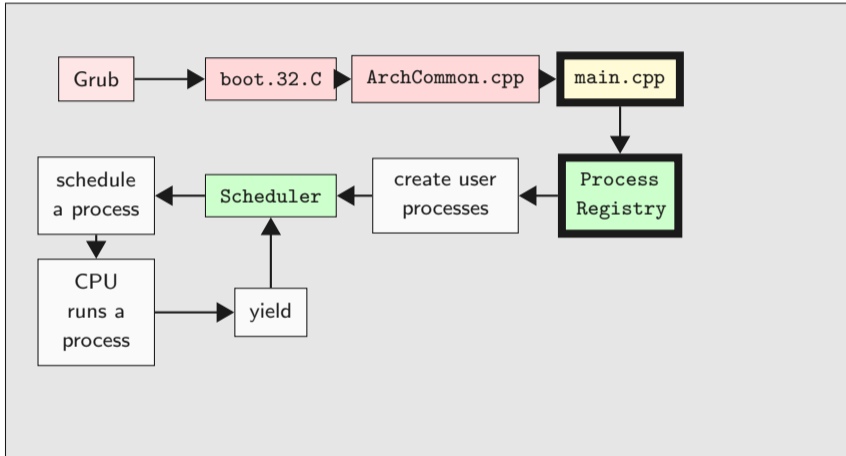

```
extern "C" void entry64()
{
    // ...
    PRINT("Switch to our own stack...\n");
    asm("mov %[stack], %%rsp\n"
        "mov %[stack], %%rbp\n" : : [stack]"i"(boot_stack + 0x4000));
    PRINT("Loading Long Mode Segments...\n");
    // ...
    PRINT("Calling startup()...\n");
    asm("jmp *%[startup]" : : [startup]"r"(startup));
}
```



```
extern "C" void startup() {
    removeBootTimeIdentMapping();
    system_state = BOOTING;
    PageManager::instance();
    writeLine("PageManager and KernelMemoryManager created\n");
    // ...
    Scheduler::instance()->addNewThread(new ProcessRegistry(...));
    // ...
    system_state = RUNNING;
    ArchInterrupts::enableInterrupts();
    Scheduler::instance()->yield();
    //not reached
}
```



- Setup kernel objects
- Setup more hardware
- Enable interrupts
- Handover control to Scheduler




```
void ProcessRegistry::Run()
{
    debug(PROCESS_REG, "mounting userprog-partition \n");
    // ...
    vfs_syscall.mount("idea1", "/usr", "minixfs", 0);
    debug(PROCESS_REG, "mount idea1\n");

    for (size_t i = 0; progs_[i]; i++)
    {
        createProcess(progs_[i]);
    }
}
```



- Kernel Thread
- Mounts hard disk
- Creates user processes
- Sleeps until last user process died, ...
- ... and unmounts the hard disk again





Status quo

- Derived from Thread
- Bad design is easier for you to improve ;)
- Executes binary code
- Has a virtual address space (Loader → ArchMemory)
- Has a userspace part and a kernel part

- Every 54.925439ms
- Implemented in InterruptUtils.cpp

```
extern "C" void irqHandler_0()  
{  
    ArchCommon::drawHeartBeat();  
    Scheduler::instance()->incTicks();  
    Scheduler::instance()->schedule();  
    ArchInterrupts::EndOfInterrupt(0);  
    arch_contextSwitch();  
}
```

- List of threads
- `schedule()` on IRQ0
- `schedule()` on IRQ65 (yield)
- `void Scheduler::addNewThread(Thread *thread);`
- Contains `IdleThread` (hlt if idle)
- Contains `CleanupThread`
 - Calls `delete` on dead threads

A process is scheduled:

- only if `switch_to_userspace_ == 1`
 - Scheduler loads register values from `ArchThreadRegisters` member variable
- Implicitly sets RIP from `ArchThreadRegisters`
- RIP initially points to binary entry point
 - CPU switches to user mode and continues with given RIP



```
userspace/libc/src/nonstd.c
```

```
extern int main();
```

```
void _start()
```

```
{
```

```
    exit(main());
```

```
}
```

libc is in userspace!



```
userspace/tests/helloworld.c
#include <stdio.h>
int main()
{
    puts("hello , world");
    return 0;
}
```

- Off-the-shelf Hello World Program
- Brian Kernigham, 1974

userspace/libc/printf.c

```
int puts(const char *output_string)
{
    // ...
    characters_written = write(STDOUT_FILENO, (void*) output_string,
    string_length);
    // ...
}
```

- In libc again
- Sanity checks
- Actually only a wrapper for syscall write

```
userspace/libc/write.c
ssize_t write(int file_descriptor, const void *buffer, size_t count)
{
    return __syscall(sc_write, file_descriptor,
                    (long) buffer, count, 0x00, 0x00);
}
```

- POSIX!
- fwrite is a simple write wrapper
- Down the rabbit hole: __syscall



```
arch/x86/64/common/userspace/syscalls.c
size_t __syscall(size_t arg1, size_t arg2, size_t arg3,
                 size_t arg4, size_t arg5, size_t arg6)
{
    asm("int $0x80\n" : "=a"(arg1)
        : "a"(arg1), "b"(arg2), "c"(arg3),
          "d"(arg4), "S"(arg5), "D"(arg6));
    return arg1;
}
```

- Copy arguments to registers
- Issue interrupt 0x80

→ switch to kernel space

```
arch/x86/64/common/source/arch_interrupts.S
```

```
arch_syscallHandler:
```

```
    pushall
```

```
    movq %rsp,%rdi
```

```
    movq $0,%rsi
```

```
    call arch_saveThreadRegisters
```

```
    call syscallHandler
```

InterruptUtils.cpp

```
extern "C" void syscallHandler() {
    thread->switch_to_userspace_ = 0;
    threadRegisters = thread->kernel_regs_;
    ArchInterrupts::enableInterrupts();

    auto ret = Syscall::syscallException(
        thread->user_regs_->rax,
        thread->user_regs_->rbx,
        thread->user_regs_->rcx,
        thread->user_regs_->rdx,
        thread->user_regs_->rsi,
        thread->user_regs_->rdi);
}
```

Syscall.cpp

```
size_t Syscall::syscallException(size_t syscall_number, size_t arg1,
    size_t arg2, size_t arg3, size_t arg4, size_t arg5) {
    switch (syscall_number)
    {
        // ...
        case sc_write:
            return_value = write(arg1, arg2, arg3);
            break;
        // ...
    }
}
```

syscallException only calls the right methods with the right number of parameters

- `sc_write` constant defined in `Syscall.h`
- `Syscall::write` method is regular in-kernel C++

```
size_t Syscall::write(size_t fd, pointer buffer, size_t size)
{
    if (fd == fd_stdout) //stdout
    {
        debug(SYSCALL, "Syscall::write: %.*s\n", (int)size, (char*) buffer);
        kprintf("%.*s", (int)size, (char*) buffer);
    }
    // ...
}
```

```
common/include/console/debug.h
const size_t MAIN          =Ansi_Red      | ENABLED;
const size_t THREAD        =Ansi_Magenta | ENABLED;
const size_t USERPROCESS  =Ansi_Cyan     | ENABLED;
const size_t PROCESS_REG  =Ansi_Yellow   | ENABLED;
const size_t BACKTRACE    =Ansi_Red      | ENABLED;
const size_t USERTRACE    =Ansi_Red      | ENABLED;
//group memory management
const size_t PM           =Ansi_Green    | ENABLED;
const size_t KMM          =Ansi_Yellow;
```

Add your own debug tags!

“This is not the printf you are looking for”



- `kprintf` → SWEB Terminal
 - `kprintfd` → Host
 - Cannot be disabled by flags
- Prefer debug


```
#include <stdio.h>
int main()
{
    puts("hello, world");
    return 0;
}
```

What does `return 0;` do?

```
userspace/libc/src/nonstd.c
```

```
extern int main();
```

```
void _start()
```

```
{
```

```
    exit(main());
```

```
}
```

AANNND



IT'S GONE

```
size_t Syscall::syscallException(...) {
    // ...
    case sc_exit:
        exit(arg1);
        break;
    // ...
}
void Syscall::exit(size_t exit_code) {
    debug(SYSCALL, "Syscall::EXIT: called, exit_code: %zd\n", exit_code)
    ;
    currentThread->kill();
}
```

```
void Thread::kill() {
    switch_to_userspace_ = 0;
    state_ = ToBeDestroyed;
    if (currentThread == this) {
        ArchInterrupts::enableInterrupts();
        Scheduler::instance()->yield();
    }
}
```

Recall: Scheduler will call delete on thread if `state_ == ToBeDestroyed`

```
UserProcess::~~UserProcess() {  
    delete loader_;  
    vfs_syscall.close(fd_);  
    delete working_dir_;  
    process_registry_ ->processExit();  
}  
Thread::~~Thread() {  
    delete user_registers_;  
    delete kernel_registers_;  
}
```



THAT'S IT, I'VE HAD ENOUGH

You and SWEB

Organizational Details

SWEB

Booting SWEB

Back to Assignment 1



- Mouse driver
- Window manager
- Network driver
- Soundblaster driver
- Gameboy emulator
- 3D game engine
- Running it on a small ARM board with only 256KB RAM



- Try out the tutorials on <http://iaik.tugraz.at/os>
 - Set up development environment
 - Implement your first syscall
- Get acquainted with the source code: Try out implementing things in SWEB
- Start with the practicals **NOW!** (or rather already a week ago)

- Base line SWEB: a user process **is a** (kernel) thread
- We want: multiple threads per user process
- What do we have to change?

- Each thread has its own instances of some resources
 - id, stack, registers, status, ...
- Other resources are shared among all threads
 - memory, files, ...

- How to “use” multithreading?
- Syscalls!
- Which ones?
- You decide - but function names and arguments have to be POSIX-compatible!
- Minimum requirements:

`pthread_create`

`pthread_exit`

`pthread_cancel`

`pthread_join`

- By definition: Operating system is written by people who know what they do
- User programs?
- System calls provide a “safe” interface
- Control flow is transmitted to kernel code
- Typical syscalls: `fork()`, `read()`, `write()`, `execve()`, `wait()`, `exit()`
- You will step through a syscall in the tutorial this week!

- `fork()` creates a new process by duplicating the calling process
- The new process (=the child), is an exact duplicate of the calling process (=the parent)
- Interesting in combination with multithreading!

- Replaces the current process image with a new process image
- `exec` with arguments → more points
- `fork()/exec()` combination often used



- `sleep()` sets a thread asleep for a given number of seconds
- `clock` returns how much cpu time the current process consumed



- `exit()` terminates the current process
- Already implemented, but ...
- ... you will break the current implementation with multithreading

- I/O syscalls are already implemented, but ...
- ... they use global (not process specific) file descriptors
- Why is that a problem?



- Threads need synchronization
- Kernel has mutexes and condition variables
- We want: mutexes, condition variables and semaphores, both in kernelspace and userspace
- Pure userspace implementation (except for initialization and for going to sleep)
- Implement test programs (Readers-Writers-Problem, Sleeping Barber, etc.)



- Own ideas are the most fun!
- See <https://www.iaik.tugraz.at/teaching/materials/os/assignments/> for suggestions
- Please note: Assignment 1 tasks will **only** be counted in Assignment 1 assessment
Assignment 2 tasks will **only** be counted in Assignment 2 assessment

Assignment 1

- Deadline Design-PoC: **See Timeline**

- The design proof-of-concept comprises your implementation so far, as is. It will be graded by the test system and you will talk 30 minutes with your tutor about your progress.

- Deadline Implementation: **See Timeline**

- Minimum requirements:

- At least **15** points from the two mandatory parts (multi threading and fork)
- At least **1** point from the design proof-of-concept
- At least **25** points in total

- 100% equals 50 points

- [Tutors Points List A1 \(not intended for student use\)](#)

In Assignment 1 you will acquire basic knowledge about the inner workings of multi threading and multi processing on modern operating systems.

Task 1: Multi Threading

You will start with the [basic SWEB kernel](#) which has neither multi threading nor multi processing implemented for userspace programs. Instead there is one thread per process. You have to figure out what to change to make SWEB support multi threading.

Threads share their address space with all other threads within the same process. However, each thread has its own stack, own set of cpu register values and maybe other resources. When the operating system interrupts a thread, it has to store these resources so that the thread execution can be continued later.

Interaction between userspace and kernelspace is only possible through syscalls. You can implement any syscall you like. However, for reasons of automated testing of your implementation you are required to **comply with the POSIX standard in terms of function names and parameters**.

You will need to implement at least the following syscalls:

1. `pthread_create` to create a new thread
2. `pthread_exit` to allow a thread to terminate itself
3. `pthread_cancel` to terminate another thread
4. `pthread_join` to wait for termination of another thread

Furthermore you are required to implement several user programs which demonstrate that your implementations works as expected.

Task 2: Fork

We want to give our user programs the ability to start new user programs. The POSIX compatible way of doing so, is to use the fork syscall (and the exec syscall). Fork creates an exact copy of the currently running process, with only the calling thread copied to the new process. The copied thread in the new process continues execution right after the fork call. We want to remind you to comply exactly to the POSIX standard for reasons of automated testing.

Additional Tasks

The two mandatory tasks are designed to give 30 points (of 50 points). We expect you to choose your own additional tasks. In the [Tutors Points List A1 \(not recommended for student use\)](#) you find a long list with numerous suggestions what you could implement. **DON'T** just use this list. Neither you nor we will be happy with the result. These are bonus tasks, come up with your own ideas, come up with your own improved variants of features.

We strongly encourage implementing your own interesting ideas, as long as they are related to this course, by giving you points. But, before you start, ask your tutor. For tasks that fall into one of the topics of Assignment 2 you can't get any points in Assignment 1 (and the other way round).



- Don't use it.
- No explanation. No details on how much you have to do for which point.
- Not suitable to choose tasks.
- Not suitable to estimate your points (points vary a lot!).

- Tutors Point List says **2 points**
 - “ah, actually cancel only works up to 5 times, then it stops working” → 0-1 points
 - “actually, we have implemented it including cancelstate, canceltype, cleanup push, pthread keys, correct cancel and join interaction” without any problems → maybe 8-12 points
- **You see?** The **2 points** doesn't tell you much! Can be a lot more, can be a lot less, all depends on how much you do.
- 2 points refers to a minimal correct working function that implements the most basic functionality of that syscall function

- “Is it mandatory to implement every sentence and everything that is mentioned on the (for example) `pthread_cancel` man page?”
 - No. Minimal correctly working cancel is 2 points. Everything beyond (cancel types, cancel states, etc.) is all elective/bonus.
- **70%** of the points you will collect in the OS exercises will be elective/bonus.



Get those tutorials done asap and start coding!