

Designing with Freescale

-


Embedded Linux for i.MX Applications Processors - Real-World Projects

Marcin Bis

<http://bis-linux.com>
martin@bis-linux.com

Warsaw, Poland - June 26th 2013

- **Marcin Bis**
- Embedded Linux (from administration to programming: system, kernel, device drivers)
- Trainings, consulting, support (<http://bis-linux.com>)
- Industrial appliances (Linux + Real-Time)





- To get familiar with Embedded Linux @ **Freescale** platforms.
- To provide an in-depth understanding of Embedded Linux, and Real-Time derivatives of Linux.
- To enable **You** to put together a working Embedded Linux system with "rich GUI" application.
- To teach how an Open Source software is developed and how to make advantage of this process.

We will also build a working **ICS** (industrial control system) - a laser.

- 1 Why to consider Linux?
- 2 Let's build ICS
- 3 What we will need (hardware + tools)?
- 4 Real-Time
- 5 ICS using PREEMPT-RT
- 6 ICS using Xenomai
- 7 ICS using multiple processors (cores)

Why to consider Linux?

1 Why to consider Linux?


Pros ...

... and cons

Why to consider Linux?

Booming usage of Linux in embedded environments. From consumer electronics to industrial control devices.

- Hardware capable of running Linux become cheaper.
- Linux supports lots of hardware, as well as communication protocols.
- Security is easy to achieve.
- Code is easy to develop and reusable.



- Android - customized Linux system.

No additional fees: no per-device licence cost, no need to upgrade constantly.

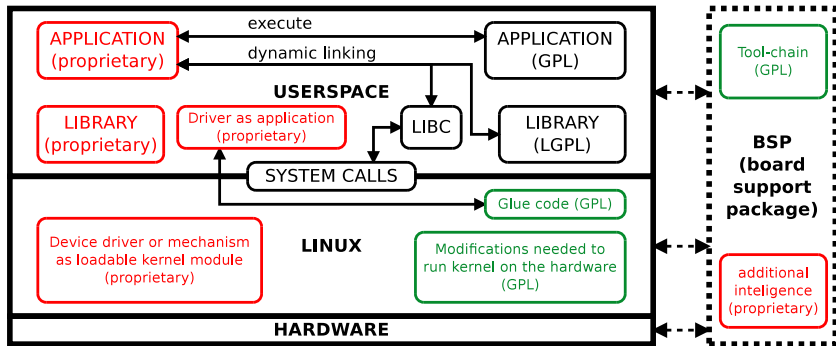
E.g. Availability of the source code, lets you easily upgrade old system/product to the latest protocols and drivers.

Linux support all embedded SoC features. Device connectivity features are implemented in universal frameworks.

- SOC drivers
- I2C, SPI, GPIO
- Flash (including flash-transaction-layers) and flash aware file systems (UBIFS, JFFS2)
- EEPROM
- I/O expanders, custom embedded parts
- networking controllers, including Wireless.

Mixing GPL with proprietary

You can add your proprietary applications, of course. . .



Do you want to build a customer-programmable device?

Let user program it in C, C++, Java, Python, Perl, PHP, Lua and even more.

Tool-chain can be provided at no additional cost (because it is open-source).

E.g. Python programmed "software PLC" device. Customer likes it to be Java programmable - not a big issue in Linux.

Stable, secure and efficient networking stack.

TCP/IP and variety of application-level protocols.

Easy encryption integration. WiFi, broadband, Bluetooth and other protocols.
VNC, Web Server, ftp, SSH...

Examples:

- Consumer device, connected to local network is recognized by Windows and Apple computers as a network neighbourhood party. In order to emulate windows behaviour it runs a cut down SMB, NetBIOS, Upnp stack.
- Device shall be easily accessible from smartphone (web browser), using secure protocol as well as a VPN connection.
- Secure (encrypted) firmware update over the network - no more STUXNET.

Design and write our software once. . . compile, run and deploy everywhere.

- Application is Linux-base rather than microcontroller-based.
- Can be developed and tested on a PC.
- There is variety of available libraries and options.
- Easy: GUI, input and display support.

You are no longer facing vendor lock-in!. Software can be easily ported to another Linux SBC.

E.g. Demonstrate application running on RaspberryPI, as well as on some more secure and manufacturing friendly platform.

If application specification needs Real-Time response. There are a few ways to transform a Linux system into Real-Time operating system.

Easily implement customers demand. Even one-time wishes.

Lower cost of developing one-time software variants:

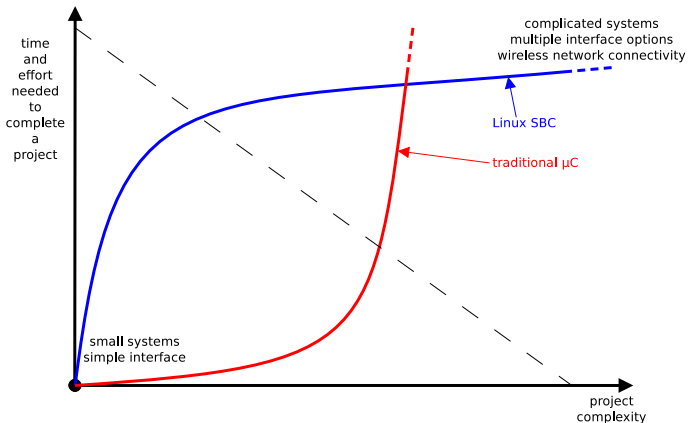
- Let's use twice as big screen.
- Or add barcode reader and rotary encoder.
- We want to control device remotely via IP connection.
- ... but customer's network is running IPv6.
- OK, and now let's display user interface in Russian/Chinese

You:

- Learning curve.
There is a great amount of knowledge to get familiar with. Variety of skills: from electronics design to system administration is needed to make a "helicopter view of a system".
- Fragmented knowledge.
Linux is a go-fast technology rather than finished project. There is no central knowledge base nor single vendor providing support (actually there are many of them). Relevant support can be difficult to find on the Internet or applicable only to specific version of source code.

Learning curve

Simple project: bare-processor or RTOS applications are easy to establish - just install BSP and get familiar with documentation.



Linux project: knowledge of the whole software stack is needed. But, afterwards, adding sophisticated features is very easy.

Your Product:

- Memory footprint.
Power and space limitations. Microcontroller systems used for many tasks has (e.g.) 16kB of RAM and 128kB of Flash in single chip. Typical universal Linux SBC with network connectivity and full user interface has 64MB of RAM and 256MB of Flash. Even small Linux system needs 32MB RAM/8MB Flash.
- Development Tools.
Most Linux system development and debugging can be done using Eclipse, although there is no install and ready solution (like tools for various microcontrollers). IDE has to be properly configured by a user.

- So: is Linux the answer for You?
 - No! Linux is a question to consider.
- Definitively worth considering.

Let's build ICS


2 Let's build ICS

Case 1 - Weight dosing

Case 1 - analysis


Case 2 - Laser

Case 2 - analysis



Weight-dosing process - specification

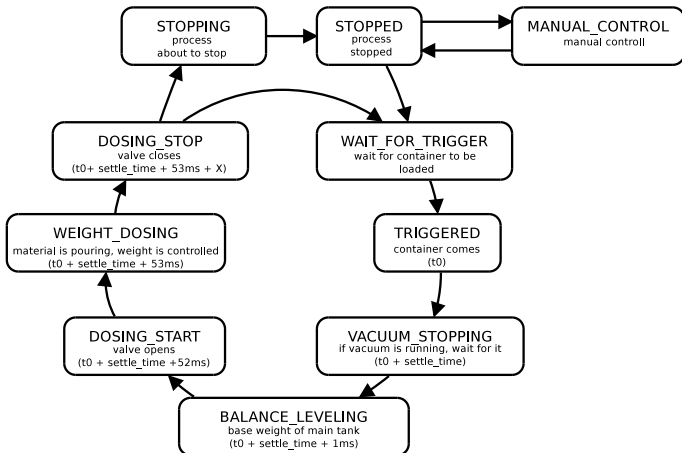
Loose material (or fluid)
is loaded into containers.



- the main tank is suspended on weight (tensometer)
- conveyor or robot provides containers, appearance of the container triggers interrupt
- the valve opens, and the material is poured into a container
- amount of material is measured by reading data from the weight
- main tank has a limited capacity, it can be replenished from main silo by turning on vacuum
- if vacuum is turned on, it has to work for some minimum time, while vacuum is working, material cannot be poured.

Weight-dosing process - analysis

Weight dosing process can be modelled as a finite state machine.



Process starts on `WAIT_ON_TRIGGER` state. If triggered, it runs on timer (1ms).

Hardware

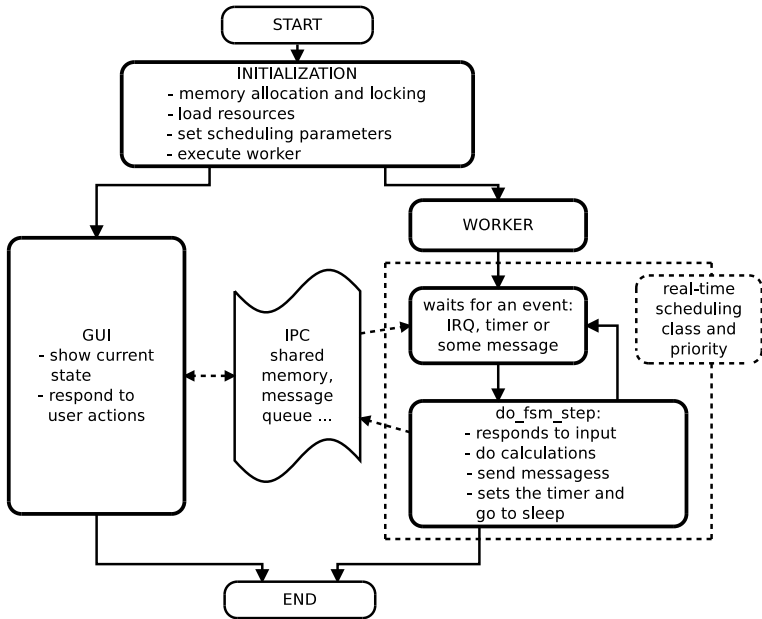
- PC for development, Ubuntu 12.04
- **Freescal**e Board

GUI part

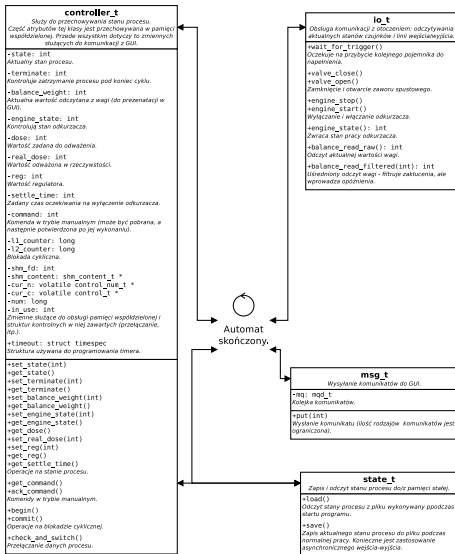
- written in QT/C++
- userspace components provided by: **Buildroot** or **LTIB** or **YOCTO** or something else...
- ext4 on SD card as primary storage or NAND flash storage

Real-Time operations

- implemented as separate process in C
- which allows easy porting
- communicates with GUI using shared memory and message queue, in a lockless way:
 - two control structures are stored in SHM
 - one is utilized by running process, other can be changed by GUI
 - then structs are switched



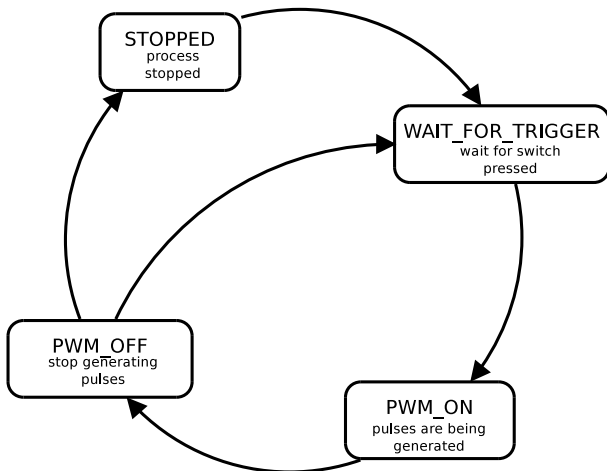
System boot constraint: must be operational under 10s (2.5s RT task, 8-9s GUI).

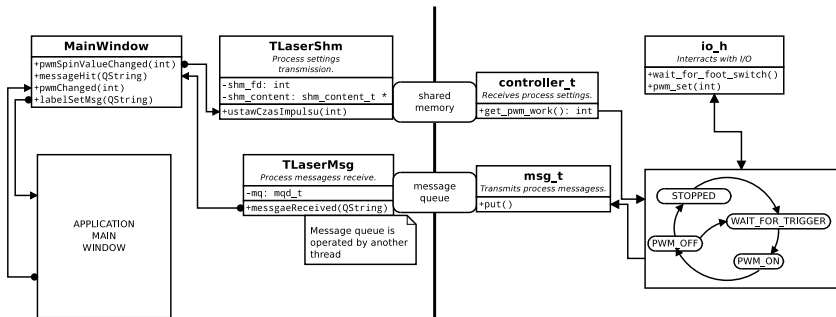


- PWM controlled laser diode (actually real lasers power circuits are far more complicated)
- period: 200 μ s
- laser must send a given set of pulses
- laser operation is triggered by user pressing button

Actual model implemented with laser pointer.

- Pressing button triggers PWM for exactly 900ms
- Button both: virtual and physical





What we will need (hardware + tools)?

3 What we will need (hardware + tools)?

- Options

- MQX Lite


- Linux

- Linux + MQX

Using Linux...




Weight dosing - pneumatics are used (actuator latency is 15ms).




Threads - share virtual memory, have different scheduling settings.

Welding machine - μ s

Medical laser controller - μ s (or even less)




Xenomai provides better latency and predictability.




Special hardware can be utilized too:


- two processors: eg. additional Cortex-M for running worker task
- multicore systems: eg. Freescale Vybrid (Cortex-A5 + Cortex-M4)


Why not?





i.MX6 Quad





Real-Time

- 4 Real-Time
 - Theory
 - Latency tests
 - Testing circuit
 - Unmodified Linux kernel
 - Real-Time - concepts
 - Real-Time - measurements
 - Not so good results ...
 - Where is latency coming from?
 - How to achieve Real-Time in Linux?

Popular definition

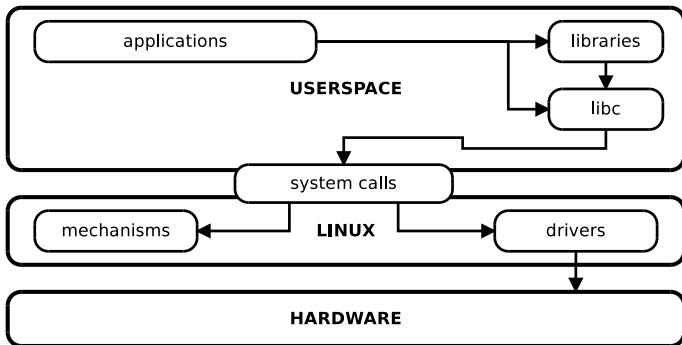
Correctness of operation depends not only on whether performed without error, but also on the time (the upper limit) in which the operation completed.

Practical definition

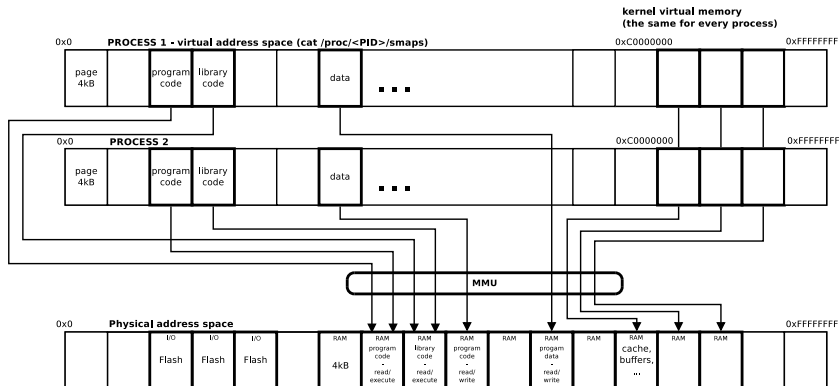
RT system is one in which can be proved that any required operation will be completed in a certain time.

- Mathematical proof would be perfect. Unfortunately systems are so complex, it is not possible.
- System is tested (TDD). If deadlines are met (under load) for all use-cases, system is Real-Time.
Note: In some cases (eg. certification for safety-critical tasks), full-code coverage would be needed!


- Linux kernel is designed to be „democratic”
 - resources are equally disposed
 - eg.: scheduler avoids process starvation
- Usually, determinism is not taken into account
- **throughput** is.




Most layers and subsystems are complex:



What are we testing?







Input is triggered on falling and rising edge, output state changes according to input.

In this case, we are using GPIO pin-s.
Input can be other external or internal: timer, camera, network PHY, ADC etc.






01_inout.c


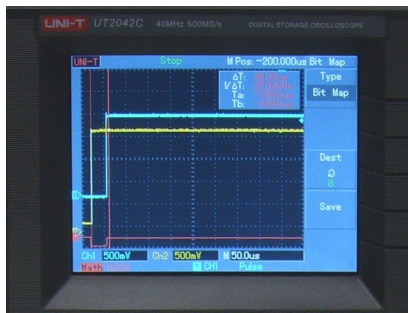


02_uinout.c






Code is on GitHub:

<https://github.com/marcinbis/mb-rt-data.git>



Results





Deadline

Point in time, before which the action (system response) must occur.

- **Hard Real-Time** - deadline must be meet (fatal error if not).
- **Firm Real-Time** - deadline should be meet (system response is useless otherwise).
- **Soft Real-Time** - deadline should be meet, but nothing critical will happen if not (eg. decreased user experience, sample drop ...).

Latency

The time between the moment in which the action was to occur, and in which, in fact, occurred.

Jitter

Undesired deviation of latency. For various reasons, latency is not constant. Too large jitter, renders system unusable for data acquisition.

Predictability

How much time, the action will take (eg. from IRQ occurred to handler finished executing).

$O(1)$ algorithms should be used.

Worst Case

Due to imperfect nature of real-world systems, we are considering the Worst Case.

We have to know the latency in worst possible case.

```
$ cat /proc/loadavg  
5.02 3.76 2.04 2/47 432
```

- I/O on SD card:

```
cat /dev/mmcblk0p1 > /dev/null
```

- sending ASCII data to serial console:

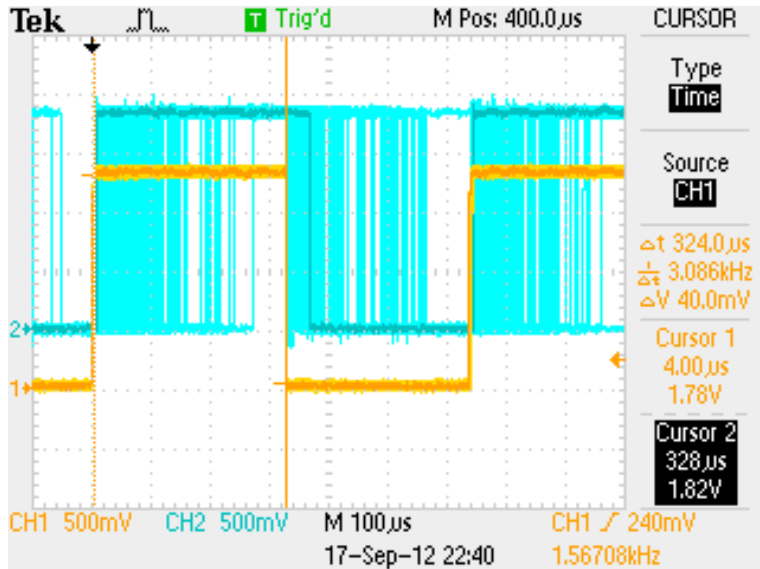
```
cat /dev/zero | od -v
```

- send network packages:

```
ping -f <ip address>
```

WARNING!: these tests just generate IRQ, they are not showing real-life load.
Use real-case tests.

Results under load

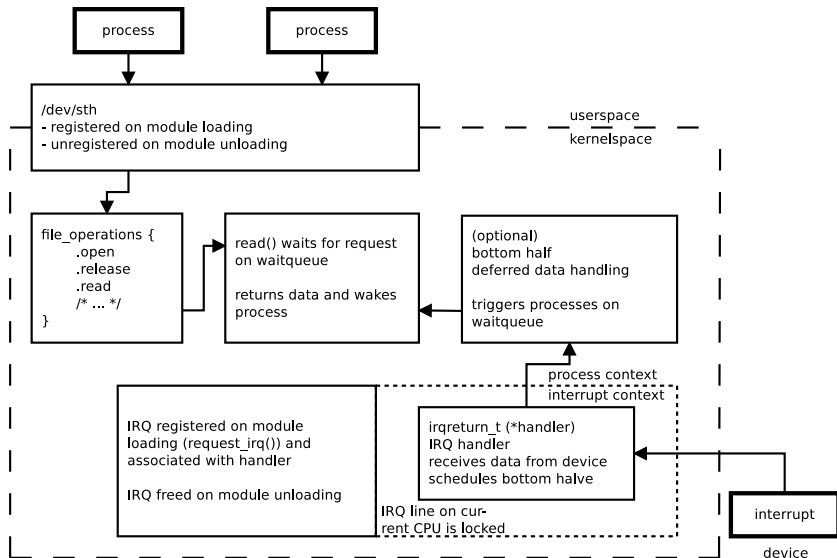


Linux:

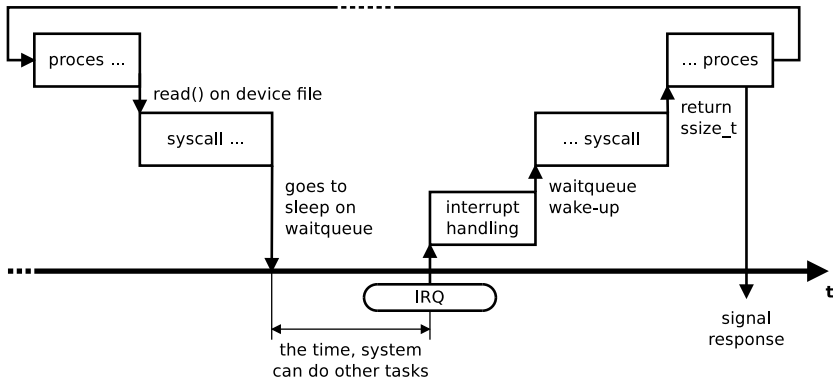
Separates:

- logic (in userspace)
- mechanisms (provided by kernel)

Interrupt-based I/O

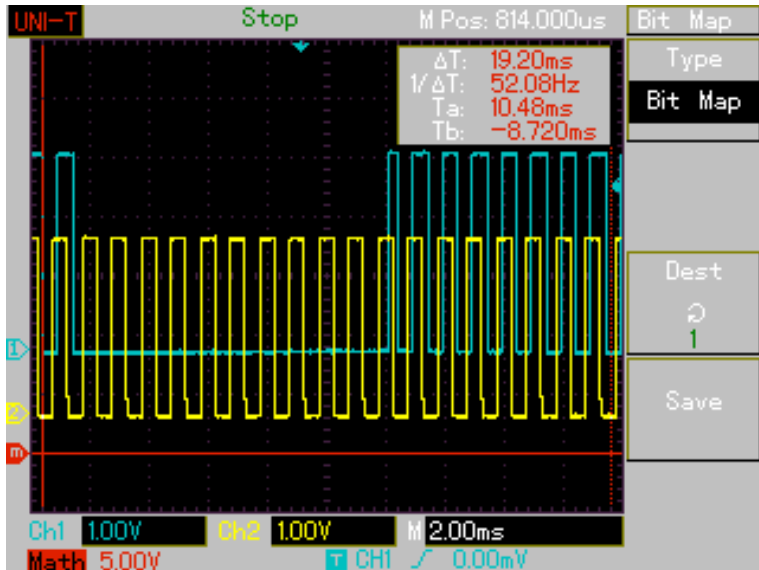


Interrupt-based I/O - another view

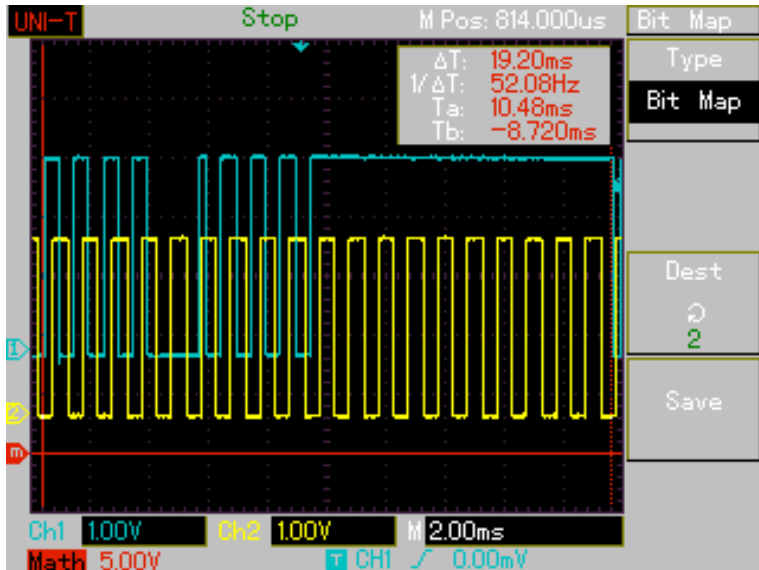


- `03_cinout.c`
- `04_real_cinout.c`
- **In case of GPIO:** `/sys/class/gpio/` can be used as well (`poll()`, `read()`, `write()`).

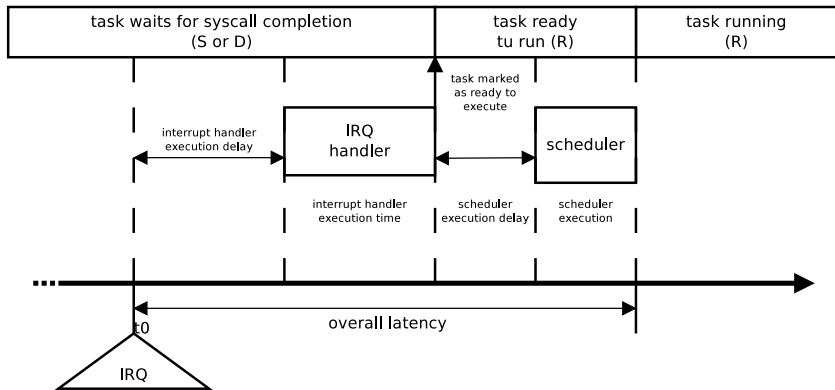
Userspace - results under load



Userspace - results under load




Where is latency coming from?



1 Micro-kernel approach:

- RTLinux - <http://en.wikipedia.org/wiki/RTLinux>, there used to be open-source version: <http://www.rtlinuxfree.com/>.
- Adeos/I-Pipe - <http://home.gna.org/adeos/> - common base.
- RTAI - <https://www.rtai.org/> - minimum possible latency.
- Xenomai - <http://www.xenomai.org/> - provides various APIs.



2 In-kernel approach:

- RT PREEMPT - <https://rt.wiki.kernel.org>
<http://www.kernel.org/pub/linux/kernel/projects/rt/>


1 Dual processor/microcontroller approach

- Dual/Quad Core i.MX6 - devote one core to Real-Time task.
- Additional μ C - add (cheap) microcontroller and run RTOS on it (ore bare-metal app).



2 Heterogenous cores:


- Freescale Vybrid



ICS using PREEMPT-RT

- 5 ICS using PREEMPT-RT
 - PREEMPT-RT
 - Implementation
 - Internal latency measurements
 - Demo

1 Standard kernel



2 Interrupts as threads


Kernel Features --->

Preemption Mode (Complete Preemption ()) --->


(X) Complete Preemption (Real-Time)

-- Thread Softirqs /* 2.6.33 */

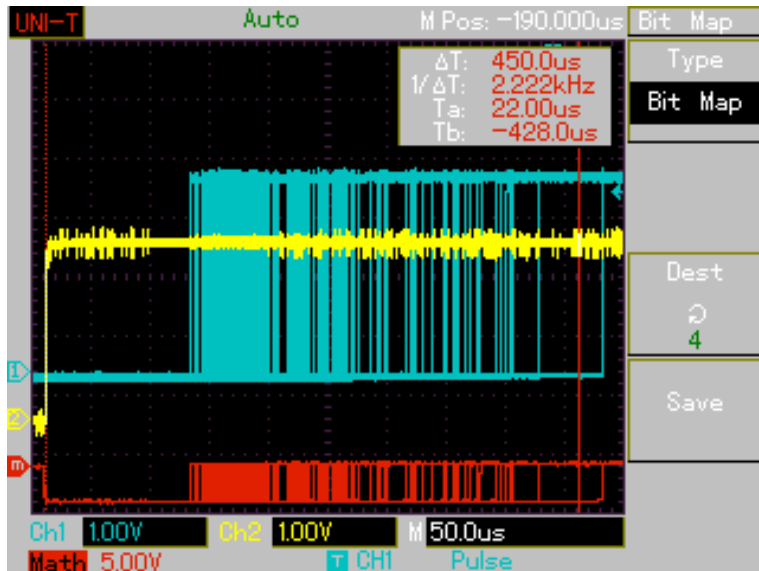
-- Thread Hardirqs



04 under load



04 under load



- **Real-Time != Real Fast**
Maximum latency (Worst Case) is limited, but minimum latency is bigger.
- **Kernel with RT-PREEMPT patch, does not make the whole system Real-Time**
- Specially designed application and POSIX RT-API should be used:
 - Defined: IEEE 1003.1b. Linux supports it.
 - Scheduler
 - Memory locking
 - Shared memory
 - RT signals
 - Semaphores (priority inheritance)
 - Timers (esp. `CLOCK_MONOTONIC`)
 - AIO

Use appropriate programming language

- **C** - but make it object-oriented (for reference - Linux kernel: buses, drivers, classes etc.)
- **C++** - would be nice too
 - cannot be utilized inside kernel or as Xenomai kernel process
 - can be executed as bare-metal μ C or in userspace
- Utilize design patterns.

Set the proper scheduler class and priority

```
struct sched_param sp;  
sp.sched_priority = MY_PRIORITY;  
ret = sched_setscheduler(0, SCHED_FIFO, &sp);
```

- Interrupts run in threads, default to: SCHED_FIFO/50.
- ...do not forget to fine-tune them.
- SCHED_DEADLINE can be helpful too.

Lock all memory (mlock)

```
mlockall(MCL_CURRENT|MCL_FUTURE);
```

Try to cause page_fault (allocated memory, data from files)

```
buf = malloc(BUF_SIZE);
memset(buf, 0, BUF_SIZE);
```

- Memory is locked, so it stays on place.

Prefault the stack (it can be shared within process we have forked from)

```


/* GCC will not inline this function */
__attribute__((noinline)) void stack_prefault(void)
{
    unsigned char tab[MAX_SAFE_STACK];
    /* GCC will omit optimizations */
    asm("");
    memset(tab, 0, MAX_SAFE_STACK);
}
/*...*/
stack_prefault();

```


Use POSIX timer to do the fsm step (in a proper way)

```
#define NSEC_IN_SEC 1000000000L
#define INTERVAL 1000000L
struct timespec timeout;

clock_gettime(CLOCK_MONOTONIC, &timeout);
while (1) {
    do_fsm_step(&some_data);
    timeout.tv_nsec += INTERVAL;
    if (timeout.tv_nsec >= NSEC_IN_SEC) {
        timeout.tv_nsec -= NSEC_IN_SEC;
        timeout.tv_sec++;
    }
    clock_nanosleep(CLOCK_MONOTONIC, TIMER_ABSTIME,
                   &timeout, NULL);
}
```



```
clock_nanosleep(CLOCK_MONOTONIC, 0, &interval, NULL);
```

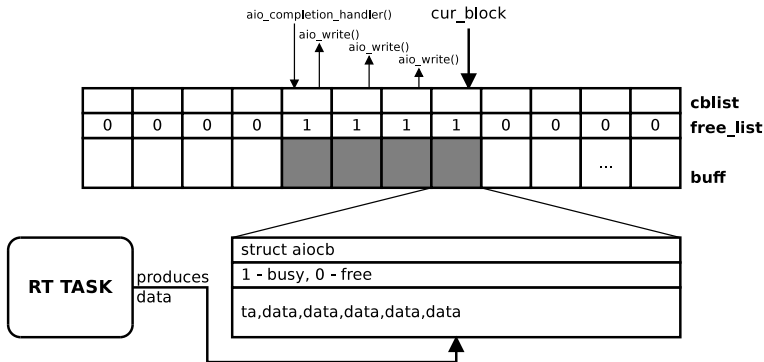


```
clock_nanosleep(CLOCK_MONOTONIC, TIMER_ABSTIME, &timeout, NULL);
```

Utilize AIO to write or read data (eg. sensor data, production logs)

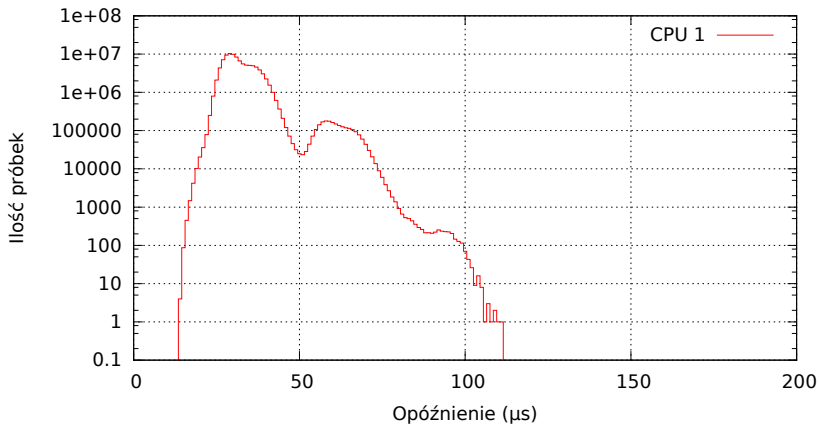
```

struct aiocb {
    int          aio_fildes //File descriptor.
    volatile void *aio_buf  //Location of buffer.
    /* ... */
};
aio_write(struct aiocb *);
aio_return(struct aiocb *);
  
```



Freescale i.MX53 3.4.25-rt36 PREEMPT RT

Opóznienia: 14 μ s min, 32 μ s avg, 111 μ s max, 0 próbek poza wykresem.



Buildroot demo.

QT rapid development demo.

LASER ICS using PREEMPT-RT

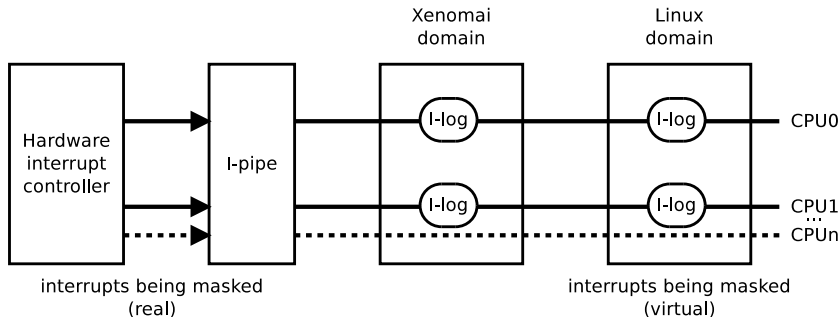
ICS using Xenomai

6 ICS using Xenomai

How does it work?

Internal latency measurements

Demo



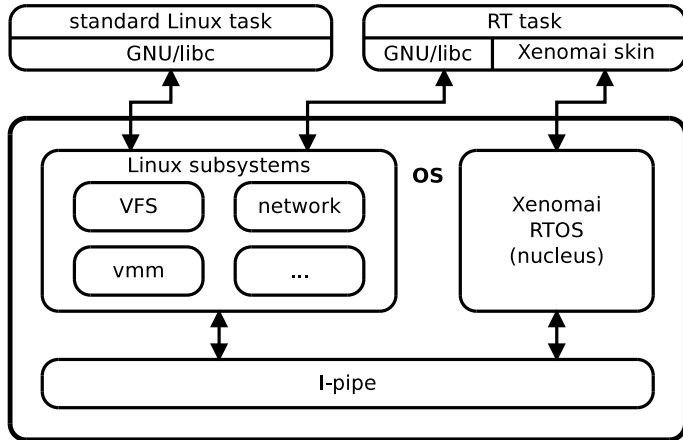
- I-Pipe take control over all hardware interrupts
- All system calls are passed through (I-Pipe)
- Events are dispatched to different I-pipe domains.

```

$ cat /proc/ipipe/Xenomai
+----- Handling ([A]ccepted, [G]rabbed,
|+---- Sticky      [W]ired, [D]iscarded)
||+--- Locked
|||+-- Exclusive
||||+- Virtual
[IRQ]  |||||
 38:  W..X.
418:  W...V
[Domain info]
id=0x58454e4f
priority=topmost

$ cat /proc/ipipe/Linux
 0:  A....
 1:  A....
...
priority=100

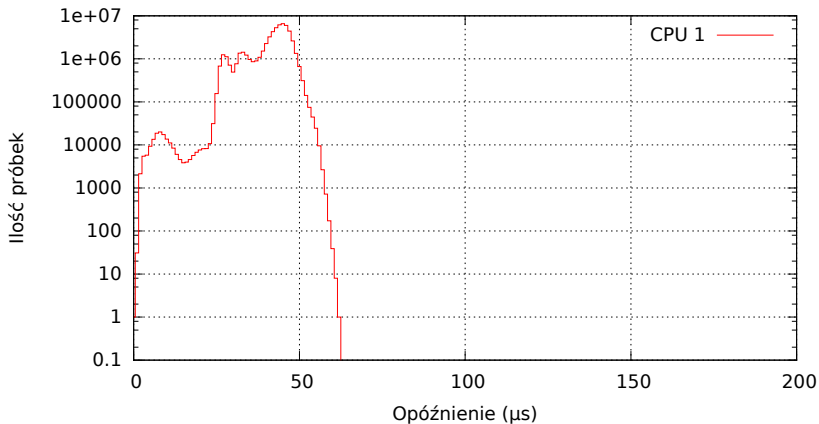
```



There are actually two kernels. Process can migrate between them:

- Xenomai
- Linux

Freescale i.MX53 3.0.36 Xenomai
 Opóźnienia: μs min, μs avg, μs max, próbek poza wykresem.





LASER ICS using Xenomai

ICS using multiple processors (cores)

7 ICS using multiple processors (cores) Demo

- Run MQX-Lite on Cortex-M0 to do the real-time stuff. Run Linux on more powerfull processor.
- **Vybrid**: Linux runs on Cortex-A5, MQX on Cortex-M4





LASER ICS using Cortex-M0.

LASER ICS using Vybrid.

Thank You!



http://bis-linux.com/dwf_waw2013

This presentation, source code and additional materials.

<http://bis.org.pl>

My books.

<http://bis-linux.com/>

Embedded Linux and Real-Time - training and support.

Contact

marcin@bis.org.pl

Questions?