

Embedded Linux Conference 2021

Understanding the structure of a Linux kernel device driver



Embedded Labworks



\$ WHOAMI

- × Embedded software developer for more than 25 years.
- × Consultant and trainer at Embedded Labworks.
<https://e-labworks.com/en>
- × Open source contributor (Buildroot, Yocto Project, Linux kernel, etc).
- × Blogger at <https://embeddedbits.org/>.





ABOUT THE TALK

- x The objective of this talk is to cover the modern architecture of a Linux kernel device driver.
- x This is an entry level talk about the main pieces of a device driver in the Linux kernel (let's connect the dots!).
- x We will not cover Linux driver model's internal structures and data representation (kobjects, ktypes, ksets, etc).
- x We will also not cover Linux API's in details.
- x In the end of this talk, you will be able to read and understand a modern Linux device driver, and why not, start writing your own!





AGENDA

1. Introduction to device drivers and character drivers
2. Hardware access (MMIO, gpiolib)
3. Introduction to the driver model
4. Frameworks
5. Buses and device tree
6. Lots of hands-on!





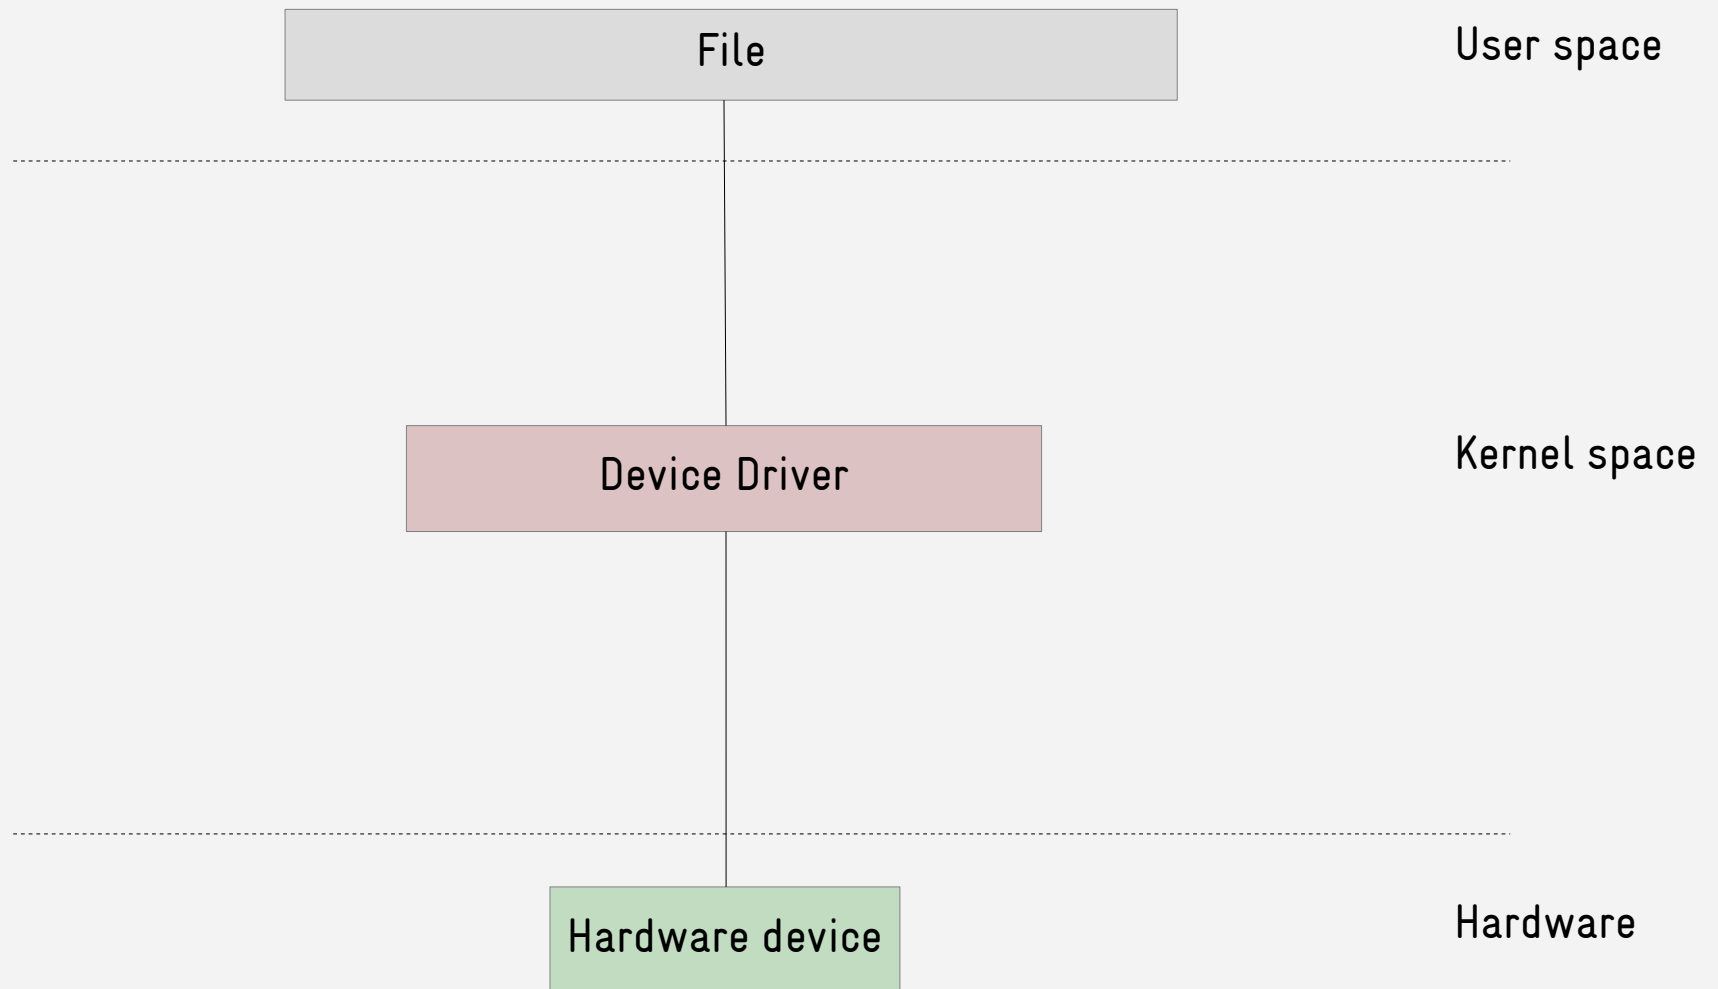
WHAT ARE DEVICE DRIVERS?

- x Device drivers are just abstractions to a piece of hardware!
- x Although you can have device drivers running in userspace (via some kernel interface like UIO or I2CDEV), it's much more common to have them running in kernel space.
- x Providing an infrastructure to write and run device drivers is part of an operating system's kernel responsibility.
- x On Unix based systems, a file is the most common abstraction to communicate with a device.





DEVICE DRIVER IS AN ABSTRACTION





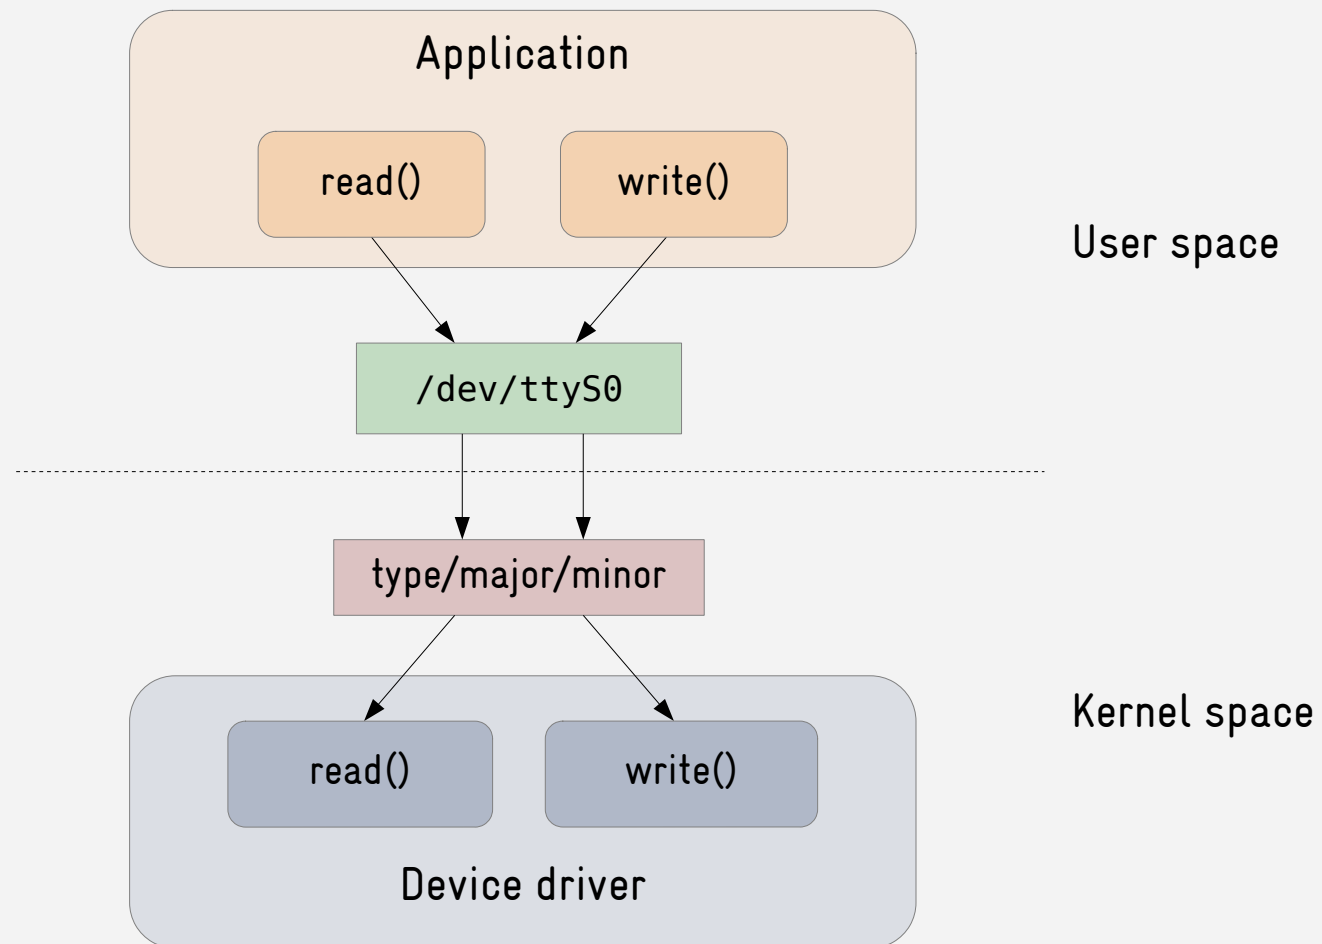
CHAR DRIVER: A SIMPLE ABSTRACTION

- x There are a few kernel APIs that could be used to export an interface to userspace in the filesystem (`/dev`, `/sys`, etc).
- x Device nodes are one of the interfaces that could be used, where files are exported to users in `/dev` in the form of **char** or **block** device files.
- x These device files have associated three basic information:
 - x Type (block or char).
 - x Major number.
 - x Minor number.





CHAR DRIVER AS A FILE ABSTRACTION





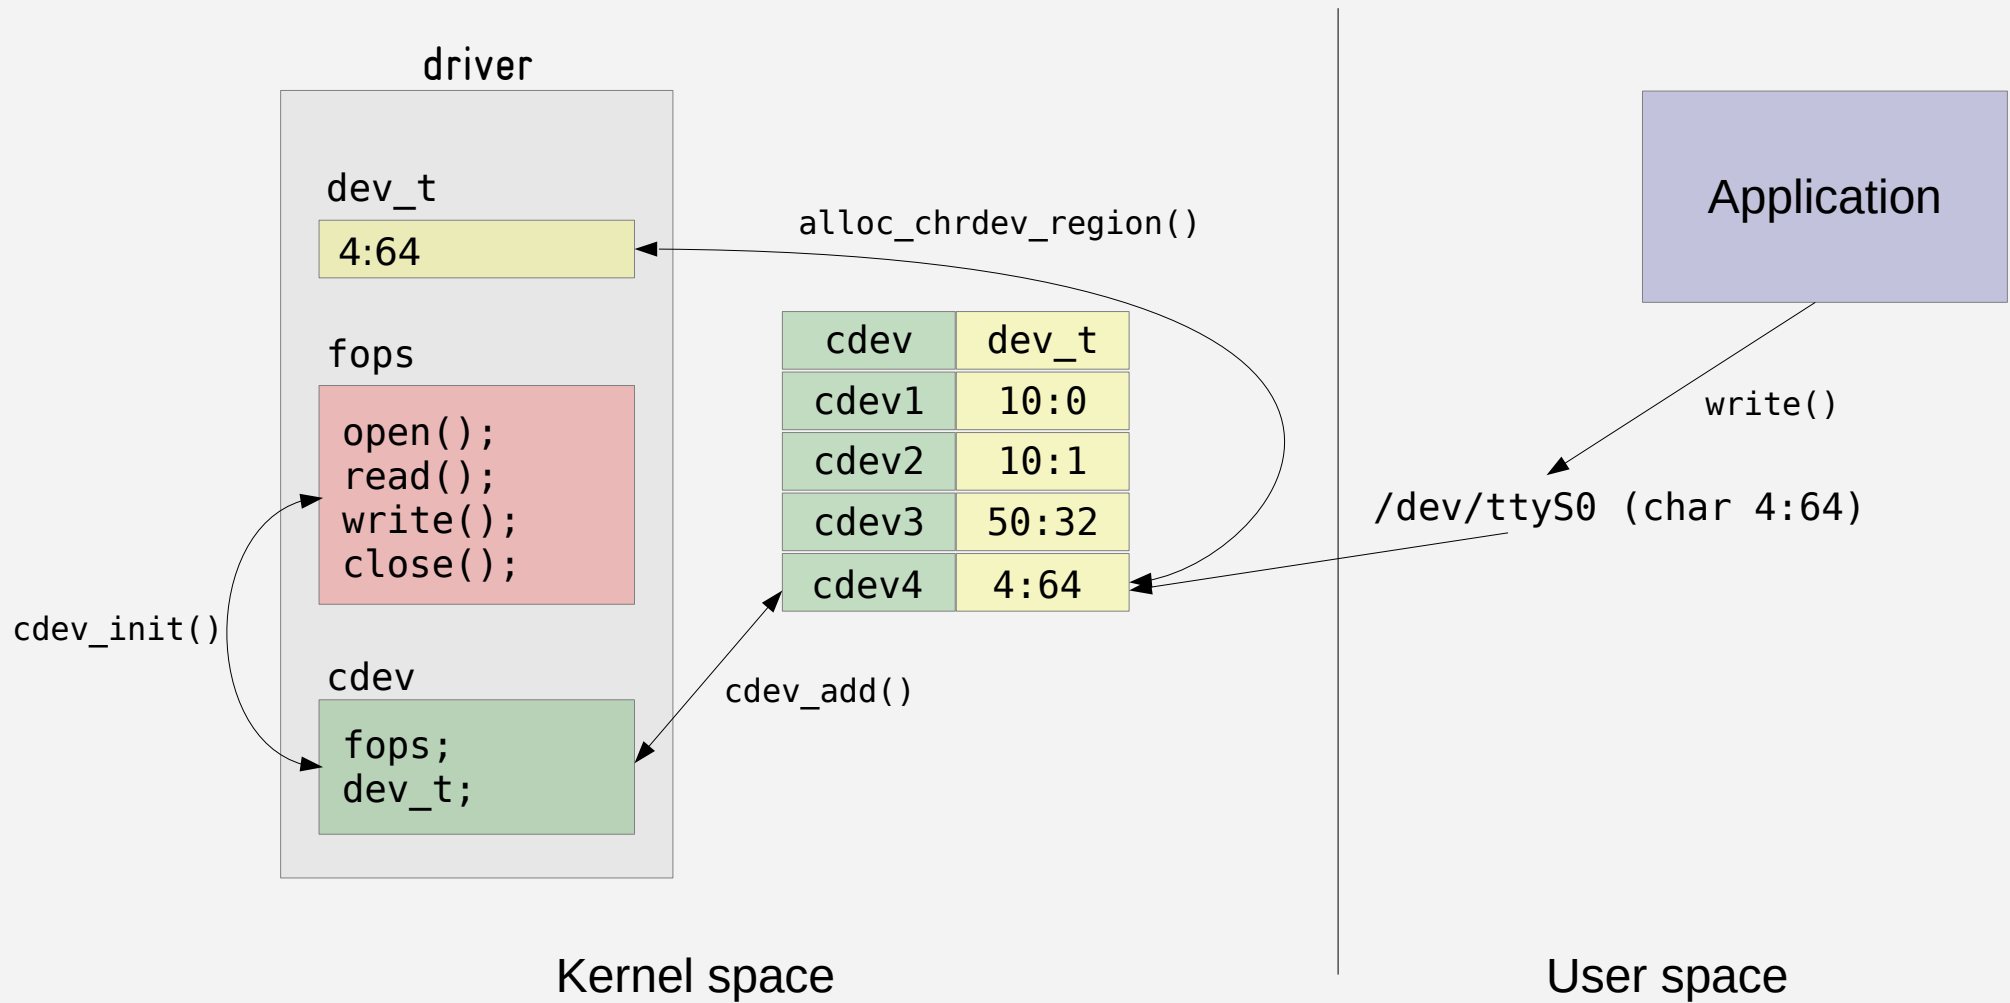
IMPLEMENTING A CHAR DRIVER

- x **Step 1:** Allocate the device number (major/minor). This can be done with `register_chrdev_region()` or `alloc_chrdev_region()`.
- x **Step 2:** Implement the file operation (open, read, write, ioctl, etc).
- x **Step 3:** Register the char driver in the kernel with `cdev_init()` and `cdev_add()`.





CHAR DRIVER



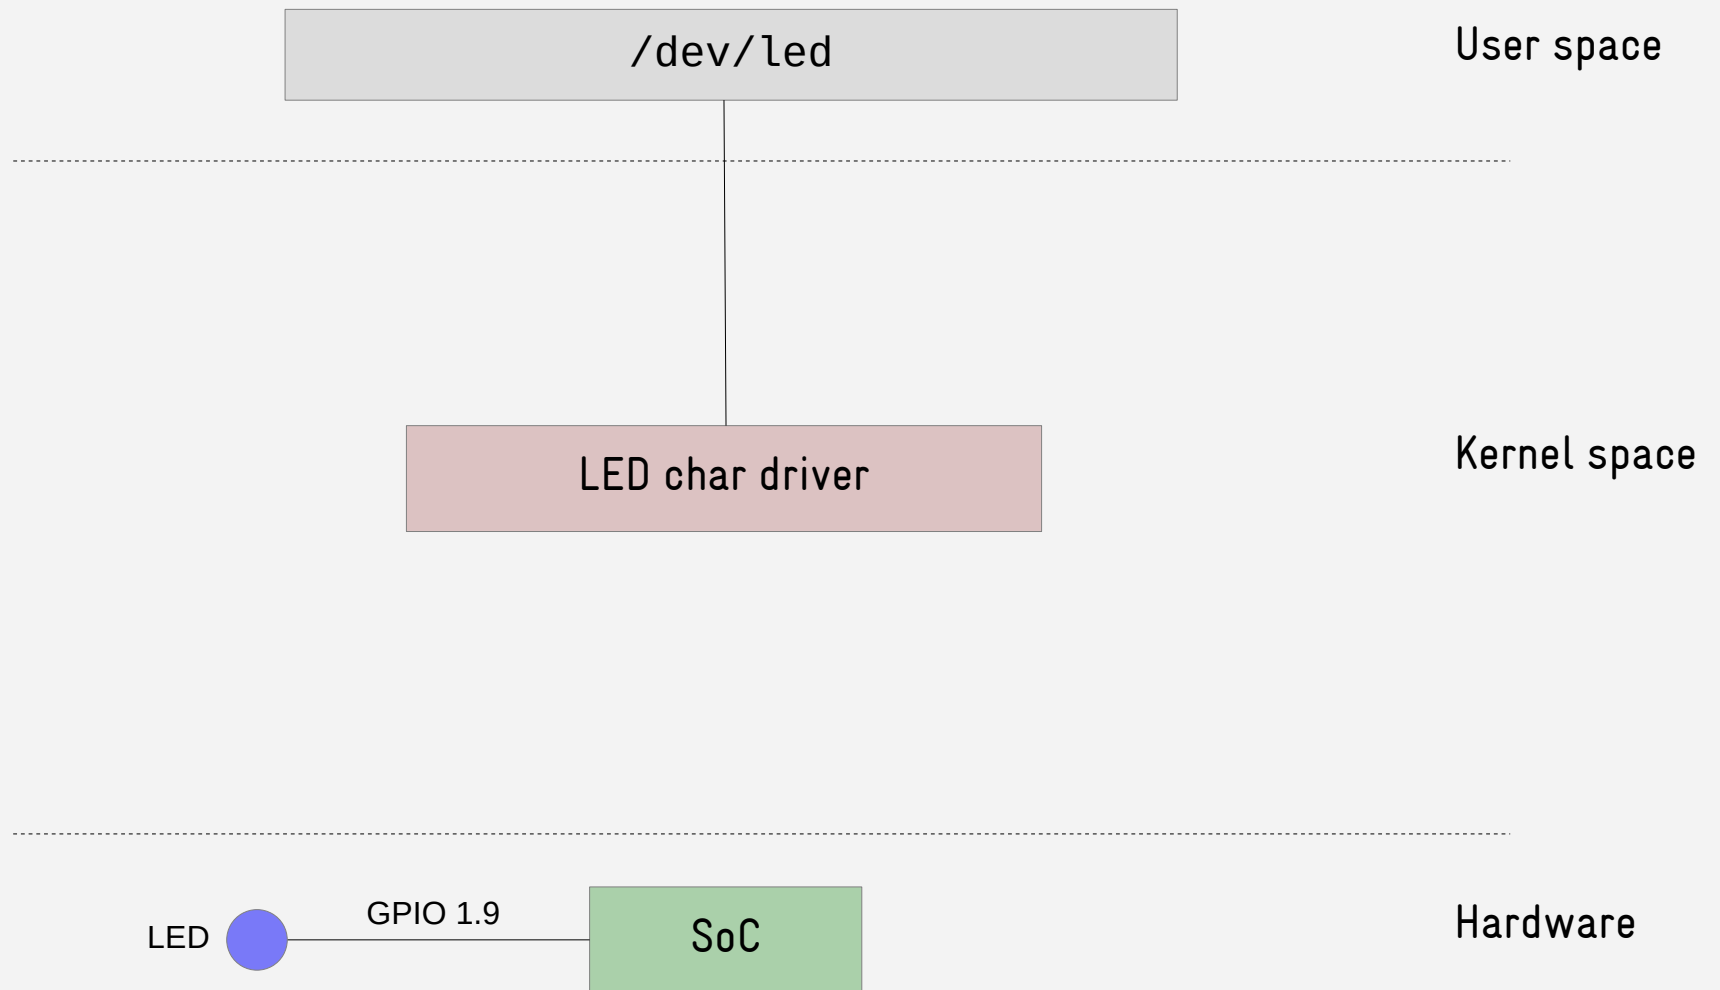


Hands-on 1

Writing a char driver



LED DRIVER





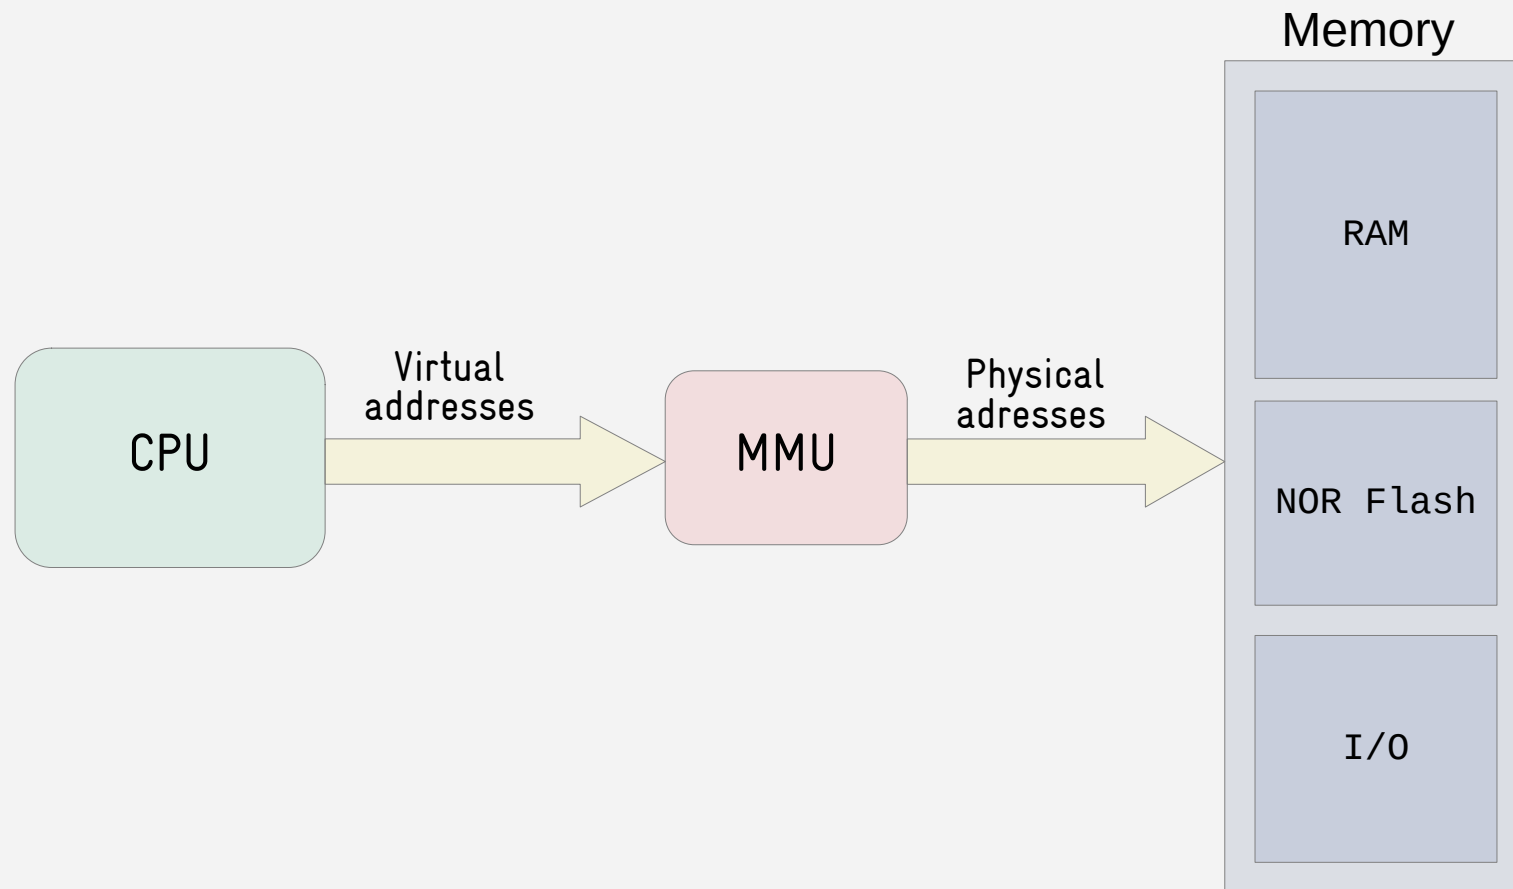
TALKING TO THE HARDWARE

- x Depending on the hardware architecture, there are a few mechanisms a CPU can use to communicate with a hardware device, including:
 - x **Port I/O:** a dedicated bus is used to communicate with hardware devices.
 - x **Memory-mapped I/O:** the memory address space is shared with hardware devices.
- x MMIO is currently the most common approach adopted by popular architectures like ARM.





MEMORY-MAPPED I/O





TALKING TO A MMIO DEVICE

- x **Step 1:** request access to MMIO registers using a few kernel APIs like `request_mem_region()`.
- x **Step 2:** Map the registers' physical addresses to virtual addresses, using functions like `ioremap()`.
- x **Step 3:** Use the kernel API to read from and write to the registers, with functions like `readl()` and `writel()`.



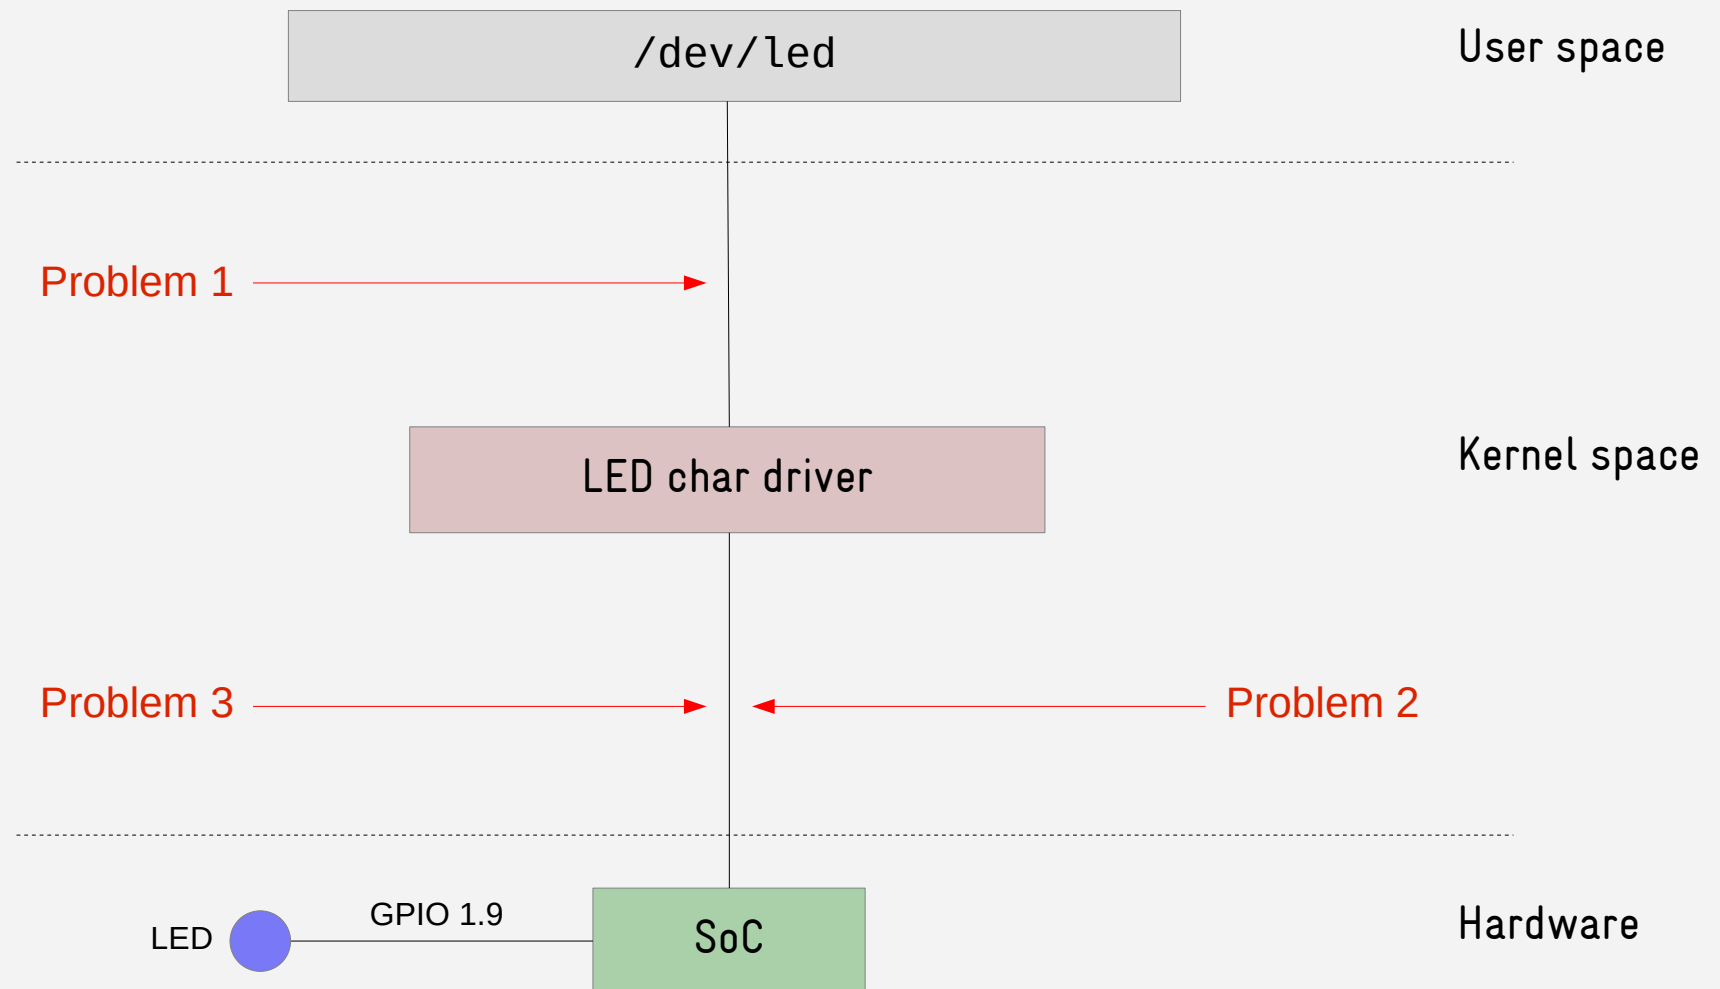


Hands-on 2

Talking to a MMIO based device



LED DRIVER





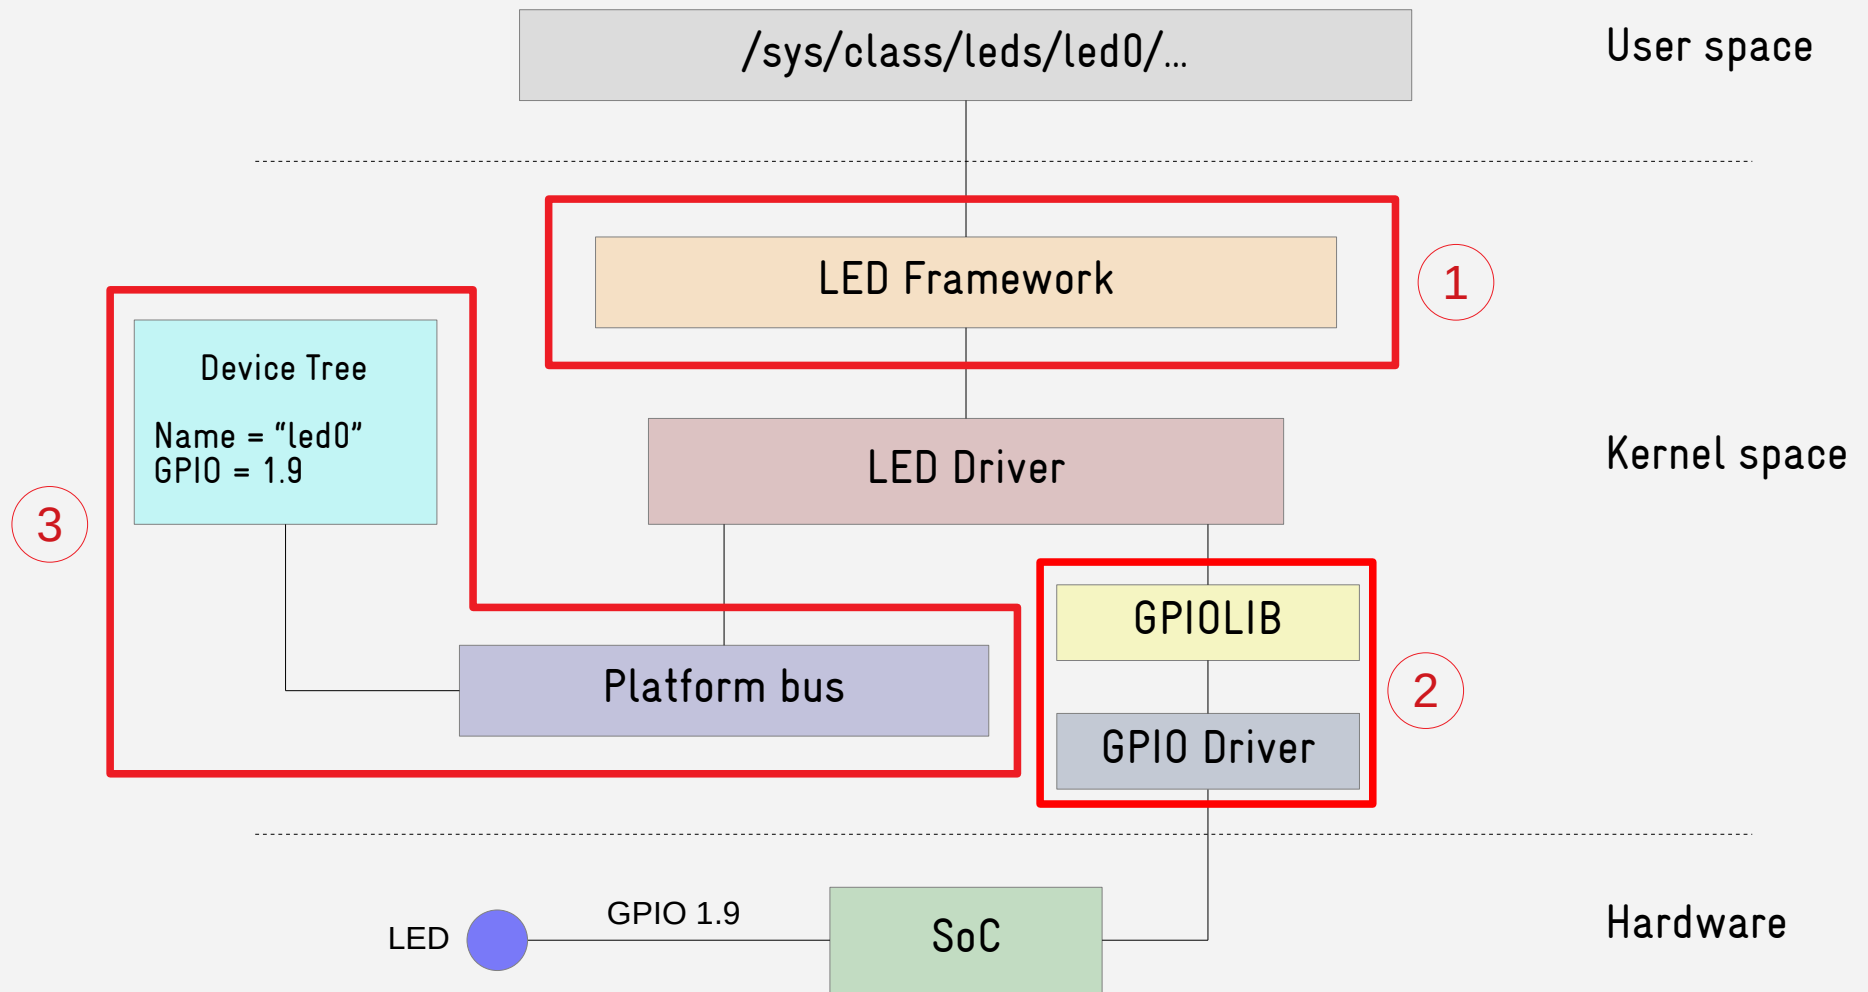
THE DRIVER MODEL

- x The Linux driver model provides several abstractions to device drivers to make the code more modular, reusable and easy to maintain.
- x Among its components, we have:
 - x **Frameworks:** the interface exported by a type or class of devices is standardized.
 - x **Buses:** information about the devices and where they are connected is abstracted away from the driver.





LED DRIVER





FRAMEWORKS

- x Frameworks make it easier to provide a standardized interface and abstraction for users.
 - x Device drivers developers don't need to think about the interface that will be exported to users.
 - x Users know beforehand the interface provided by a driver based on its class or type.
- x There are frameworks for the most common types of devices: input, IIO, ALSA, V2L2, RTC, watchdog, etc!
- x And of course, this is a framework for LED drivers!





USING THE LEDES FRAMEWORK

- x **Step 1:** Initialize an structure of type `led_classdev`.
- x **Step 2:** Provide (at least) a callback function to change the status of the LED.
- x **Step 3:** Register the driver in the framework with the function `led_classdev_register()`.



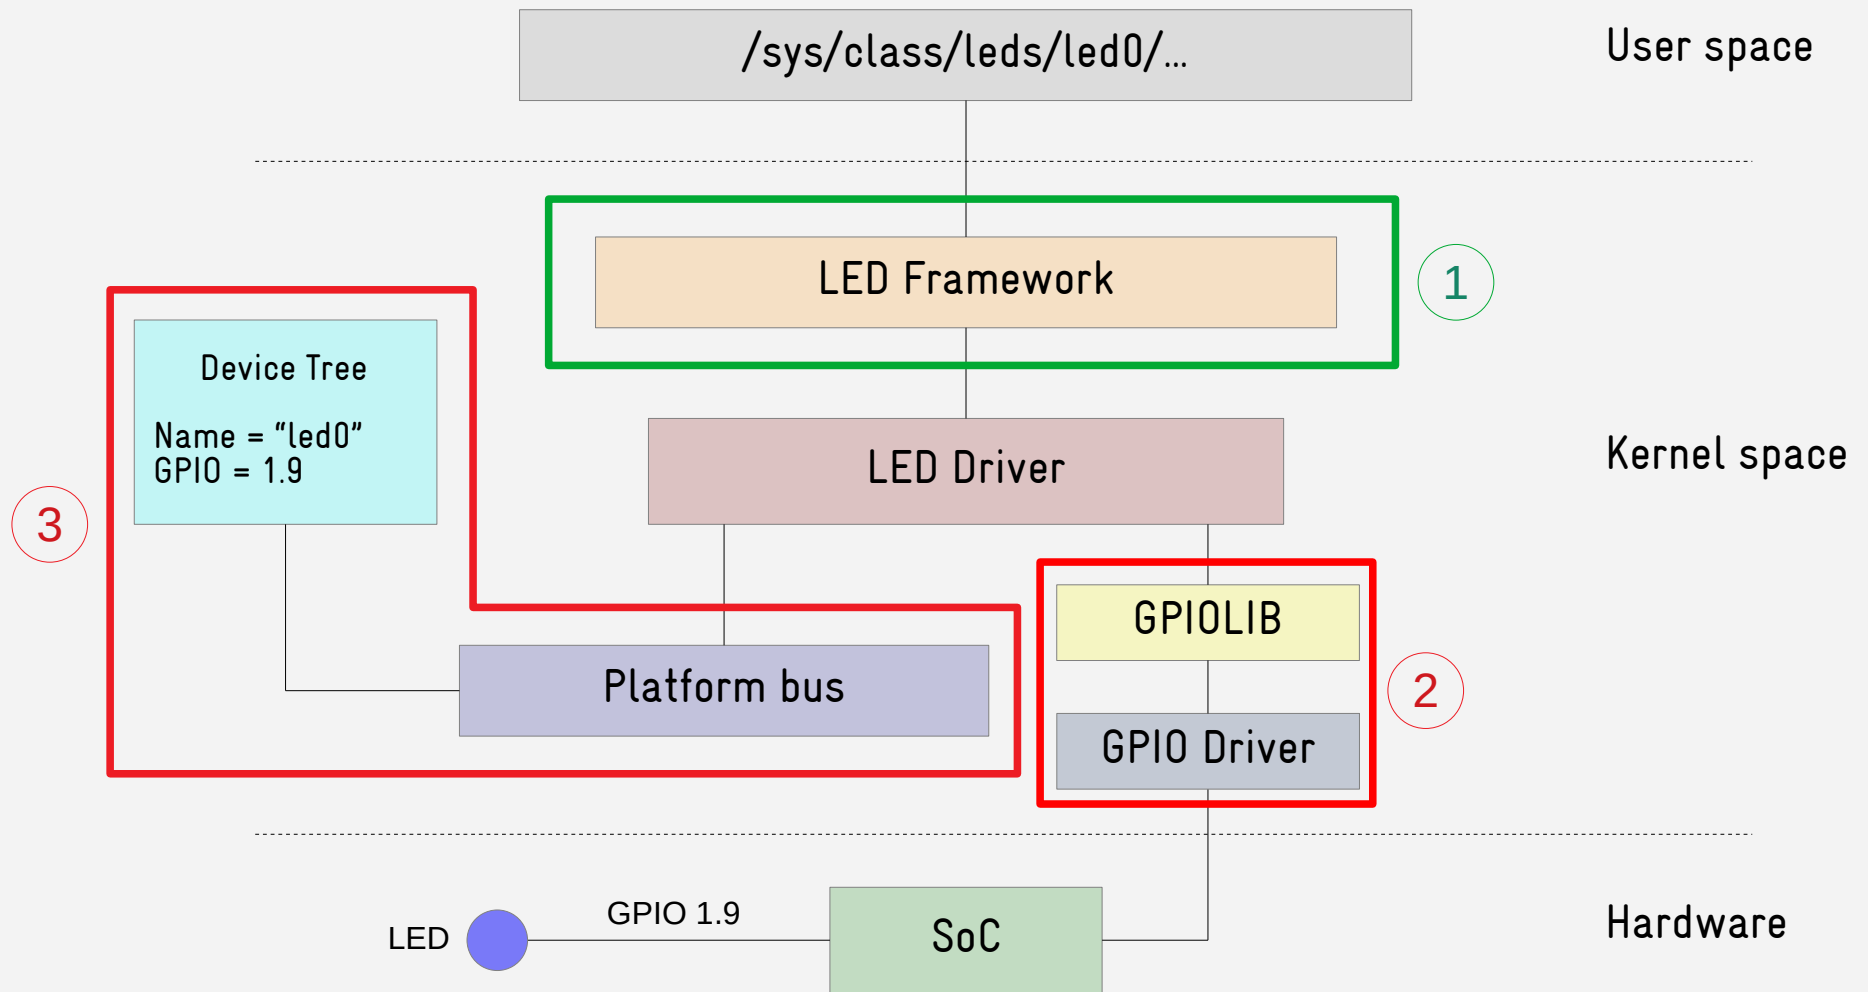


Hands-on 3

Using the LEDS framework



LED DRIVER





GPIOLIB

- x The Linux kernel implements a **producer/consumer** model for GPIO management:
 - x There are GPIO "producers" like GPIO controllers drivers.
 - x There are GPIO "consumers" like a LED driver, keyboard driver or a userspace application.
- x **gpiolib** is the API provided by the Linux kernel for GPIO producers and consumers.





GPIOLIB API

```
#include <linux/gpio.h>
#include <linux/gpio/consumer.h>

struct gpio_desc *gpiod_get(struct device *dev,
                           const char *con_id,
                           enum gpiod_flags flags);
void gpiod_put(struct gpio_desc *desc);

int gpiod_direction_input(struct gpio_desc *desc);
int gpiod_direction_output(struct gpio_desc *desc, int value);

void gpiod_set_value(struct gpio_desc *desc, int value);
int gpiod_get_value(const struct gpio_desc *desc);

struct gpio_desc *gpio_to_desc(unsigned gpio);
int desc_to_gpio(const struct gpio_desc *desc);

int gpiod_to_irq(const struct gpio_desc *desc);
```



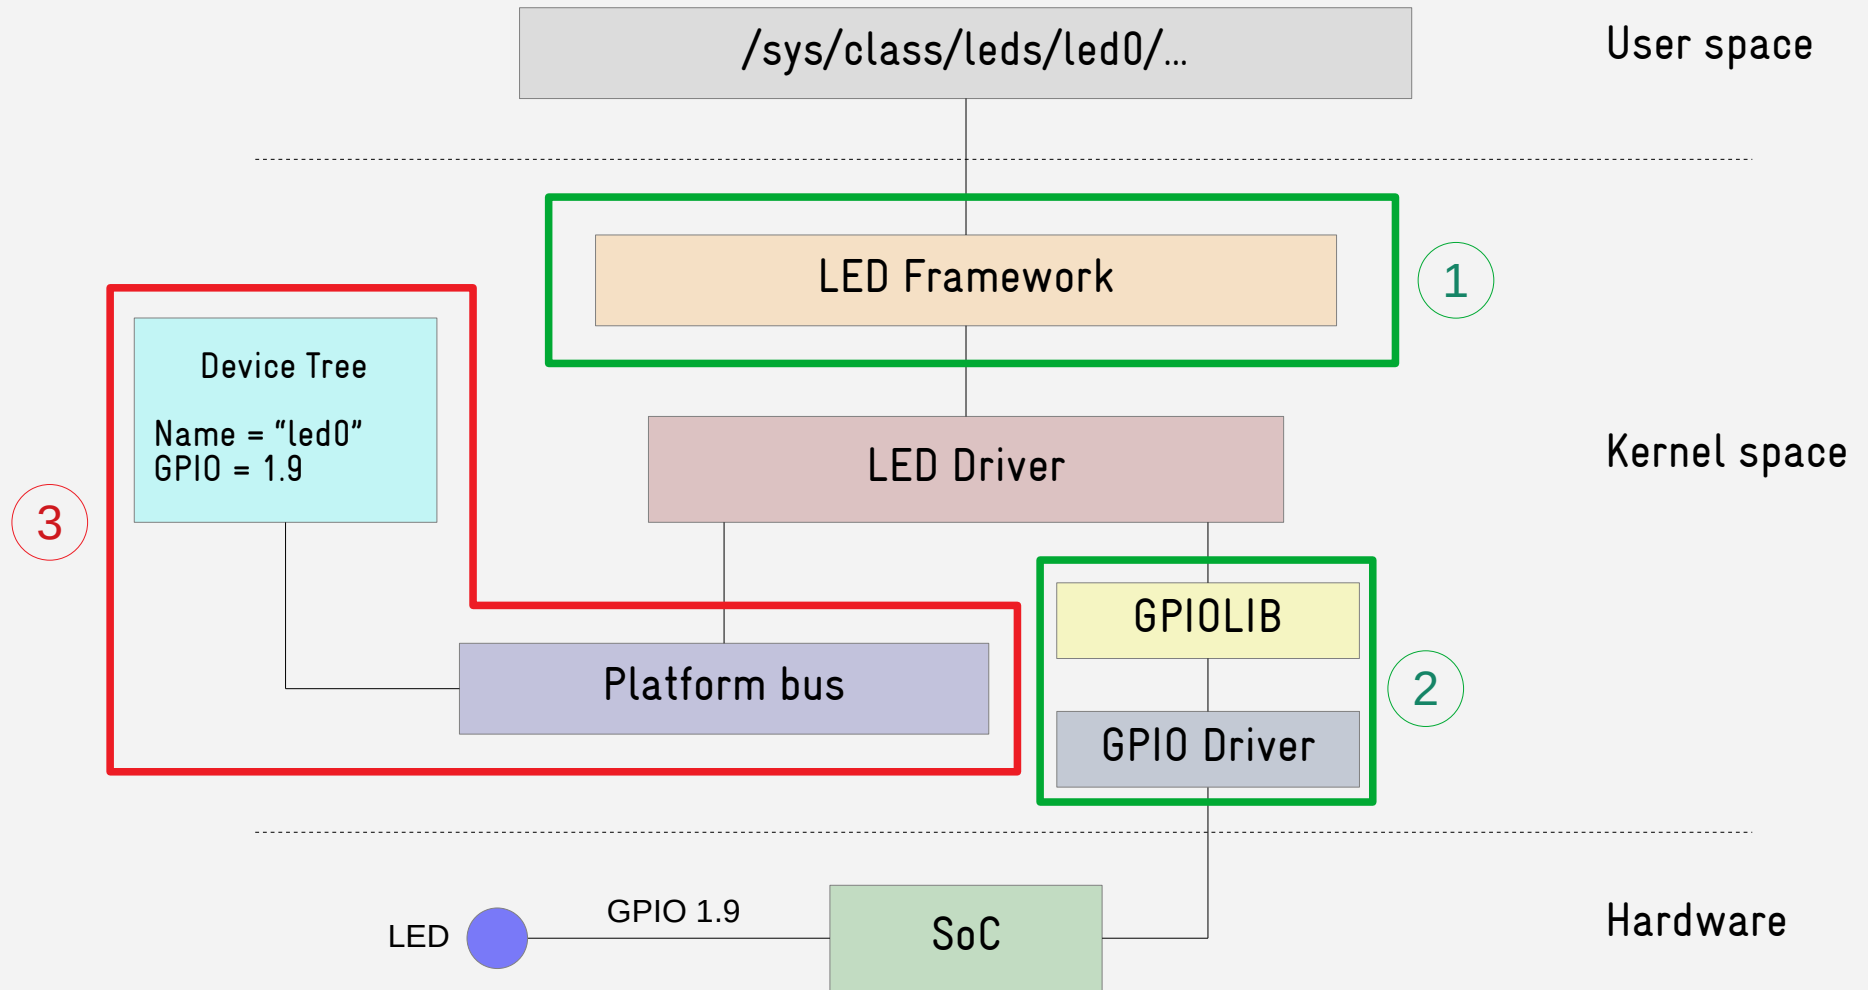


Hands-on 4

Accessing the hardware via GPIO LIB



LED DRIVER





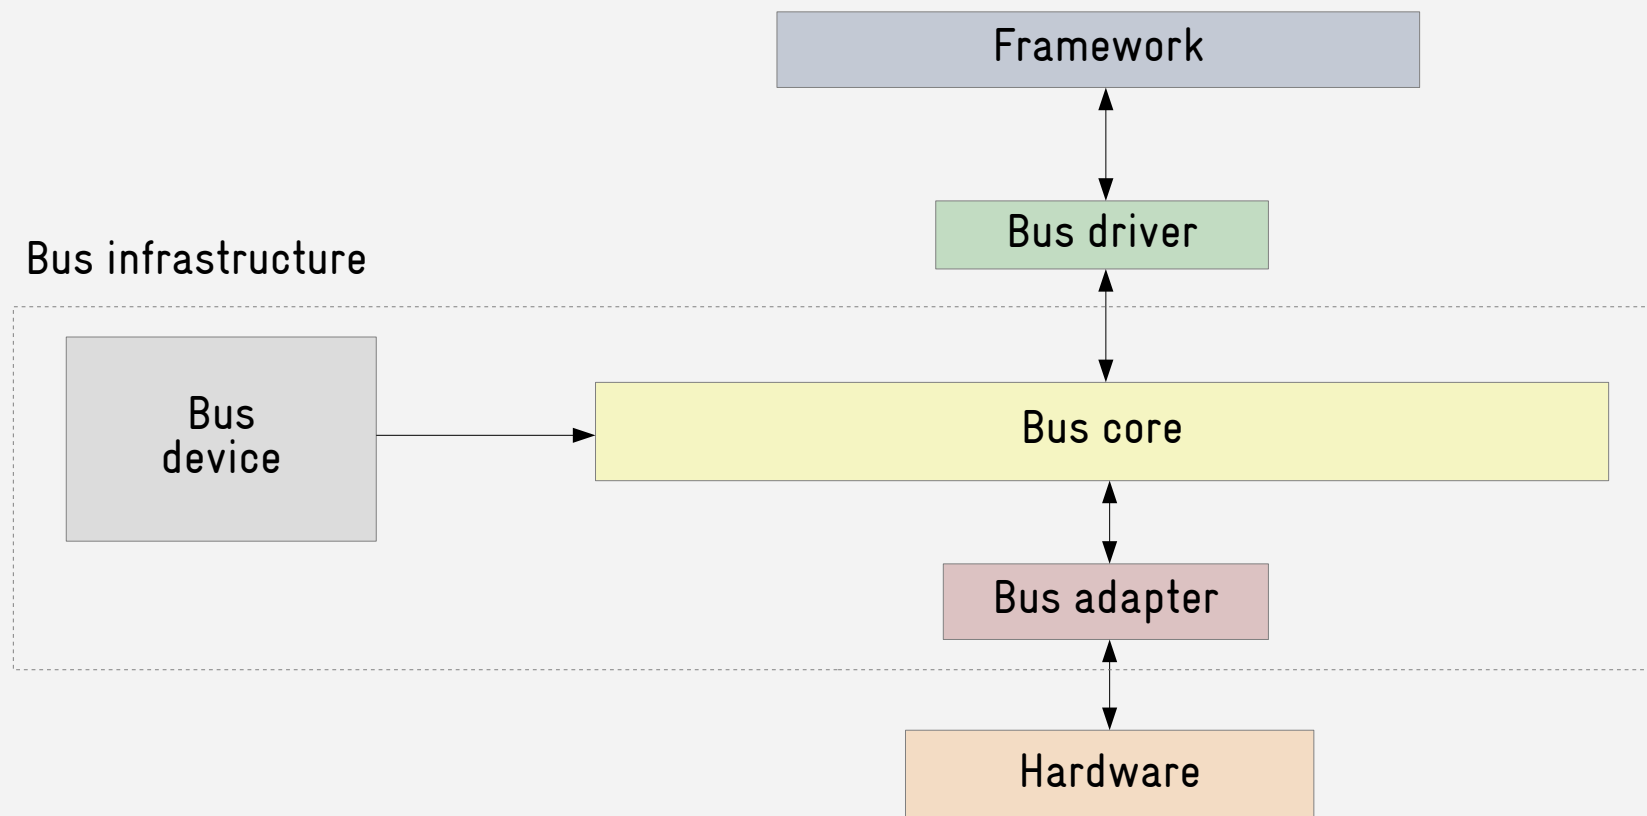
BUS INFRASTRUCTURE

- x **Bus core:** API implementation for a given bus (USB core, SPI core, I2C core, PCI core, etc), represented in the kernel by the `bus_type` structure.
- x **Bus adapters:** bus controller drivers, represented in the kernel by the `device_driver` structure.
- x **Bus drivers:** drivers responsible for managing a device connected to the bus, represented in the kernel by the `device_driver` structure.
- x **Bus devices:** devices connected to the bus, represented in the kernel by the structure `device`.





BUS INFRASTRUCTURE (cont.)





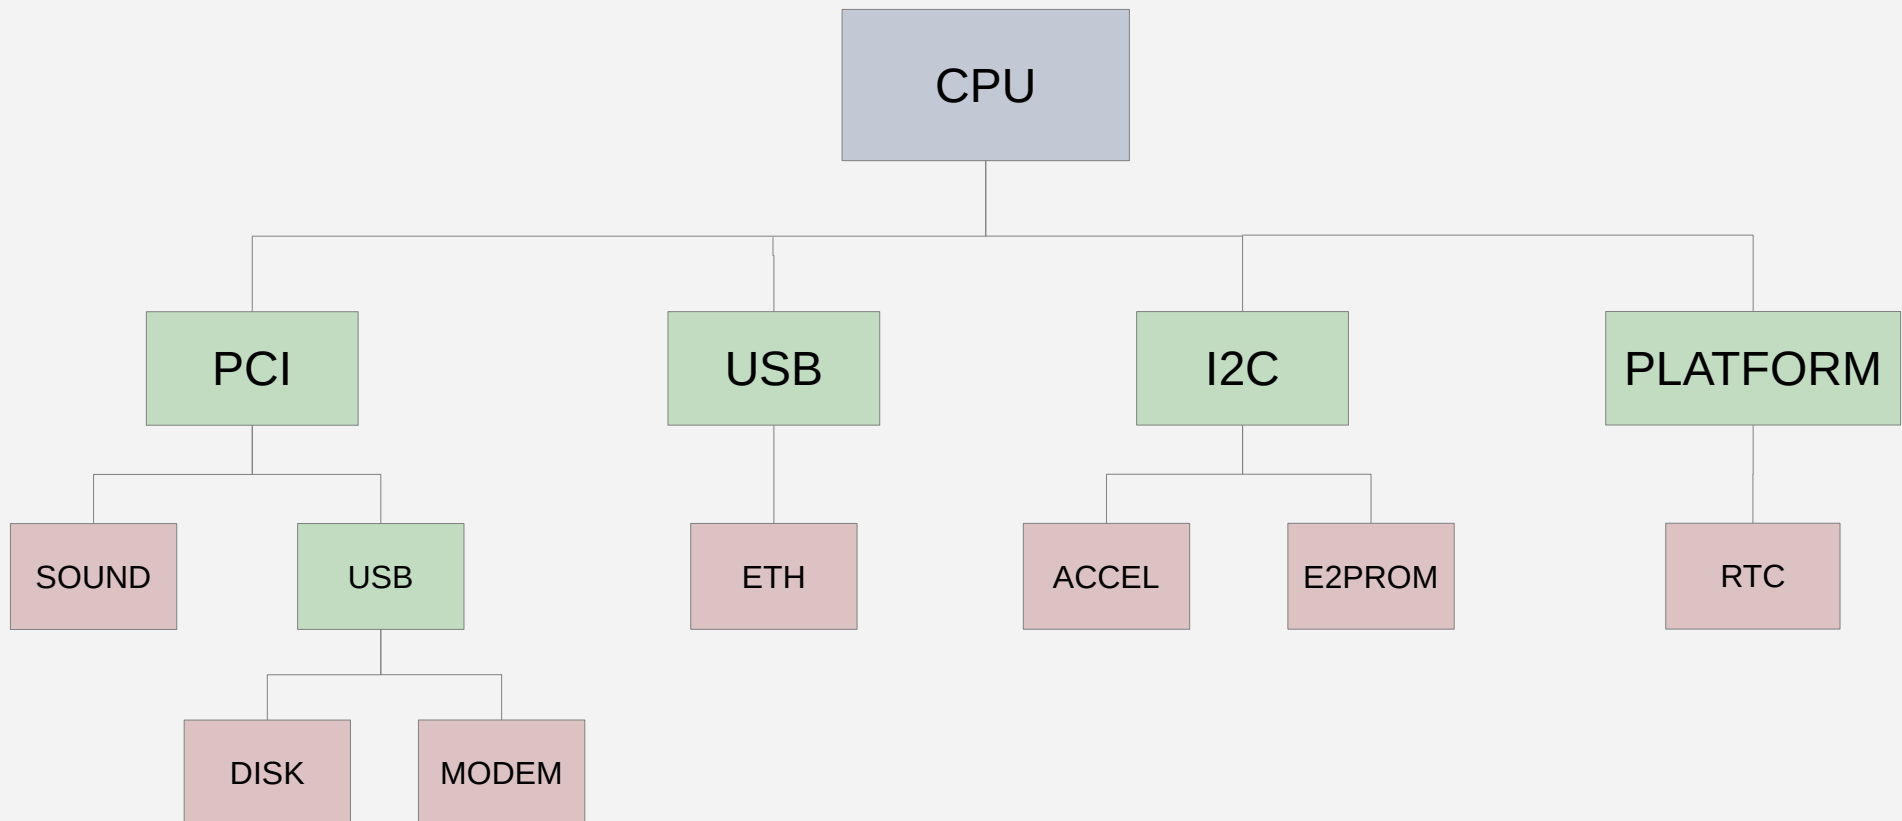
ADVANTAGES

- x More control over the access on a given bus.
- x Separate the driver code from the device description.
- x Easier to identify the hierarchy of devices and buses.
- x Improves power management.
- x Improves usability and modularity.



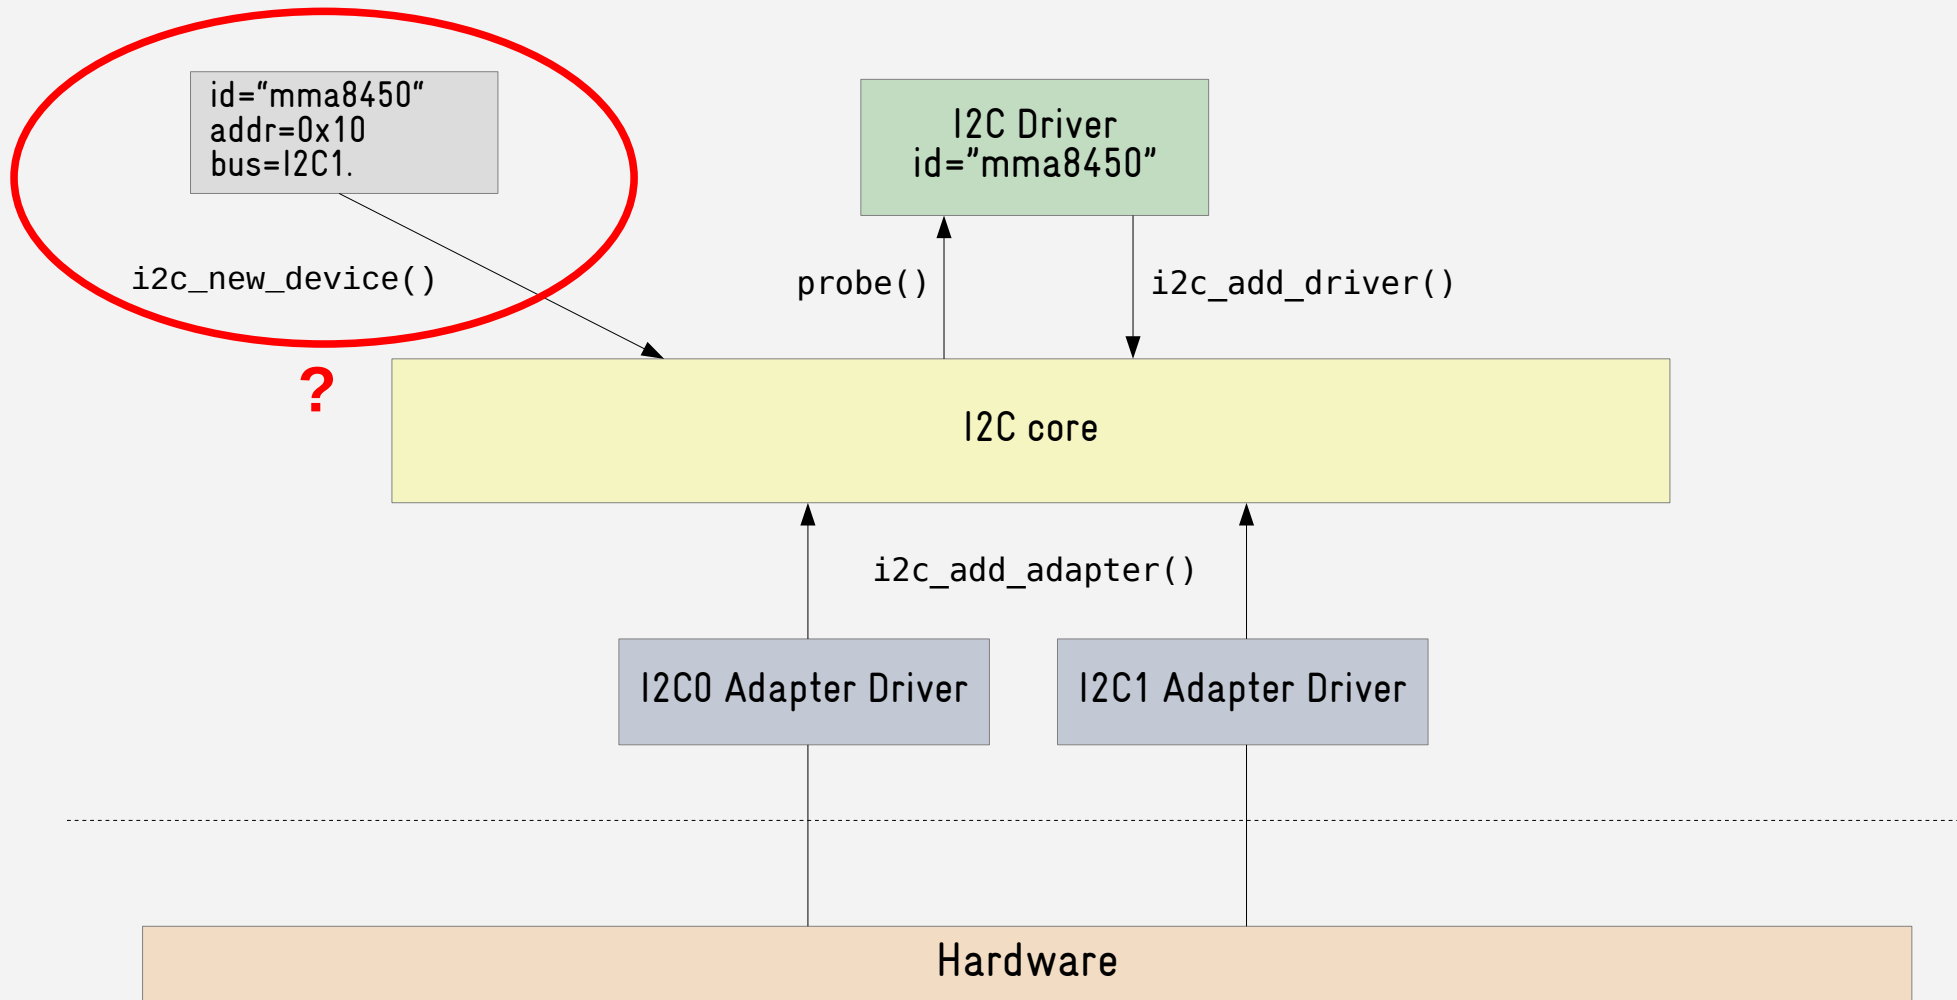


BUSES AND POWER MANAGEMENT



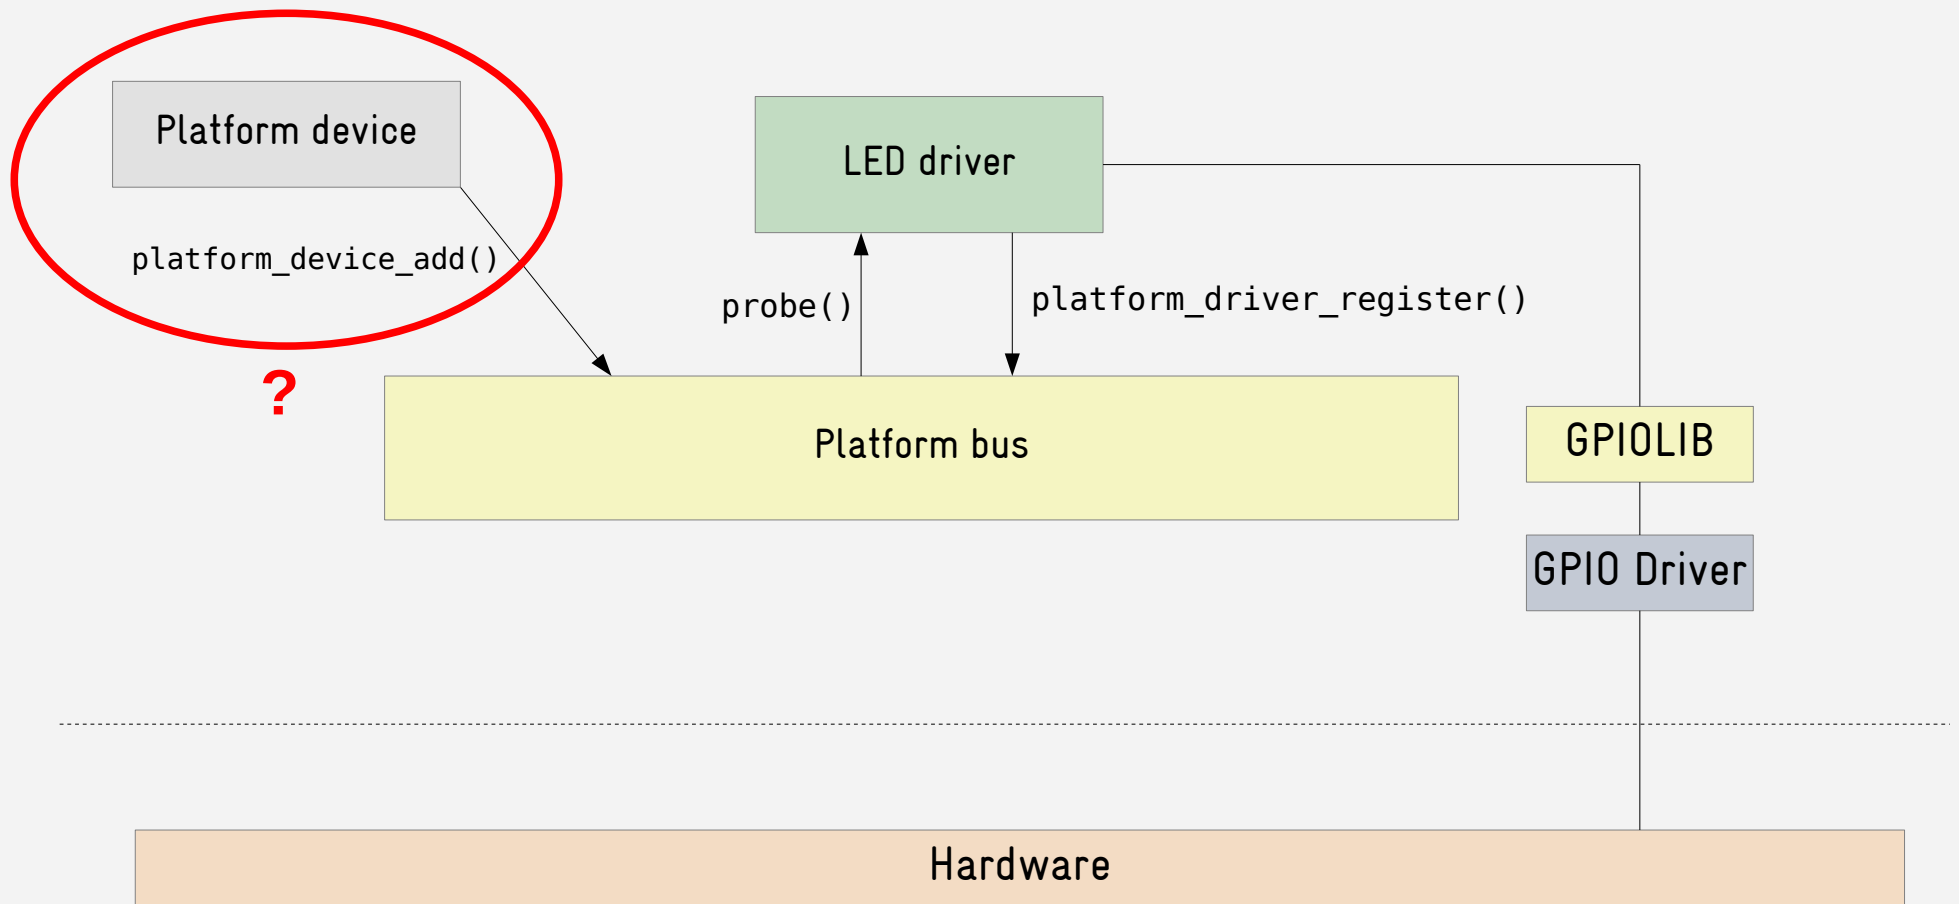


I2C BUS





PLATFORM BUS





REGISTERING A DEVICE

- x A device can be statically registered in the code through an API provided by the bus core, such as `i2c_register_board_info()` or `platform_device_register()` (this mechanism is deprecated).
- x Devices can be described and registered via a mechanism provided by the hardware platform (e.g. ACPI on x86).
- x Devices can be described and registered via a device tree (standard mechanism in some architectures such as PowerPC and ARM).
- x Buses that support device enumeration can automatically identify and register devices (e.g. USB and PCI).



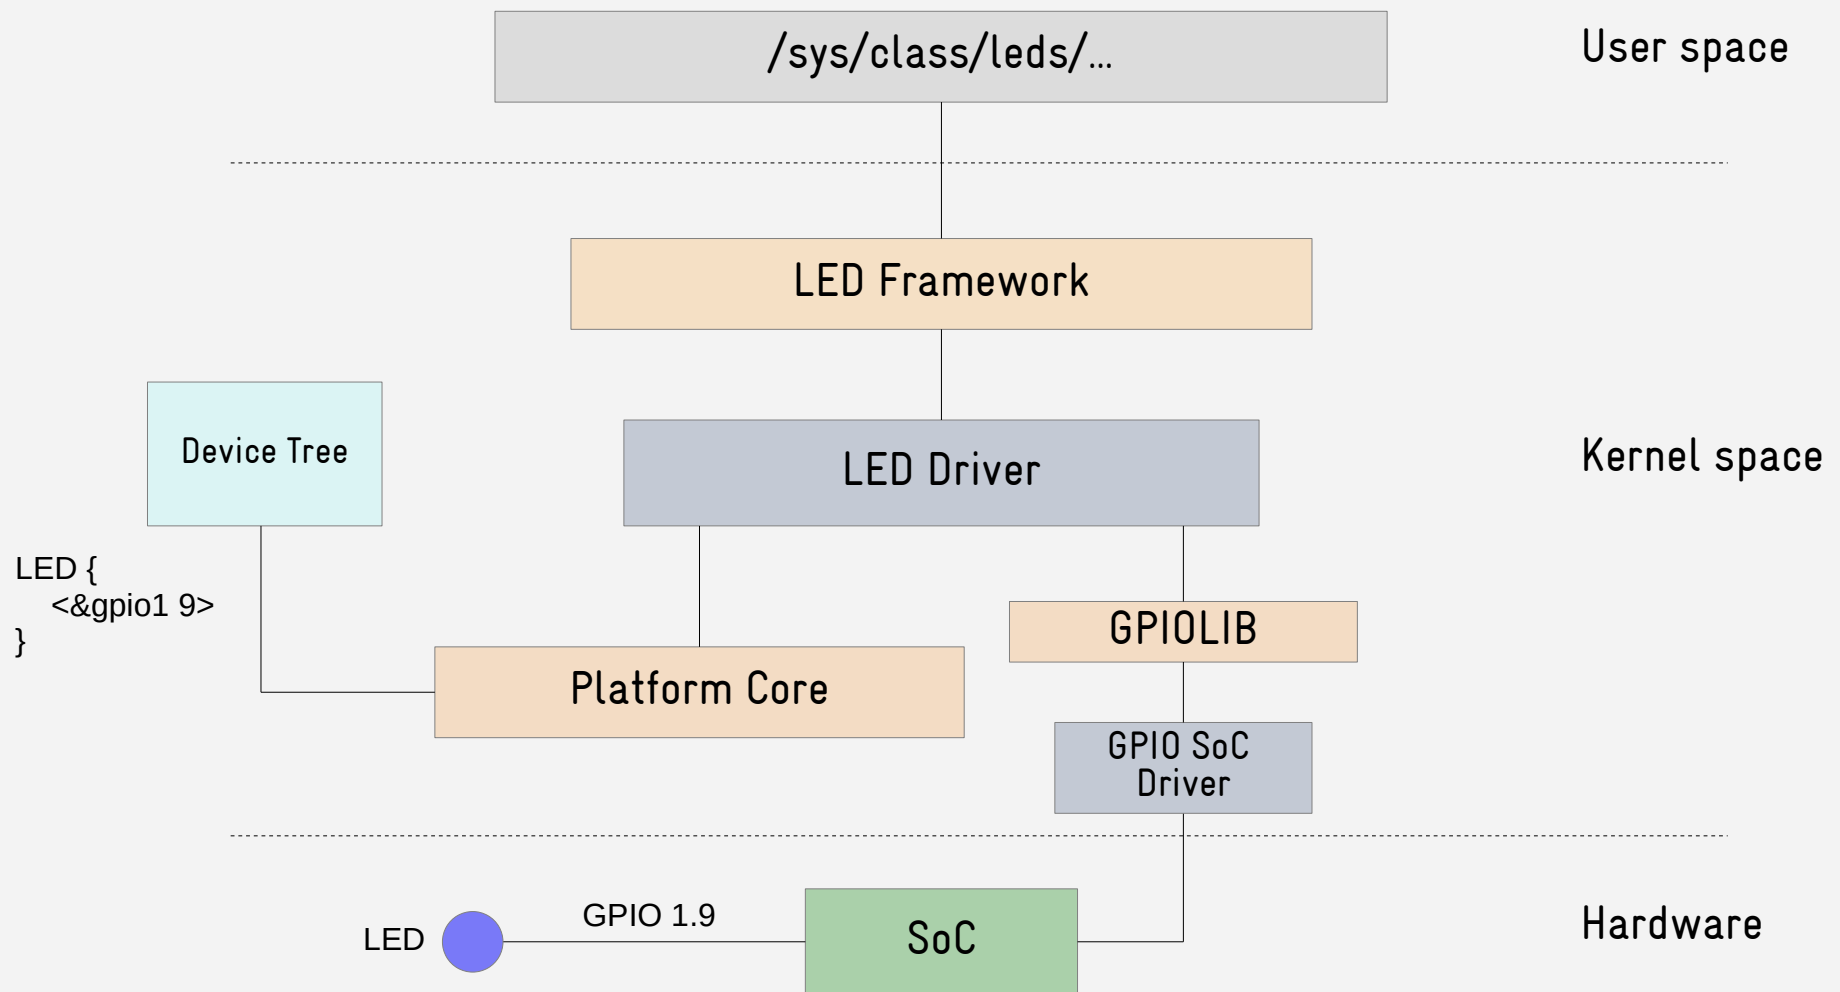


Hands-on 5

Platform bus and device tree

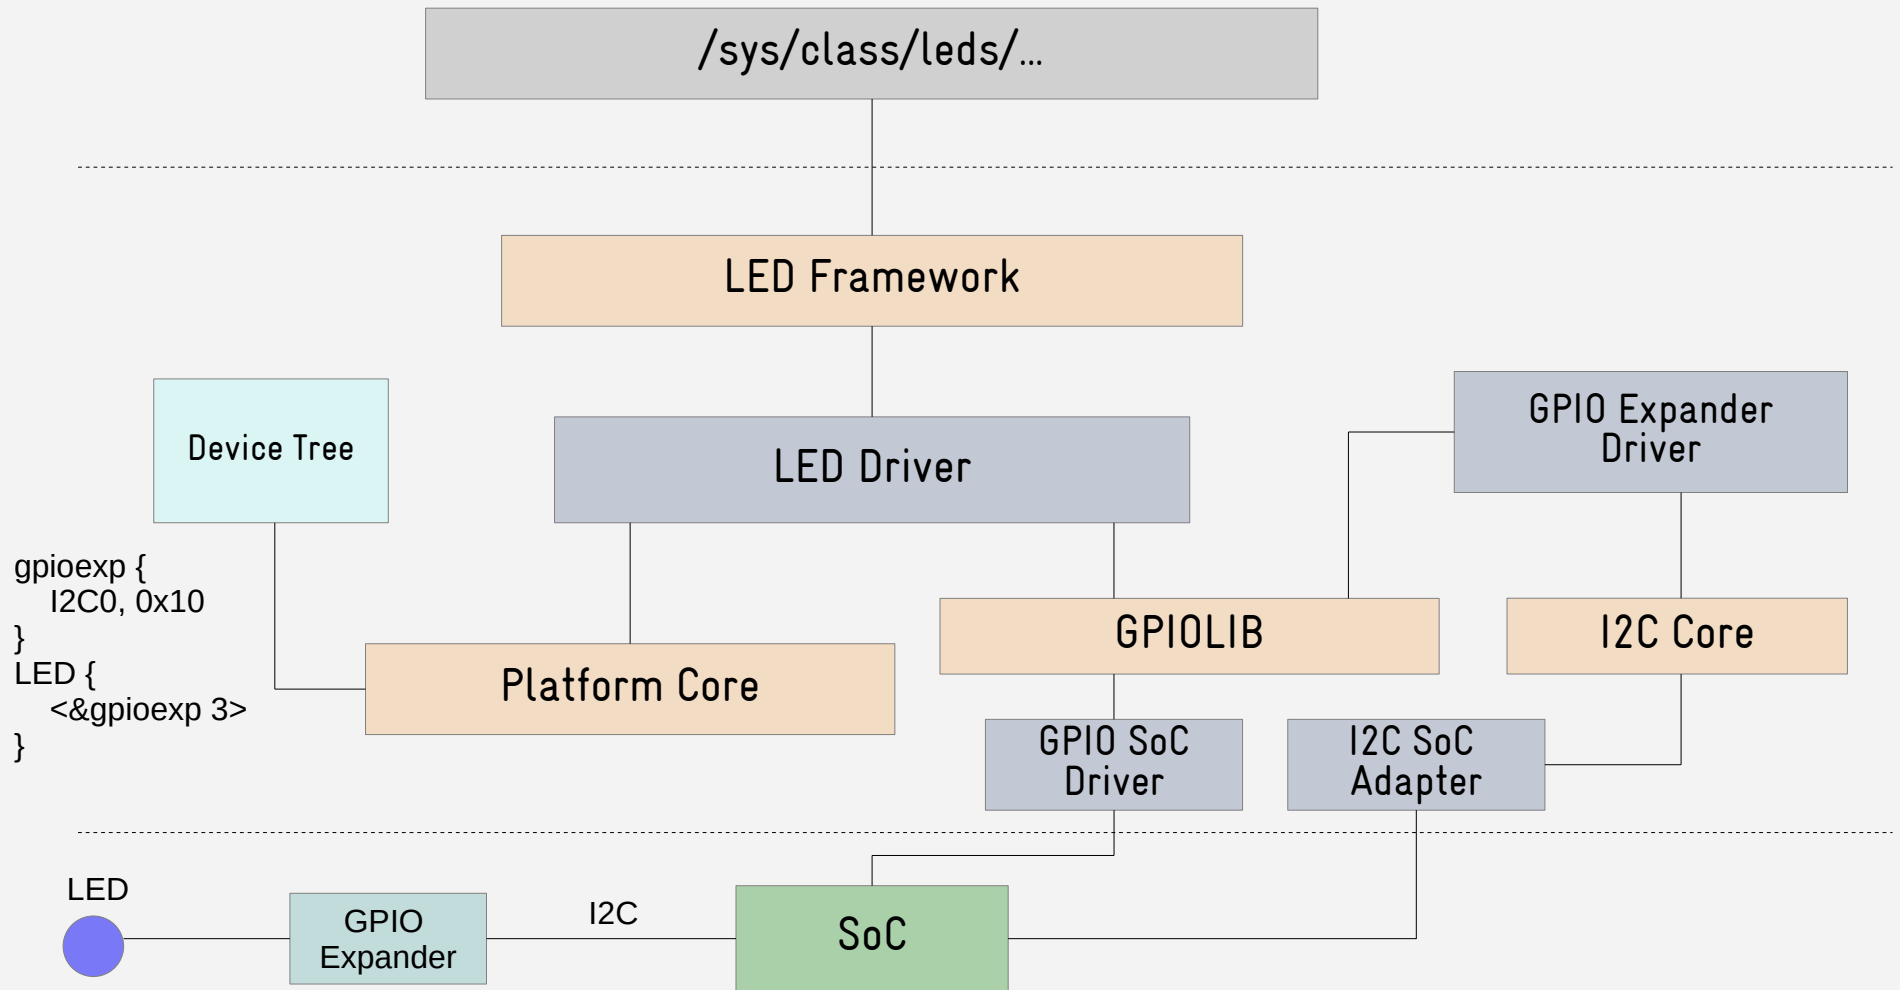


A FLEXIBLE MODEL





A FLEXIBLE MODEL (cont.)



Thanks!

E-mail sergio.prado@e-labworks.com
Website <https://e-labworks.com/en>
Linkedin <https://www.linkedin.com/in/sprado>
Twitter <https://twitter.com/sergioprado>
Blog <https://embeddedbits.org>



Embedded Labworks

