

Microkernel Construction

M³ Tutorial

Viktor Reusch, Nils Asmussen

07/02/2026

Preparation



- Install *qemu* or *VirtualBox* (or anything else that supports *qcow2/VDI*)
- Download VM image matching your setup from the course website
- Extract the image (≈ 15 GiB): `tar xzf m3-tutorial....tar.gz`



- Install *qemu* or *VirtualBox* (or anything else that supports *qcow2/VDI*)
- Download VM image matching your setup from the course website
- Extract the image (≈ 15 GiB): `tar xzf m3-tutorial....tar.gz`
- Create a virtual machine (with 4 cores and 8 GiB of memory):
 - `qemu-system-x86_64 -enable-kvm -smp 4 -m 8G -drive if=virtio,format=qcow2,file=m3-tutorial.qcow2`
 - `qemu-system-aarch64 -enable-kvm -smp 4 -m 8G -drive if=virtio,format=qcow2,file=m3-tutorial.qcow2 -bios /usr/share/qemu-efi-aarch64/QEMU_EFI.fd -cpu max -M virt -device virtio-gpu-pci -device qemu-xhci -device usb-tablet -device usb-kbd`
 - Create new appliance in VirtualBox with VDI



- Install *qemu* or *VirtualBox* (or anything else that supports *qcow2/VDI*)
- Download VM image matching your setup from the course website
- Extract the image (≈ 15 GiB): `tar xzf m3-tutorial....tar.gz`
- Create a virtual machine (with 4 cores and 8 GiB of memory):
 - `qemu-system-x86_64 -enable-kvm -smp 4 -m 8G -drive if=virtio,format=qcow2,file=m3-tutorial.qcow2`
 - `qemu-system-aarch64 -enable-kvm -smp 4 -m 8G -drive if=virtio,format=qcow2,file=m3-tutorial.qcow2 -bios /usr/share/qemu-efi-aarch64/QEMU_EFI.fd -cpu max -M virt -device virtio-gpu-pci -device qemu-xhci -device usb-tablet -device usb-kbd`
 - Create new appliance in VirtualBox with VDI
- Open terminal, go to `$HOME/M3`
- If required, user name: `m3`

SSH Access



- You can access the VM via SSH for remote developing / *SFTP* / *SSHFS*
- Add to the QEMU command line: `-nic user,hostfwd=:127.0.0.1:2222-:22`
- Connect via: `ssh -p 2222 m3@127.0.0.1`

The Hands-on Part



- What you have to expect:
- Develop M³ applications
- Learn how to ...
 - Send messages via the TCU
 - Use system services
 - Communicate with accelerators
- Inside the ready-made VM



- Introduction
- **Assignments**
 - **Hello World**
 - IPC
 - Filesystem Service
 - Accelerators
 - Sessions

Build System



- *ninjabie*
 - Originally created for M³
 - Primarily required for C/C++, but Rust is integrated as well
 - Open source: <https://github.com/Barkhausen-Institut/Ninjabie>
 - Simple Python scripts



- *ninjabie*
 - Originally created for M³
 - Primarily required for C/C++, but Rust is integrated as well
 - Open source: <https://github.com/Barkhausen-Institut/Ninjabie>
 - Simple Python scripts
- Heavy lifting done by *ninja*
 - Very fast
 - Build files not designed for humans
 - Generated from Python scripts by *ninjabie*



- *ninjabie*
 - Originally created for M³
 - Primarily required for C/C++, but Rust is integrated as well
 - Open source: <https://github.com/Barkhausen-Institut/Ninjabie>
 - Simple Python scripts
- Heavy lifting done by *ninja*
 - Very fast
 - Build files not designed for humans
 - Generated from Python scripts by *ninjabie*
- Convenience script: `b`
 - Builds everything
 - Provides commands for running, debugging, analyzing, ...

Working with M³



- Environment variables
 - **M3_ISA**: `x86_64`, `riscv`
 - **M3_BUILD**: `release`, `debug`, `bench`
 - **M3_GEM5_CPU**: `Deriv03CPU`, `TimingSimpleCPU`
 - Already set for you in `.envrc`

Working with M³



- Environment variables
 - **M3_ISA**: `x86_64`, `riscv`
 - **M3_BUILD**: `release`, `debug`, `bench`
 - **M3_GEM5_CPU**: `Deriv03CPU`, `TimingSimpleCPU`
 - Already set for you in `.envrc`
- Commands
 - `./b`: build everything (no command)
 - `./b run <script>`: run given boot script (e.g., `boot/hello.xml`)
 - `./b -n clippy [path]`: Rust linting
 - `./b -n fmt`: format source code
 - `./b -n doc`: generate documentation
 - `./b -h`: detailed help

Build Scripts




- `/build.py`
 - Main file of Ninjapie
 - Generates ninja build file
 - Called by `b`



- /build.py
 - Main file of Ninjapie
 - Generates ninja build file
 - Called by b
- .../app/build.py
 - Called by /build.py
 - Describes how to build a specific program/library/...
 - Example for Rust application:

```
1 def build(gen, env):  
2     env.m3_rust_exe(gen, out = 'hello')
```

 Python

Building Rust With Cargo



.../app/Cargo.toml:

```
1  [package] toml
2  name = "hello"
3  version = "0.1.0"
4  edition = "2021"
5
6  # The M3 build system will link the final binary.
7  [lib]
8  path = "src/hello.rs"
9  crate-type = ["staticlib"]
10
11 [dependencies]
12 m3 = { path = "../..libs/rust/m3" }
```

Boot Scripts



/boot/...xml:

```
1  <config>
2    <kernel args="kernel"/> <!-- Start kernel -->
3    <dom> <!-- On another tile -->
4      <app args="root"> <!-- Start init process -->
5        <dom>
6          <app args="hello"/> <!-- Start hello application -->
7        </dom>
8      <dom>
9        <app args="hello"/>
10     </dom>
11   </app>
12 </dom>
13 </config>
```

XML

```
1 use m3::errors::{Code, Error}; // Imports
2 use m3::println;
3 pub fn my_func(param: &bool) -> Result<bool, Error> { // Function head
4     let mut my_var: isize = 128; // Declare and define mutable integer
5     other_func(&mut my_var); // Call function with mutable reference
6     let my_other = calc_func(&my_var, 12 as i32); // Integer conversion
7     if my_other.value < 128 { // Condition
8         println!("Got integer {}!", my_other.info()); // Print macro
9         return Ok(true); // Return enum variant
10    }
11    let ok_val = fallible_func().unwrap(); // Assert Result is Ok
12    Err(Error::new(Code::Abort)) // Construct error object and return
13 }
```



Application Code



/src/apps/.../src/... rs:

```
1  #![no_std] Rust
2
3  use m3::errors::Error;
4  use m3::println;
5
6  #[no_mangle]
7  pub fn main() -> Result<(), Error> {
8      println!("Hello world!");
9      Ok(())
10 }
```

Execution Demo

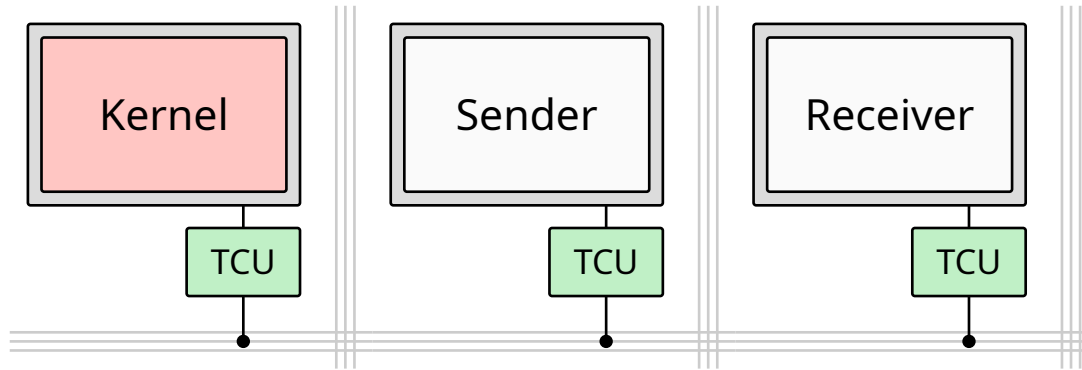


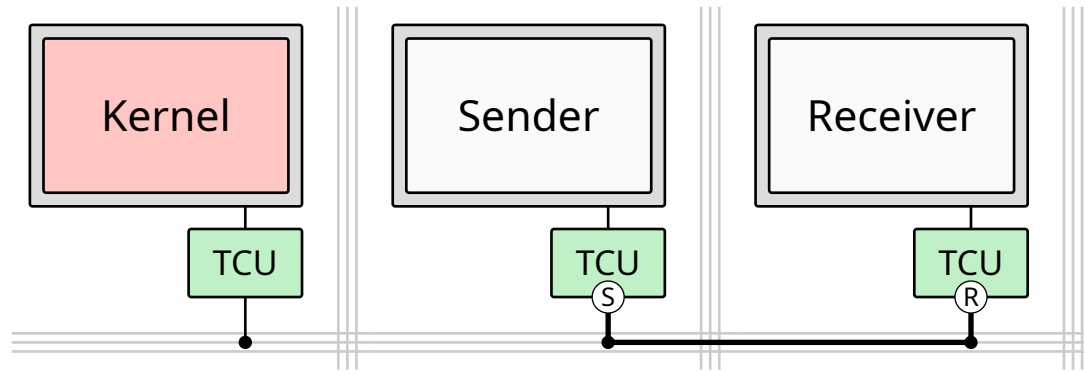
Assignment #1: Hello World




- Run the `rust-hello.xml` boot script
- It should print `Hello world!` in the terminal
- Modify the boot script to start two instances of the hello world application
- Modify the source code to print `Hello tutorial!` instead
- Run it again

- Introduction
- **Assignments**
 - Hello World
 - **IPC**
 - Filesystem Service
 - Accelerators
 - Sessions






Message Passing




Message Passing



SendGate

- Allows to **send messages** to the connected RecvGate
- If sender specifies a “reply gate”, receiver can reply to each message


Message Passing




SendGate

- Allows to **send messages** to the connected RecvGate
- If sender specifies a “reply gate”, receiver can reply to each message
- Has **credits**: each credit allows to send one message
- Sending removes one credit, receiving a reply adds one credit

Message Passing





RecvGate

- Allows to **receive messages** from multiple SendGates
- Has a receive buffer of N **slots** with M **bytes each**
- Sends fail if the message is larger than M or no slot is free
- If the credit system is used, the latter case can be avoided

Message Passing: ACKing messages




- Each received message is stored in a slot of the receive buffer
- The TCU needs to know **when slots can be reused**

Message Passing: ACKing messages



- Each received message is stored in a slot of the receive buffer
- The TCU needs to know **when slots can be reused**
- Replying on a message implicitly **acknowledges** this message
- If no reply is used, messages need to be acknowledged explicitly:

```
1 impl RecvGate {
2     fn ack_msg(&self, msg: &tcu::Message) -> Result<(), Error> {}
3 }
```



... or automatically when using GateIStream



Unmarshalling: GateIStream

```
1 let mut gate_stream = recv_msg(&recv_gate).expect("receive  
   failed");  
2 let val = gate_stream.pop::3 let text = gate_stream.pop::
```



Unmarshalling: GateIStream

```
1 let mut gate_stream = recv_msg(&recv_gate).expect("receive
  failed");
2 let val = gate_stream.pop::().expect("unable to get value");
3 let text = gate_stream.pop::().unwrap();
```



Marshalling: send macros

```
1 send_vmsg!(send_gate, &reply_gate, arg1, /*...*/);
2 reply_vmsg!(gate_stream, arg1, /*...*/);
3 let gate_stream = send_recv!(send_gate, &reply_gate, arg1, /*...*/);
```



Message Passing: Creation



- Programmatically:
 - `RecvGate::new(order, msg_order)`
 - `SendGate::new(&recv_gate)`
 - Receiver needs to delegate `SendGate` to sender

Message Passing: Creation



- Programmatically:
 - `RecvGate::new(order, msg_order)`
 - `SendGate::new(&recv_gate)`
 - Receiver needs to delegate `SendGate` to sender
- Configuration-based:
 - `SendGate::new_named(&name)`
 - `RecvGate::new_named(&name)`
 - Names and parameters defined in boot script

Message Passing: Creation



- Programmatically:
 - `RecvGate::new(order, msg_order)`
 - `SendGate::new(&recv_gate)`
 - Receiver needs to delegate `SendGate` to sender
- Configuration-based:
 - `SendGate::new_named(&name)`
 - `RecvGate::new_named(&name)`
 - Names and parameters defined in boot script

!! The sender also needs a receive gate for replies. !!
Create this receive gate programmatically!

Boot Script Example



```
1  <!-- ... -->
2  <app args="root">
3    <dom>
4      <app args="receiver">
5        <rgate name="chan" msgsize="64" slots="2"/>
6      </app>
7    </dom>
8    <dom>
9      <app args="sender">
10       <sgate name="chan" credits="2"/>
11     </app>
12   </dom>
13 </app>
14 <!-- ... -->
```

XML

Boot Script Example




XML

```
1  <!-- ... -->
2  <app args="root">
3    <dom>
4      <app args="receiver">
5        <rgate name="chan" msgsize="64" slots="2"/>
6      </app>
7    </dom>
8    <dom>
9      <app args="sender">
10       <sgate name="chan" credits="2"/>
11     </app>
12   </dom>
13 </app>
14 <!-- ... -->
```

chan

chan



Boot Script Example



XML

```
1  <!-- ... -->
2  <app args="root">
3    <dom>
4      <app args="receiver">
5        <rgate name="chan" msgsize="64" slots="2"/>
6      </app>                               RecvGate::new_named("chan")
7    </dom>
8    <dom>
9      <app args="sender">                   SendGate::new_named("chan")
10       <sgate name="chan" credits="2"/>
11     </app>
12   </dom>
13 </app>
14 <!-- ... -->
```

Wait! Debugging



- Logging:
 - Pass flags via `M3_LOG=...` when starting the system
 - Example: `M3_LOG=Info,Error,KernSysc,LibServ`
 - Hardcoded for `M3_BUILD=bench` to `Info,Error`
 - All flags are defined in `src/libs/rust/base/src/io/logflags.rs`

More Debugging



- Gem5 flags:
 - Can be specified via `M3_GEM5_LOG=...`
 - Tcu: show send, recv, ...
 - TcuCmd: show TCU command executions
 - ...



- Gem5 flags:
 - Can be specified via `M3_GEM5_LOG=...`
 - `Tcu`: show send, recv, ...
 - `TcuCmd`: show TCU command executions
 - ...
- Instruction trace:
 - Set `M3_GEM5_LOG=Exec`
 - Run your scenario in `gem5`
 - `grep C0T04 run/gem5.log | ./b trace tilemux,<program>`
shows you a beautified version of the instructions executed on `C0T04`

Assignment #2: IPC



- Build a program that sends items
- Build a program that receives items
- Use the credit system
- Simple version: communicate synchronously
- Fill in implementation at `src/apps/ipc` and `boot/ipc.xml`

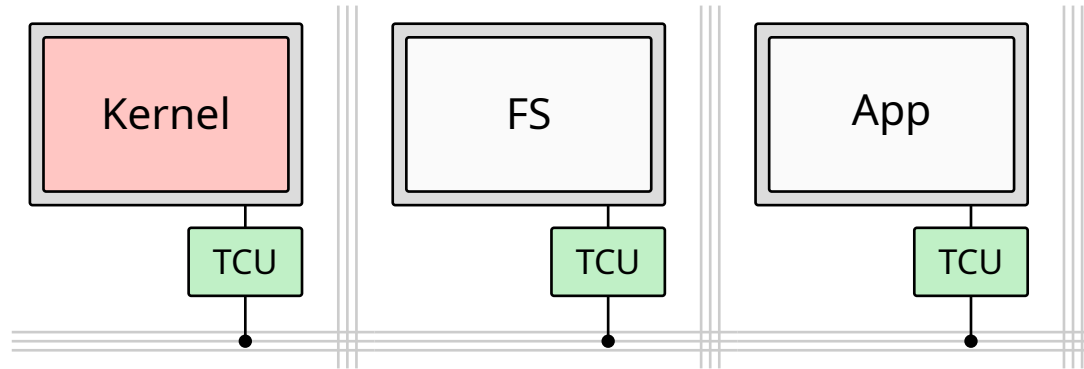
Assignment #2: IPC



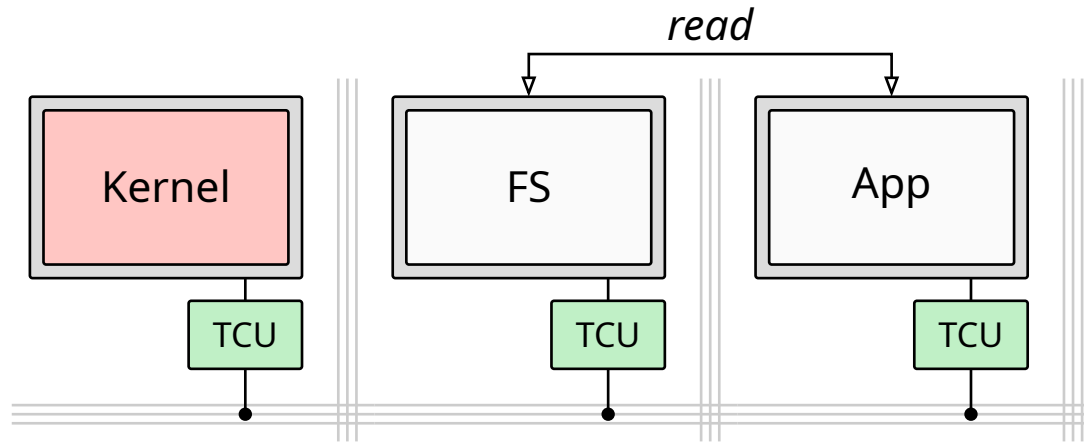
- Build a program that sends items
- Build a program that receives items
- Use the credit system
- Simple version: communicate synchronously
- Fill in implementation at `src/apps/ipc` and `boot/ipc.xml`
- Advanced version: communicate asynchronously
- Let the sender keep multiple messages in flight
- You must wait for a reply when credits/receive buffers depleted

- Introduction
- **Assignments**
 - Hello World
 - IPC
 - **Filesystem Service**
 - Accelerators
 - Sessions

Accessing Files



Accessing Files



Services



- A *server* offers a *service* for clients
- A **connection between client and server** is represented as a *session*
- Who can offer and use which service is defined via boot script

Boot Script Example (simplified)



```
1 <mods><mod name="fs" file="default.img"/></mods>
2 <!-- ... -->
3 <app args="m3fs mem" daemon="1">
4   <serv name="m3fs"/>
5   <mod name="fs"/>
6 </app>
7 <!-- ... -->
8 <app args="pager">
9   <sess name="m3fs"/>
10  <mod name="fs" perm="r"/>
11  <app args="/bin/myapp">
12    <mount fs="m3fs" path="/" />
13  </app>
14 </app>
```

XML



Virtual File System

```
1 mod VFS {  
2     // Open `path` with given permissions  
3     fn open(path: &str, flags: OpenFlags)  
4         -> Result<FileRef<dyn File>, Error> {}  
5     // Create directory at `path` with given mode  
6     fn mkdir(path: &str, mode: FileMode) -> Result<(), Error> {}  
7 }
```



FileRef and File

```
1 // Read into `buffer` (supports non-blocking mode)
2 fn read(&mut self, buf: &mut [u8]) -> Result<usize, Error> {}
3 // Write from `buffer` (supports non-blocking mode)
4 fn write(&mut self, buf: &[u8]) -> Result<usize, Error> {}
```





Analogous to Rust's Standard Library

```
1 let first_arg: &str = m3::env::args().skip(1).next().unwrap();
```



Look at the generated documentation for type information.

Assignment #3: Read a file



- **Write a program** that reads the contents of a given file
- Print the contents to stdout (`m3::println`)
- Modify `src/apps/rdfile`



Assignment #3: Read a file

- **Write a program** that reads the contents of a given file
- Print the contents to stdout (`m3::println`)
- Modify `src/apps/rdfile`

- You can **test the program** in the shell via `boot/shell.xml`
- All programs are already on the "PATH".
- Use `ls` in the shell to find a suitable file
- You can exit the shell via `Ctrl+D` (takes a while)
- Or write a boot script with a hard-coded path



Assignment #3: Read a file

- **Write a program** that reads the contents of a given file
- Print the contents to stdout (`m3::println`)
- Modify `src/apps/rdfile`

- You can **test the program** in the shell via `boot/shell.xml`
- All programs are already on the "PATH".
- Use `ls` in the shell to find a suitable file
- You can exit the shell via `Ctrl+D` (takes a while)
- Or write a boot script with a hard-coded path

- **Advanced version:** use the Rust standard library
- See `src/apps/ruststdtest` for an example

- Introduction
- **Assignments**
 - Hello World
 - IPC
 - Filesystem Service
 - **Accelerators**
 - Sessions

Accelerator Tile

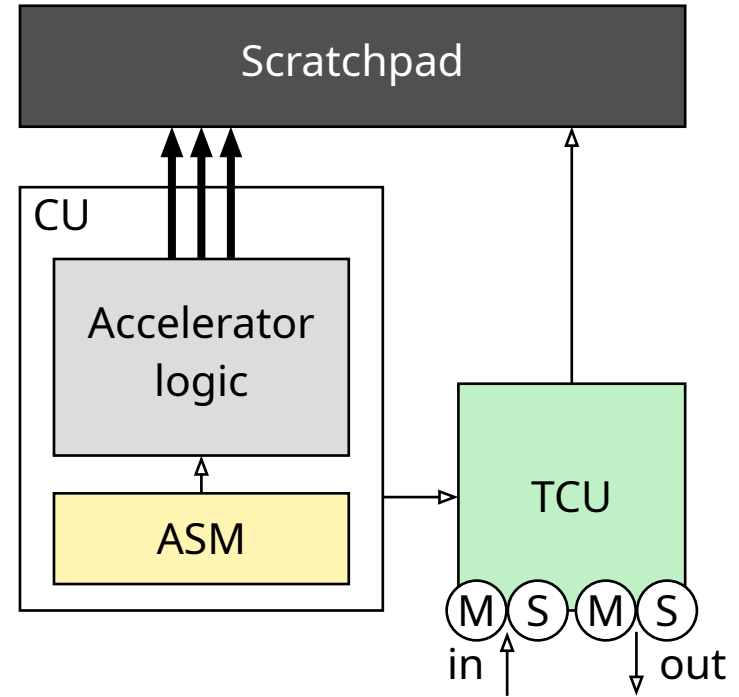


- Accelerators can sit on their own tiles
- Activities on other tiles **stream data** in/out
- Communication via **file protocol**
- We use the **pipes service** to offload buffer & notification management

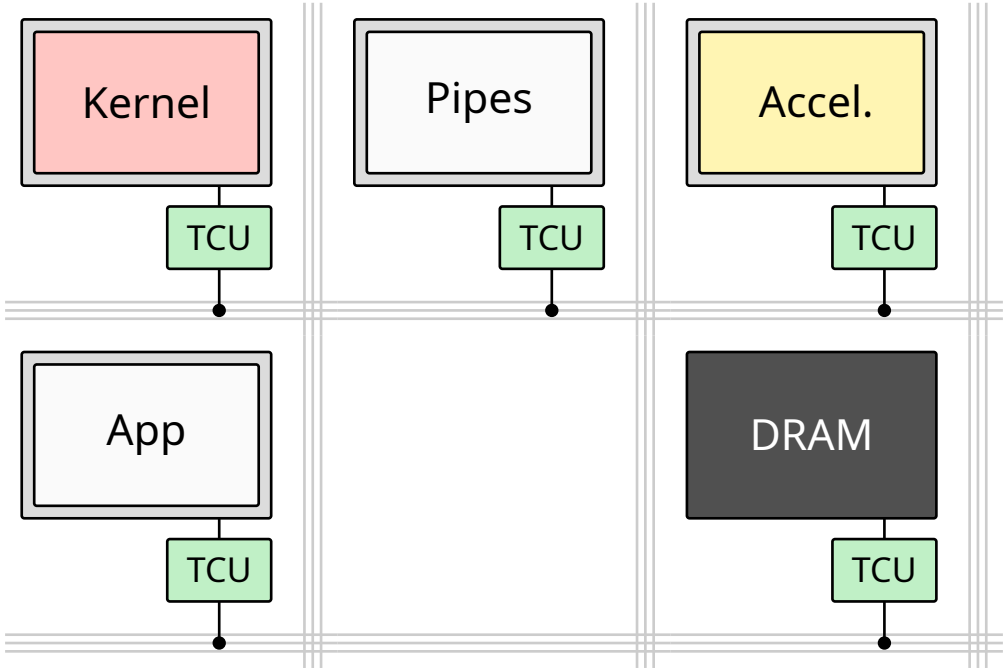
Accelerator Tile



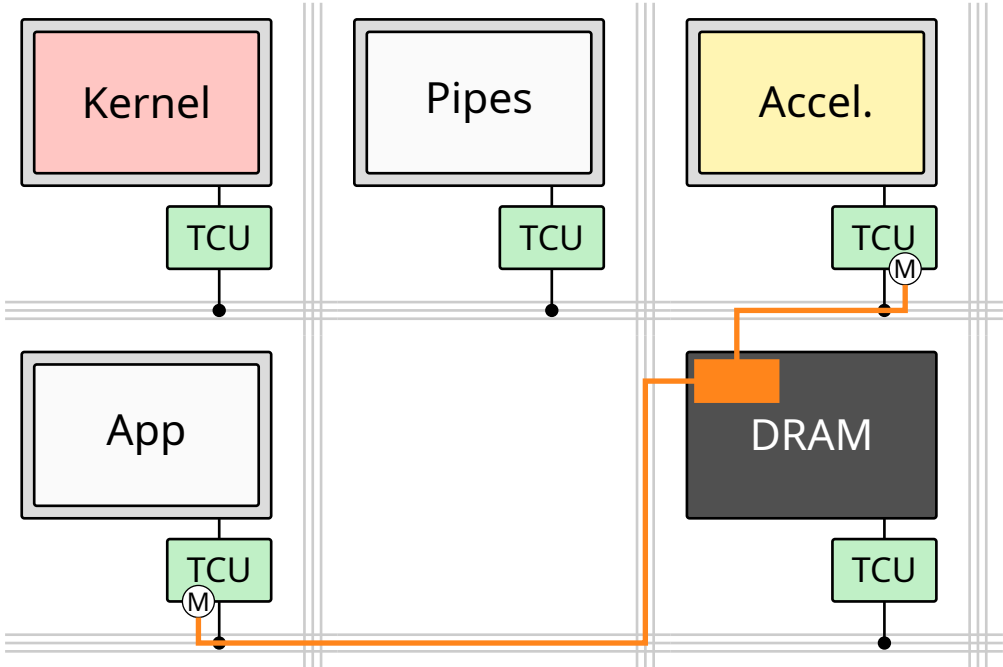
- Accelerators can sit on their own tiles
- Activities on other tiles **stream data** in/out
- Communication via **file protocol**
- We use the **pipes service** to offload buffer & notification management



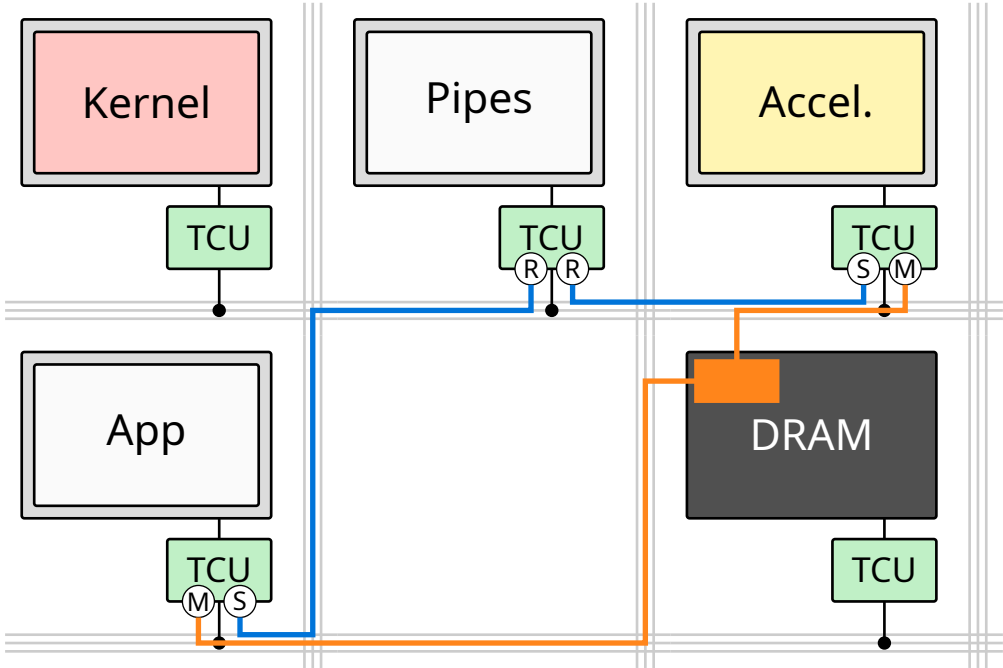
System Setup With Accelerators



System Setup With Accelerators



System Setup With Accelerators



Boot Script Example (Simplified)



```
1  <!-- ... -->
2  <app args="pipes" daemon="1">
3    <serv name="pipes"/>
4  </app>
5  <!-- ... -->
6  <app args="app">
7    <sess name="pipes"/>
8    <!-- Give the app control over the accelerator tile: -->
9    <tiles type="accel"/>
10 </app>
11 <!-- ... -->
```

XML



Abstractions

Pipes, IndirectPipe, Tile, ChildActivity, StreamAccel

And of course the file abstractions you already know.



Abstractions

Pipes, IndirectPipe, Tile, ChildActivity, StreamAccel

And of course the file abstractions you already know.

Rust Slices

```
1 let input_slice = &input[from_this_index..];  
2 let output_slice = &mut output[from_this_index..];
```



Assignment #4: Use an accelerator



- Write a program that uses the *rot13* accelerator
- Use the program arguments as input
- Print the accelerator output to stdout
- Modify `boot/acceldemo.xml` and `src/apps/acceldemo`
- Use M³'s Rust documentation for learning how to use the abstractions

- Introduction
- **Assignments**
 - Hello World
 - IPC
 - Filesystem Service
 - Accelerators
 - **Sessions**



L4-style microkernels:
Delegate access rights to memory, files, etc. via IPC



L4-style microkernels:

Delegate access rights to memory, files, etc. via IPC

M³ 's IPC bypasses the kernel ...

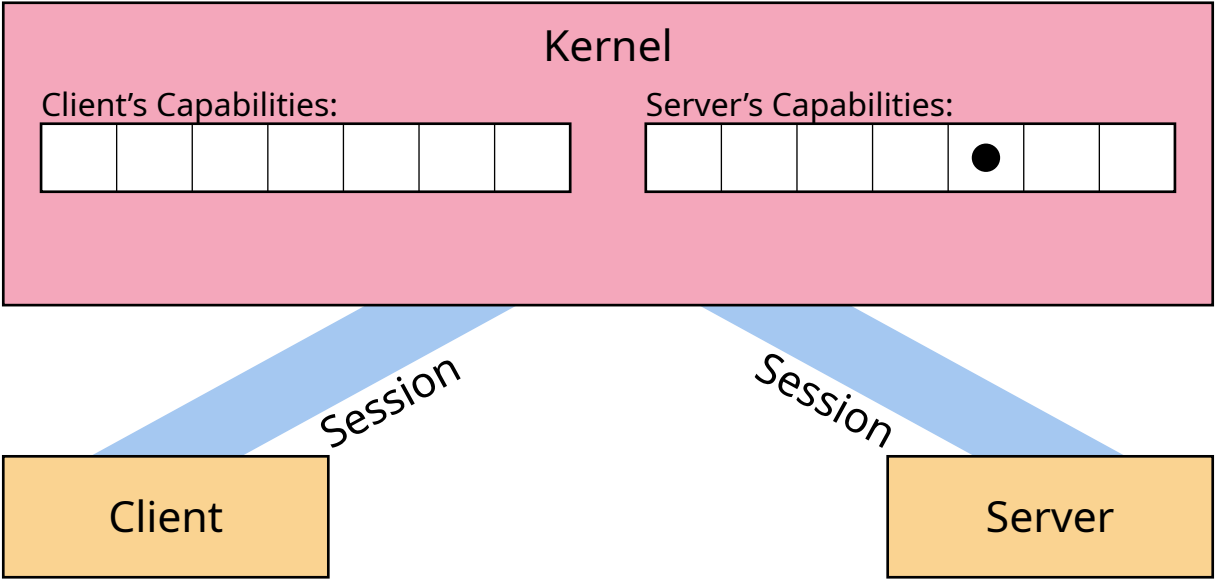


L4-style microkernels:

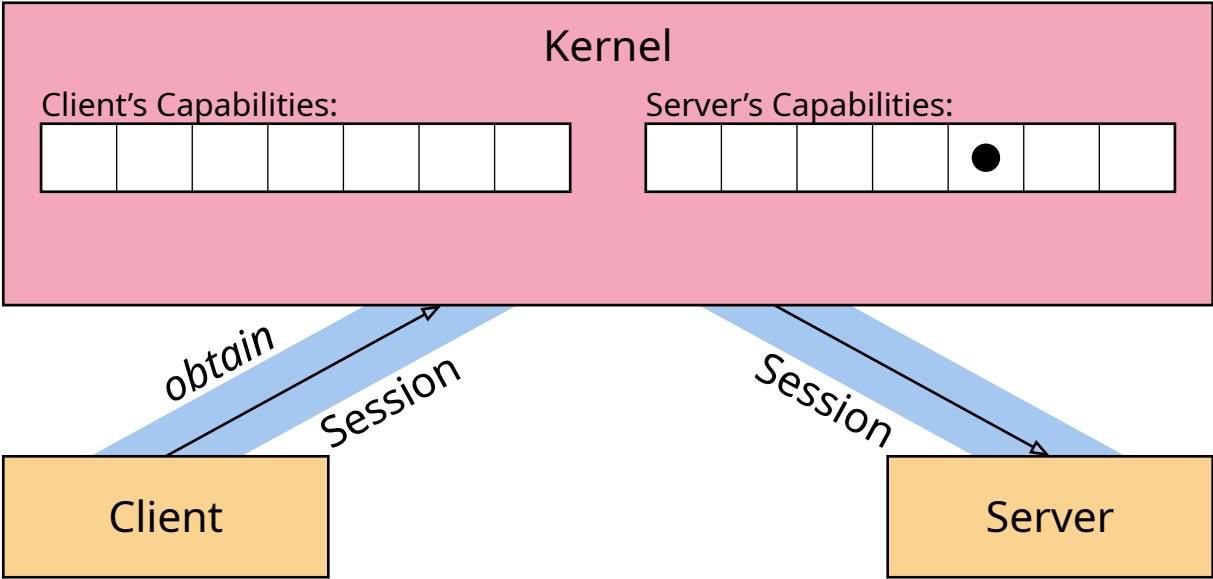
Delegate access rights to memory, files, etc. via IPC

M³ 's IPC bypasses the kernel ...
How to exchange capabilities?

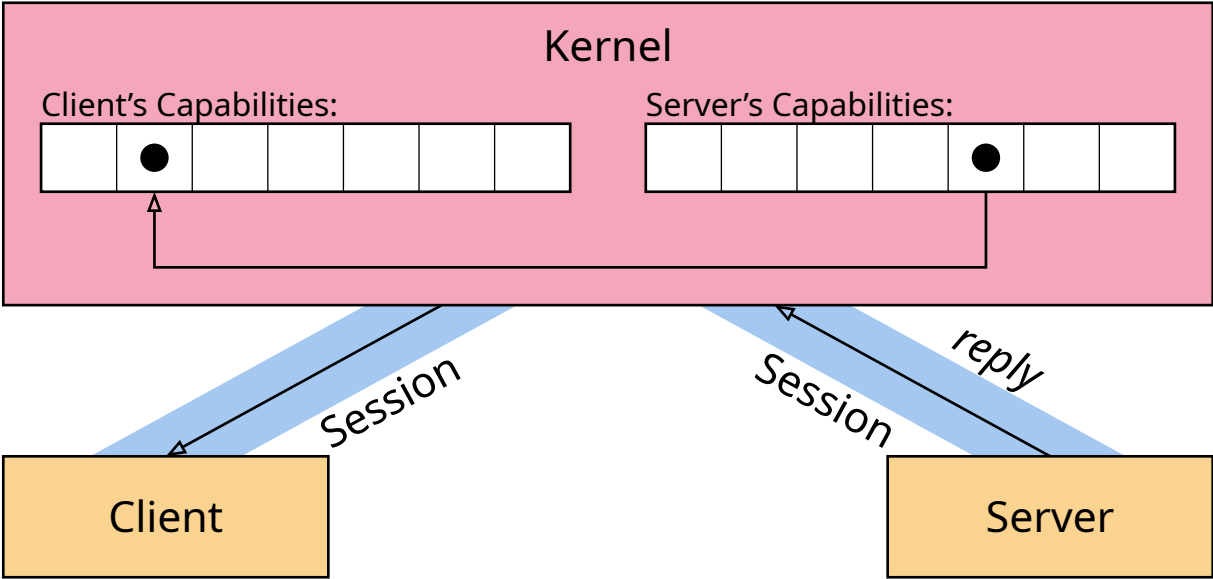
Servers and Sessions via the Kernel



Servers and Sessions via the Kernel



Servers and Sessions via the Kernel



Client Side of Obtain



```
1 let crd = session.obtain(  
2     1, // Number of capabilities to obtain  
3     |os| { // Prepare arguments for the server  
4         // First argument is the opcode:  
5         os.push(42usize);  
6         os.push("my argument");  
7     },  
8     |is| { // Handle reply from the server  
9         let some_reply_arg = is.pop::<bool>();  
10        Ok(())  
11    },  
12 )
```



Server Side of Obtain (Simplified)



```
1 // Register the handler my_handler for opcode 42:  
2 hdl.reg_cap_handler(42usize, ExcType::Obt(1), Session::my_handler);
```



Server Side of Obtain (Simplified)



```
1 // Register the handler my_handler for opcode 42: 🔒 Rust  
2 hdl.reg_cap_handler(42usize, ExcType::Obt(1), Session::my_handler);
```

```
1 impl Session { 🔒 Rust  
2     fn my_handler(xchg: &mut CapExchange<'_>) {  
3         /* ... */  
4         // Instruct the kernel to transfer the capability my_cap:  
5         xchg.out_caps(CapRngDesc::new_single(  
6             CapType::Object,  
7             my_cap.sel(),  
8         ));  
9     }  
10 }
```

Assignment #5: Servers and sessions



- Server allocates memory capability containing message
- The server provides access to this memory **capability via sessions**
- The **client obtains** copy of capability via a session
- The client prints the message stored behind the memory capability
- Modify `src/apps/sessions`
- **Fill in** the handler registration, capability exchange, and obtain logic

Conclusion

- M³ is available as **open source** (both hardware and software):
<https://github.com/Barkhausen-Institut/M3>



- README.md explains how to set it up

Conclusion



- M³ is available as **open source** (both hardware and software):
<https://github.com/Barkhausen-Institut/M3>



- README.md explains how to set it up
- **If you think M³ is interesting, contact us for thesis topics!**
nils.asmussen or viktor.reusch@barkhauseninstitut.org