

Porting a safety-critical industrial application on a mixed-criticality enabled real-time operating system

Antonio Paolillo

Paul Rodriguez, Vladimir Svoboda, Olivier Desenfans, Joël Goossens, Ben Rodriguez, Sylvain Girbal, Madeleine Faugère, Philippe Bonnot

5th December 2017

5th International Workshop on Mixed Criticality Systems (WMC 2017)

Joint work HIPPEROS and Thales R&T



THALES

Mixed-criticality?



The IMICRASAR project


Isolated MIXed CRITICALity
Avionics System ARchitecture

Goals:

- create an isolated
mixed-criticality
hard real-time platform


platform = OS + hardware
- support of an industrial application on
the platform
- results: retrieve unused CPU resource

In short, Thales brings us two components...



In short, Thales brings us two components...


and asks us to make a system out of it



Possible solutions



Possible solutions



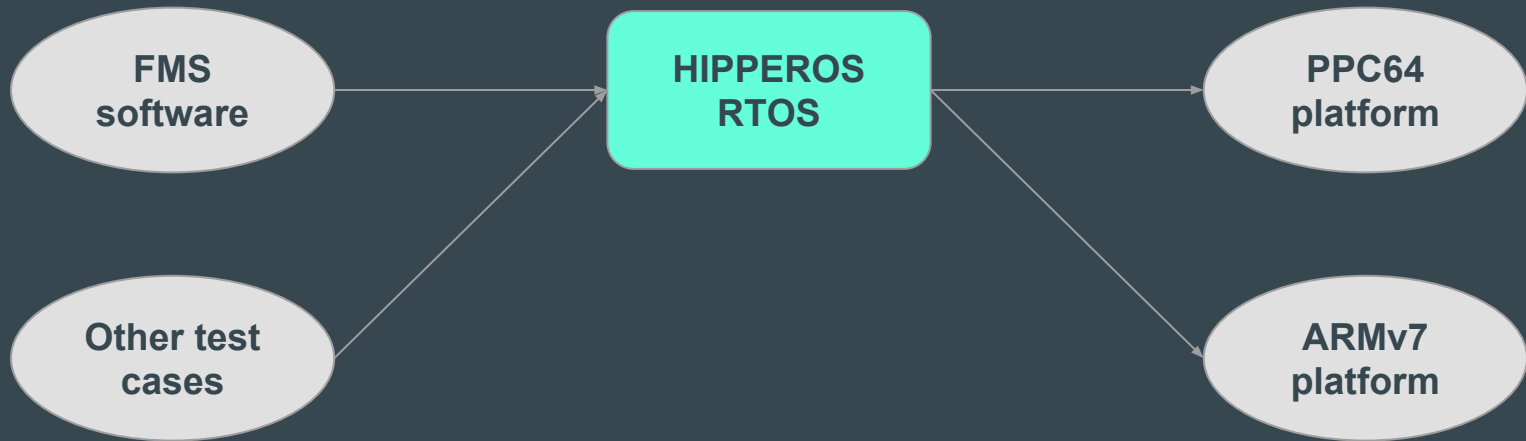
Possible solutions



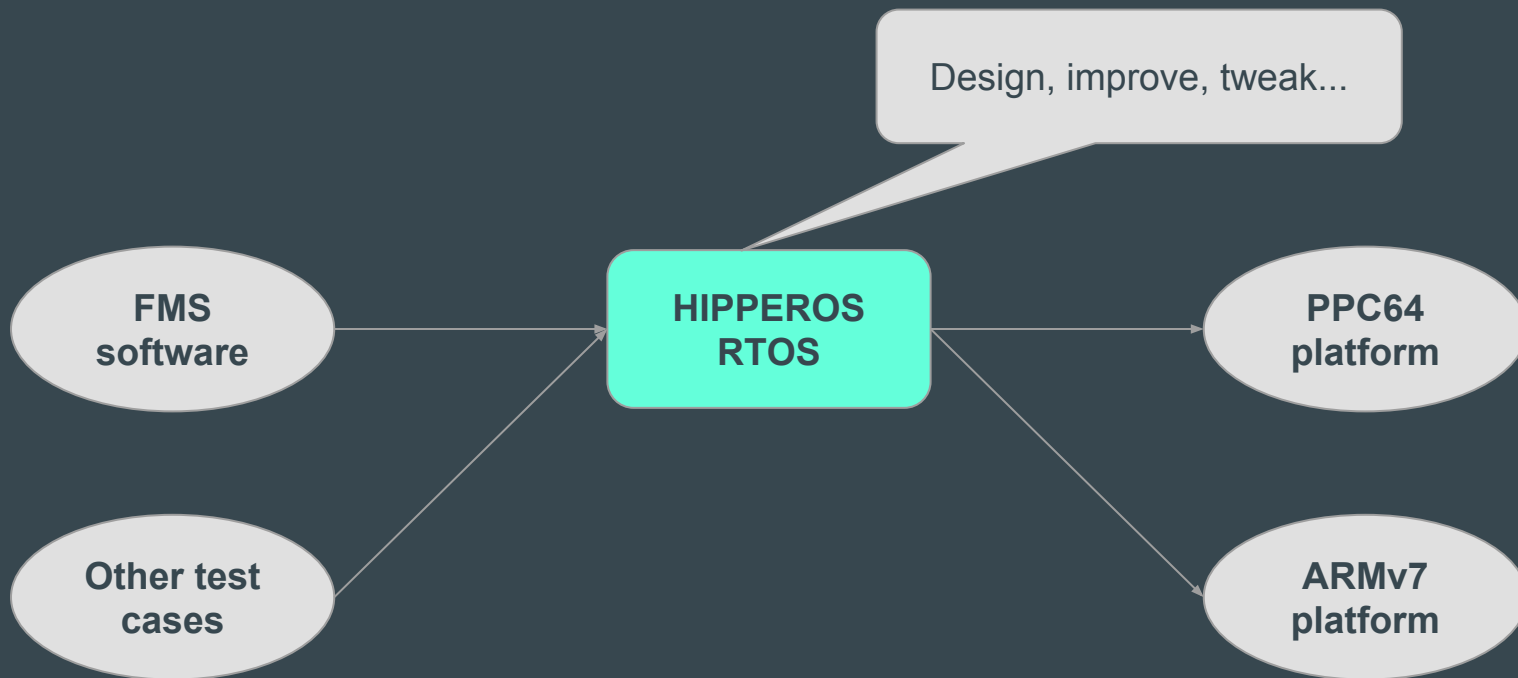
Possible solutions



Possible solutions



Possible solutions



Project roadmap


Project roadmap

1. Support of the PPC64 architecture for HIPPEROS



Project roadmap

1. Support of the PPC64 architecture for HIPPEROS
2. Port the FMS application on HIPPEROS



Project roadmap

1. Support of the PPC64 architecture for HIPPEROS



2. Port the FMS application on HIPPEROS



3. Extend HIPPEROS with Mixed-Criticality scheduling capabilities



Project roadmap

1. Support of the PPC64 architecture for HIPPEROS



2. Port the FMS application on HIPPEROS



3. Extend HIPPEROS with Mixed-Criticality scheduling capabilities



4. Validate the setup with experiments



Project roadmap

1. Support of the PPC64 architecture for HIPPEROS



2. Port the FMS application on HIPPEROS



3. Extend HIPPEROS with Mixed-Criticality scheduling capabilities



4. Validate the setup with experiments



Presentation Agenda

- A. (Short) Presentation of the Thales application
- B. Description of the RTOS
- C. Description of the platform(s)
- D. Experiments

A. The Application

The application use case

The application use case

- Provided by Thales as a case study

The application use case

- Provided by Thales as a case study
- “Autopilot”:
 - Inputs: sensors (GPS signal), database
 - Outputs: display, direction instructions

The application use case

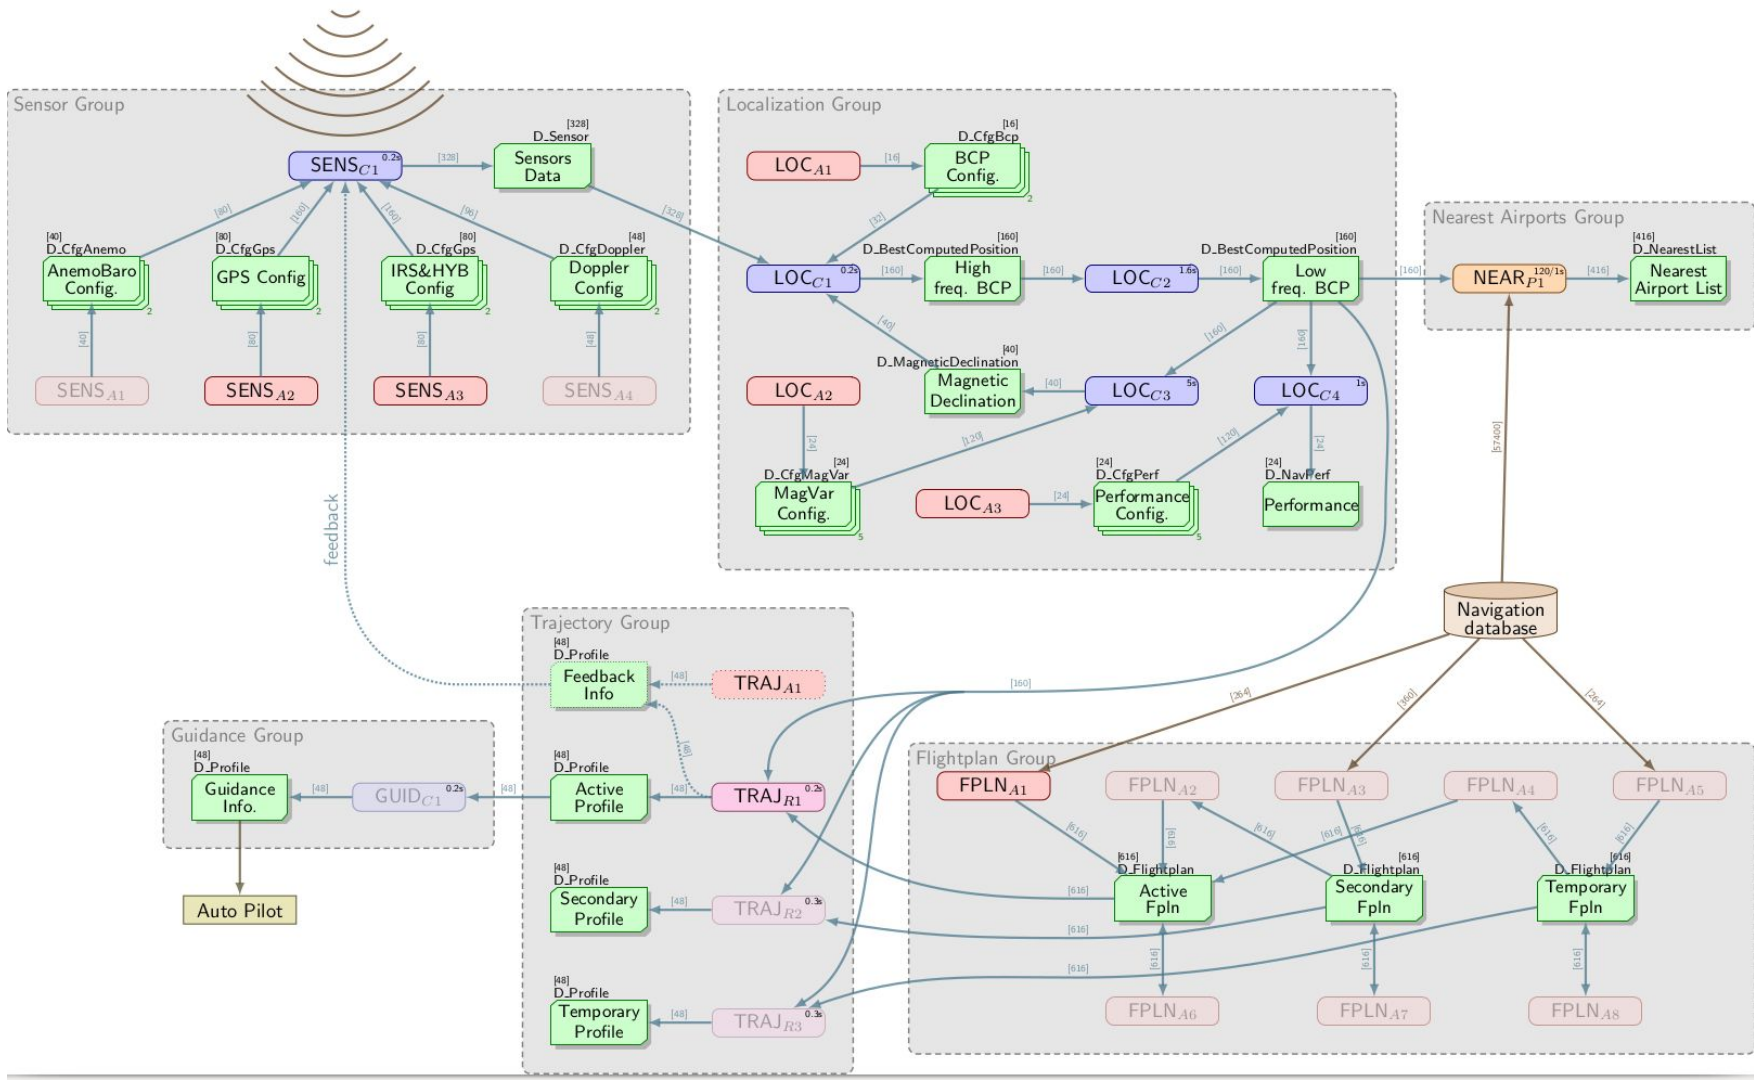
- Provided by Thales as a case study
- “Autopilot”:
 - Inputs: sensors (GPS signal), database
 - Outputs: display, direction instructions
- 18 cooperating high criticality tasks with different periodicity attributes

The application use case

- Provided by Thales as a case study
- “Autopilot”:
 - Inputs: sensors (GPS signal), database
 - Outputs: display, direction instructions
- 18 cooperating high criticality tasks with different periodicity attributes
- Programming model: Acquisition - Execution - Restitution

The application use case

- Provided by Thales as a case study
- “Autopilot”:
 - Inputs: sensors (GPS signal), database
 - Outputs: display, direction instructions
- 18 cooperating high criticality tasks with different periodicity attributes
- Programming model: Acquisition - Execution - Restitution
- The sampled use case takes 90 seconds to run



B. The Operating System

Context: the HIPPEROS company builds HIPPEROS OSeS

- Development of the kernel started in June 2013
- Spin-off company, from Université Libre de Bruxelles, created in January 2014
- Today: ~15 people among them 5 OS developers & researchers
- The goal is to ship certifiable OSeS to safety-critical software industries

Main architecture and design choices

Main architecture and design choices

- Hard real-time operating system

Main architecture and design choices

- Hard real-time operating system
- Embedded targets: **ARMv7**, ARMv8, IA32, **PowerPC 64**

Main architecture and design choices

- Hard real-time operating system
- Embedded targets: **ARMv7**, ARMv8, IA32, **PowerPC 64**
- A new micro-kernel written from scratch
 - Built for user needs, i.e. small footprint and adapted policies
 - Multi-core architecture based on an asymmetric kernel
 - Real-time model for user applications
 - MMU support and virtual address space
 - Resource sharing & IPC protocols (mutexes, semaphores, message passing, etc.)
 - Usual OS services (timers, etc.)

Main architecture and design choices

- Hard real-time operating system
- Embedded targets: **ARMv7**, ARMv8, IA32, **PowerPC 64**
- A new micro-kernel written from scratch
 - Built for user needs, i.e. small footprint and adapted policies
 - Multi-core architecture based on an asymmetric kernel
 - Real-time model for user applications
 - MMU support and virtual address space
 - Resource sharing & IPC protocols (mutexes, semaphores, message passing, etc.)
 - Usual OS services (timers, etc.)
- The OS is highly configurable

Real-time

- User processes have real-time requirements
- Determinism and bounded guarantees
- On time as opposed to fast
- Real-time scheduling policies
- Resource usage bounded and checked

New micro-kernel

New micro-kernel

- No “Linux legacy” or other previous mono-core design

New micro-kernel

- No “Linux legacy” or other previous mono-core design
- Design for SMP platforms

New micro-kernel

- No “Linux legacy” or other previous mono-core design
- Design for SMP platforms
- Asymmetric kernel design
 - One core for heavy scheduling operations
 - Other cores working to service tasks
 - Remote system call mechanism

New micro-kernel

- No “Linux legacy” or other previous mono-core design
- Design for SMP platforms
- Asymmetric kernel design
 - One core for heavy scheduling operations
 - Other cores working to service tasks
 - Remote system call mechanism
- Most services & drivers in user space

New micro-kernel



- No “Linux legacy” or other previous mono-core design
- Design for SMP platforms
- Asymmetric kernel design
 - One core for heavy scheduling operations
 - Other cores working to service tasks
 - Remote system call mechanism
- Most services & drivers in user space
- Multi-core IPC protocol to handle it

New micro-kernel



- No “Linux legacy” or other previous mono-core design
- Design for SMP platforms
- Asymmetric kernel design
 - One core for heavy scheduling operations
 - Other cores working to service tasks
 - Remote system call mechanism
- Most services & drivers in user space
- Multi-core IPC protocol to handle it

NB: we also support FPGA :-)

OS Modules


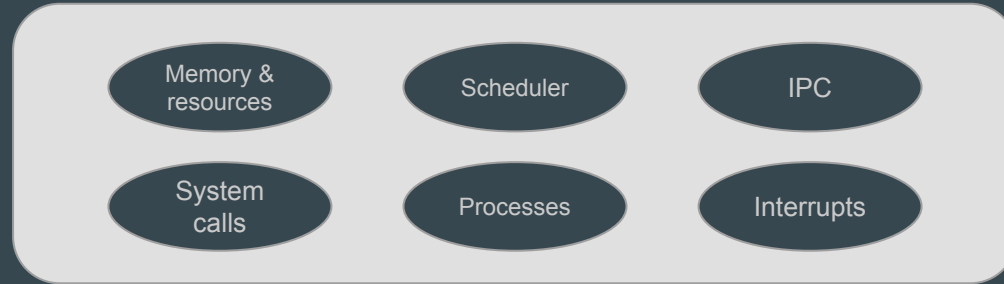


HARDWARE



OS Modules

KERNEL
SPACE


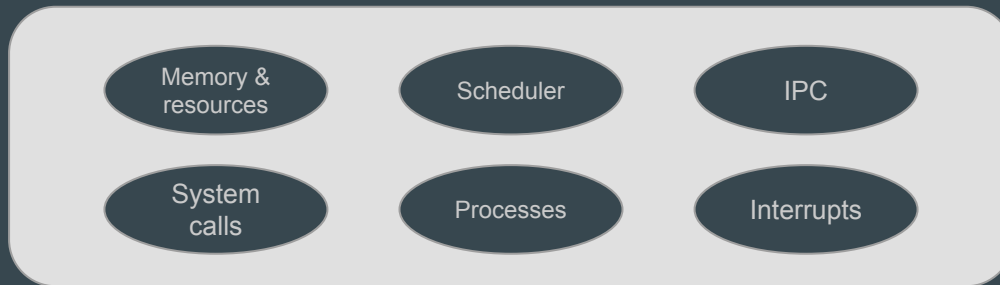


OS Modules

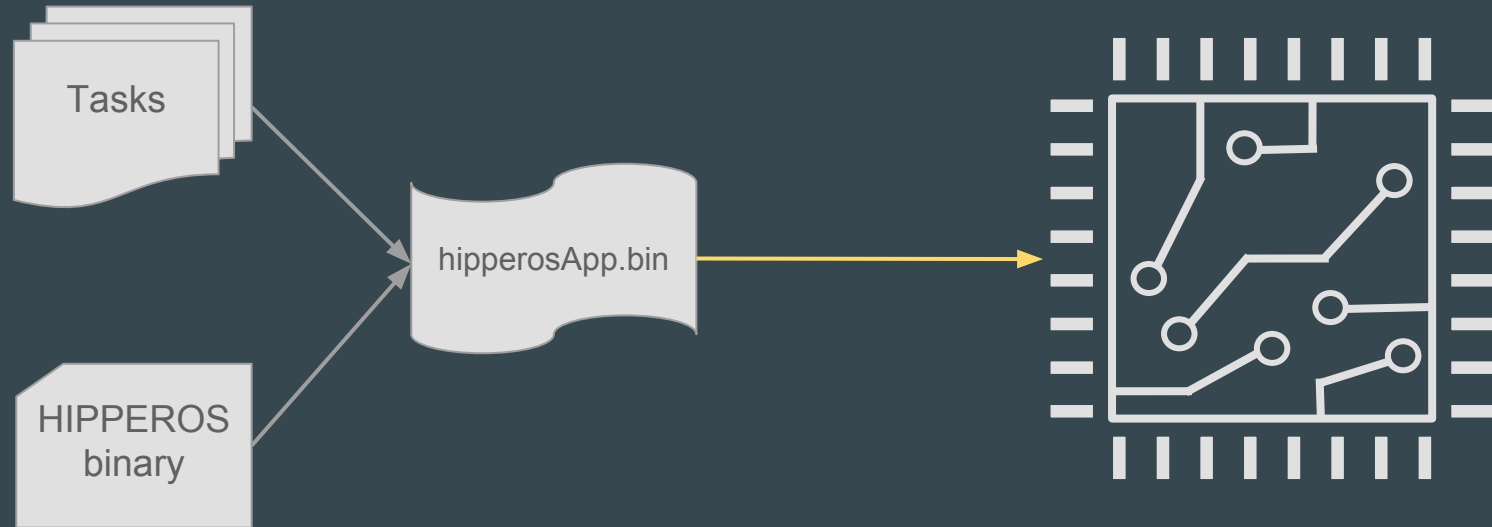
USER
SPACE



KERNEL
SPACE



In practice: build an application and deploy it on target



Task set defines real-time behaviour and specification

- Timing parameters
- Periodicity
- Code
- Core affinities
- ...

```
<task>
  <identifier>1</identifier>
  <name>task1</name>
  <stackSize>8192</stackSize>
  <recurrence>UNIQUE</recurrence>
  <entryPoint>task1_main</entryPoint>
  <coreAffinity>0 1</coreAffinity>
</task>
<task>
  <identifier>2</identifier>
  <name>task2</name>
  <stackSize>8192</stackSize>
  <recurrence>PERIODIC</recurrence>
  <timingInformation>
    <wcet>42000</wcet>
    <deadline>20000</deadline>
    <period>20000</period>
  </timingInformation>
  <flags>REALTIME</flags>
  <coreAffinity>0</coreAffinity>
</task>
<task>
  <identifier>3</identifier>
  <name>task3</name>
  <stackSize>8192</stackSize>
  <recurrence>SPORADIC</recurrence>
  <timingInformation>
    <offset>10000</offset>
    <wcet>57000</wcet>
    <deadline>120000</deadline>
    <period>120000</period>
  </timingInformation>
  <flags>REALTIME</flags>
  <coreAffinity>1</coreAffinity>
</task>
```

For more information

- Seminal paper: OSPERT 15
- We can work together
 - HIPPEROS Academic Partner Program
 - academic@hipperos.com
- Use HIPPEROS for commercial application
 - contact us: info@hipperos.com
- You will soon be able to play with it for free!
 - HIPPEROS community edition
 - Expected release date: mid 2018
- For any information, contact me: antonio.paolillo@hipperos.com



Mixed-criticality operating system?

Mixed-criticality operating system?


Vestal model

Mixed-criticality operating system?


Vestal model*

* actually, the elastic task model


3 types of tasks




3 types of tasks



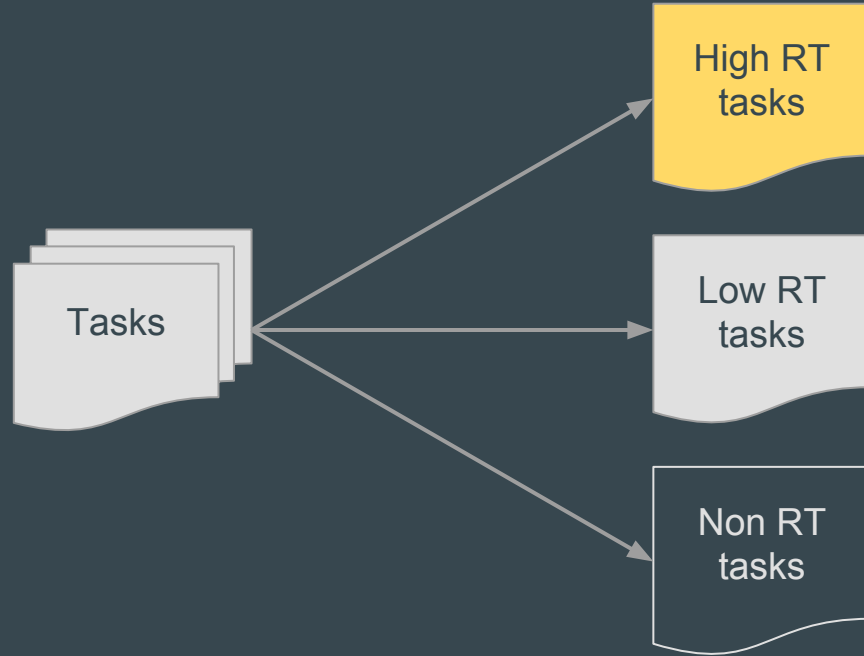
3 types of tasks



3 types of tasks

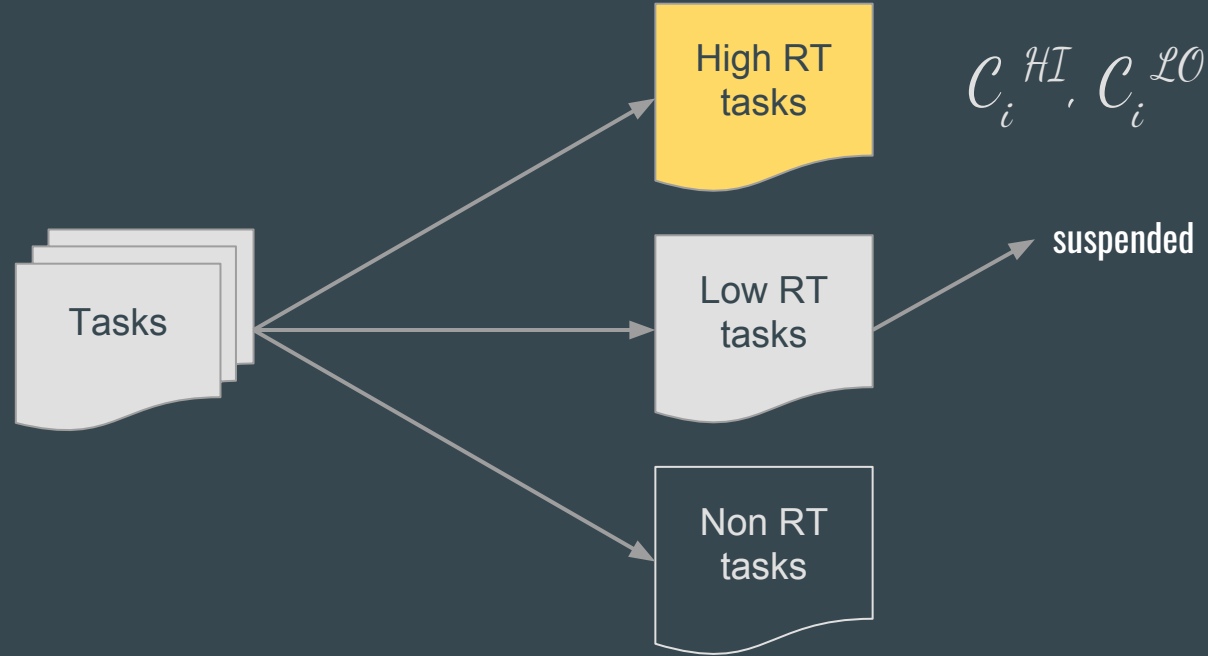


3 types of tasks

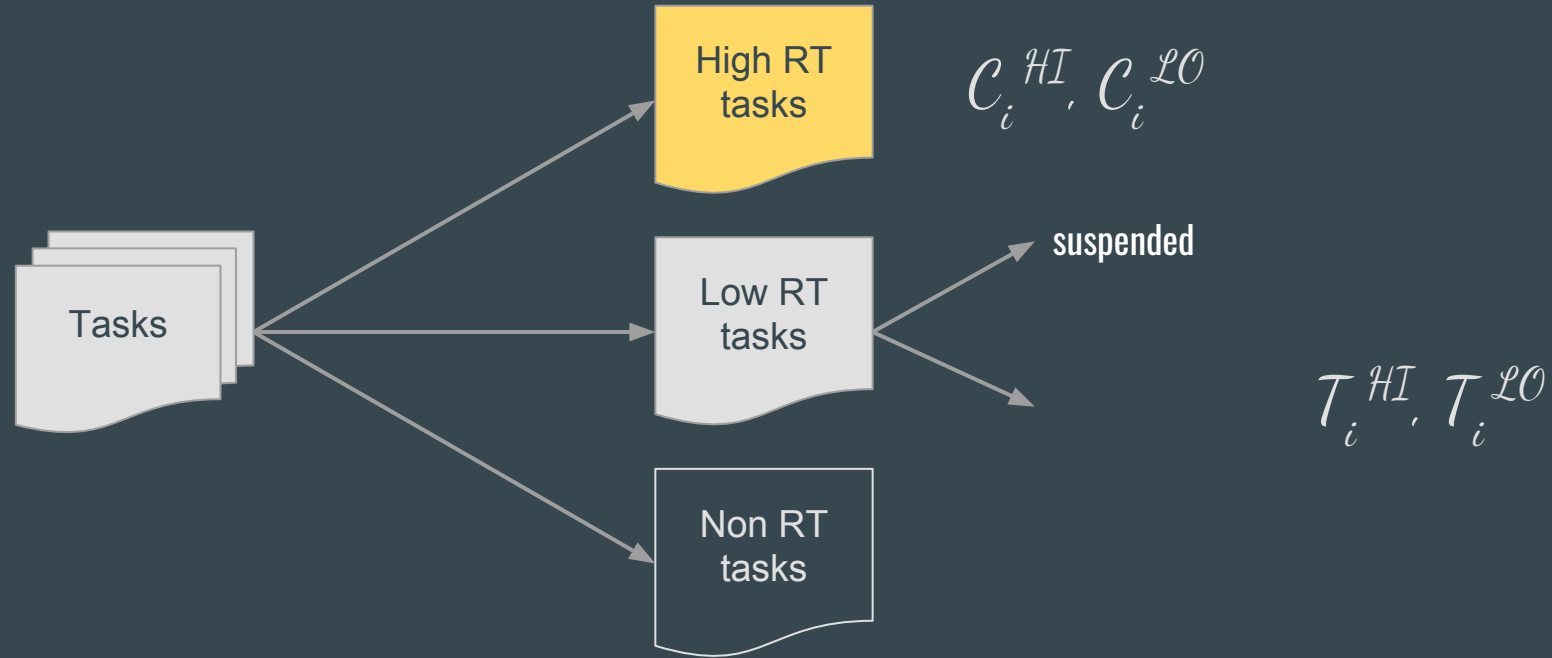


$$C_i^{HI}, C_i^{LO}$$

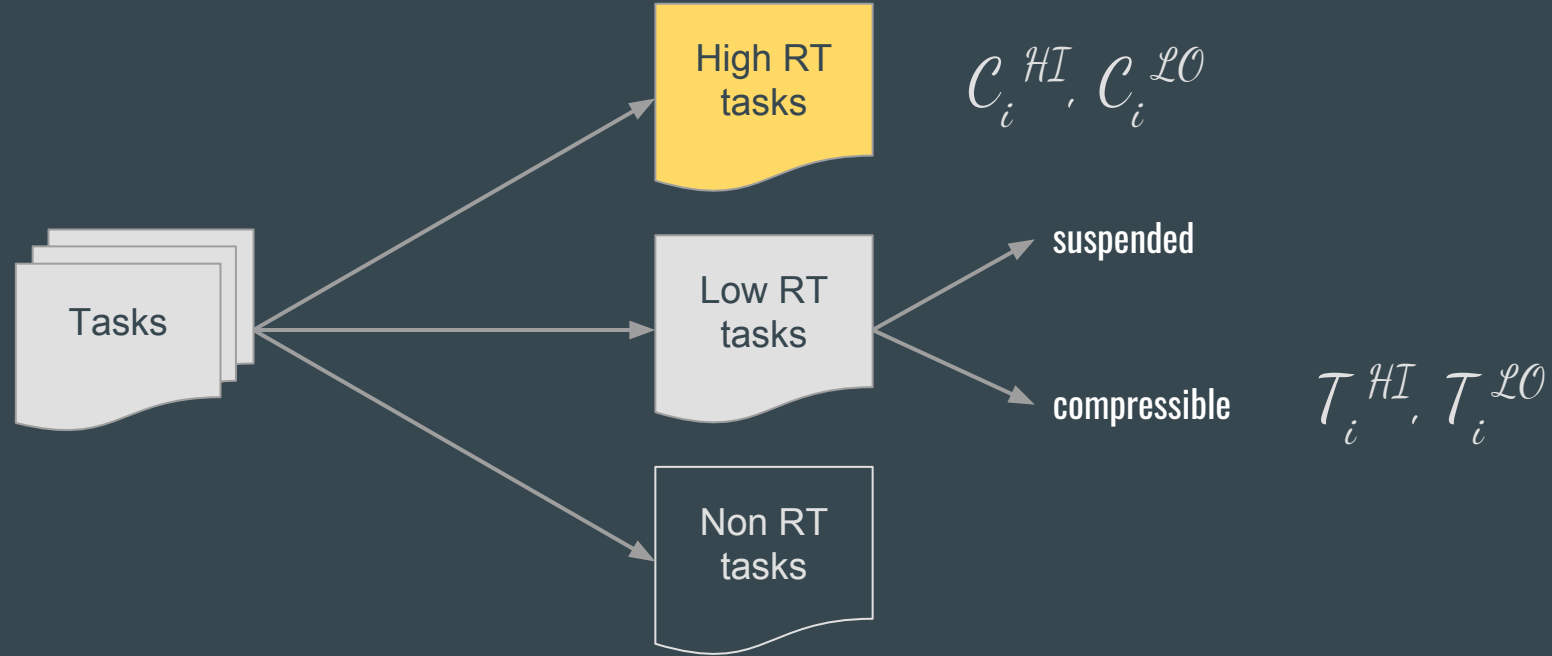
3 types of tasks



3 types of tasks



3 types of tasks



Extend the task set with Mixed-Criticality

```
<task>  
  <timingInformation>  
    <offset>0</offset>  
    <wcet>100000</wcet>  
    <deadline>200000</deadline>  
    <period>500000</period>  
  </timingInformation>  
</task>
```

Extend the task set with Mixed-Criticality

```
<task>
  <timingInformation>
    <offset>0</offset>
    <wcet>100000</wcet>
    <deadline>200000</deadline>
    <period>500000</period>
  </timingInformation>
</task>
```

$\{O_i, C_i, D_i, T_i\}$

Extend the task set with Mixed-Criticality

```
<task>
  <timingInformation>
    <offset>0</offset>
    <wcet>100000</wcet>
    <deadline>200000</deadline>
    <period>500000</period>
    <mcHigh>
      <wcetLow>25000</wcetLow>
    </mcHigh>
  </timingInformation>
</task>
```

$$\{O_i, C_i^{HI}, D_i, T_i\}$$
$$C_i^{LO}$$

Extend the task set with Mixed-Criticality

```
<task>
```

```
  <timingInformation>
```

```
    <offset>0</offset>
```

```
    <wcet>100000</wcet>
```

```
    <deadline>200000</deadline>
```

```
    <period>500000</period>
```

```
    <mcLow>
```

```
      <periodHigh>1000000</periodHigh>
```

```
      <onModeSwitch>LET_FINISH</onModeSwitch>
```

```
    </mcLow>
```

```
  </timingInformation>
```

```
</task>
```

$\{O_i, C_i, D_i, T_i^{LO}\}$

T_i^{HI}

Mode switch event - a low WCET overrun

Low tasks job instance:

```
<onModeSwitch>  
  LET_FINISH  
</onModeSwitch>
```

OR

```
<onModeSwitch>  
  KILL  
</onModeSwitch>
```

Tasks:

```
<periodHigh>  
  1000000  
</periodHigh>
```

OR

```
<suspended/>
```

OR

```
<unaffected/>
```


HI task LO WCET overrun?

HI task LO WCET overrun?

Global mode switch
(all jobs of all cores)

Switch back to LO mode?

Switch back to LO mode?

First idle instant

C. Evaluated platforms

NXP T2080RDB

- QorIQ T2080 platform:
 - 4 dual-threaded e6500 cores (1.8 GHz)
 - PowerPC 64 bits architecture
- 4 GB RAM
- OS support for caches:
 - L1
 - Partitioning of L2 (\approx private)
 - No L3
 - TLB miss software handler



BD SABRE Lite

- NXP i.MX 6Quad processor:
 - 4 Cortex-A9 cores (800 MHz)
 - ARMv7-A 32 bits architecture
- 1 GB RAM
- OS support for caches:
 - Private L1 enabled
 - No L2
- Not in IMICRASAR but “control board”



D. Experiments

What we evaluated

What we evaluated

- Idea: how to set the LO wcet such that we balance safety and efficiency?

What we evaluated

- Idea: how to set the LO wcet such that we balance safety and efficiency?
- Experiment execution time measurements → approximate HI WCET bounds
 - HI WCET: 4x max observed execution time for each task

What we evaluated

- Idea: how to set the LO wcet such that we balance safety and efficiency?
- Experiment execution time measurements → approximate HI WCET bounds
 - HI WCET: 4x max observed execution time for each task
- Run the Thales HI application concurrently with a dummy LO application

What we evaluated

- Idea: how to set the LO wcet such that we balance safety and efficiency?
- Experiment execution time measurements → approximate HI WCET bounds
 - HI WCET: 4x max observed execution time for each task
- Run the Thales HI application concurrently with a dummy LO application
 - Add dummy LO tasks and measure contributions
 - CPU bound, no I/O
 - Suspended when a mode switch occurs
 - No deadline miss for HI tasks (avoid HI/LO interferences)

What we evaluated

- Idea: how to set the LO wcet such that we balance safety and efficiency?
- Experiment execution time measurements → approximate HI WCET bounds
 - HI WCET: 4x max observed execution time for each task
- Run the Thales HI application concurrently with a dummy LO application
 - Add dummy LO tasks and measure contributions
 - CPU bound, no I/O
 - Suspended when a mode switch occurs
 - No deadline miss for HI tasks (avoid HI/LO interferences)
- Limitations:
 - partitioned fixed priority scheduling
 - no specific MC scheduler is implemented (future work)

Summary: evaluation scheme

- A. Measure job execution time to bound WCET

- B. Run the HI use case with a LO application

Summary: evaluation scheme

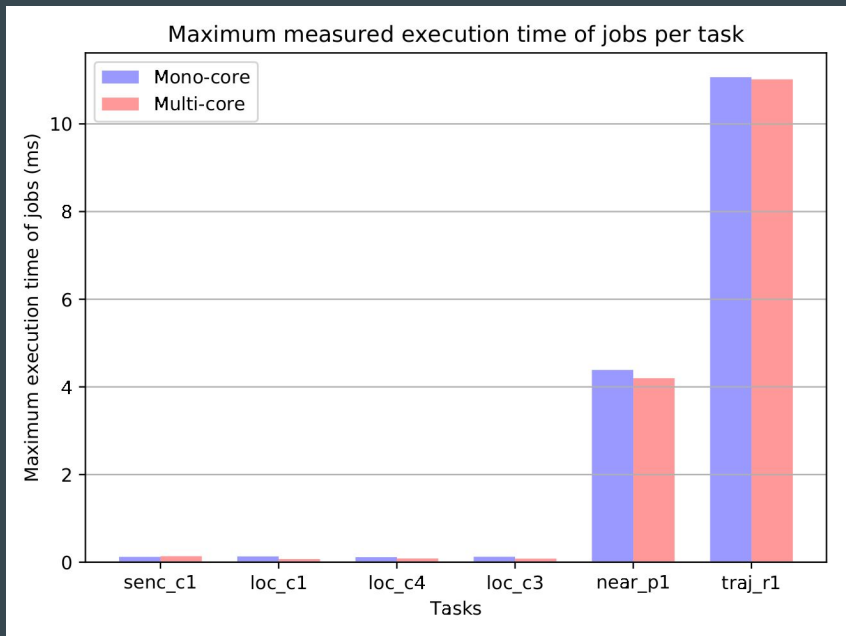
- A. Measure job execution time to bound WCET
- B. Run the HI use case with a LO application

Job execution time measurements

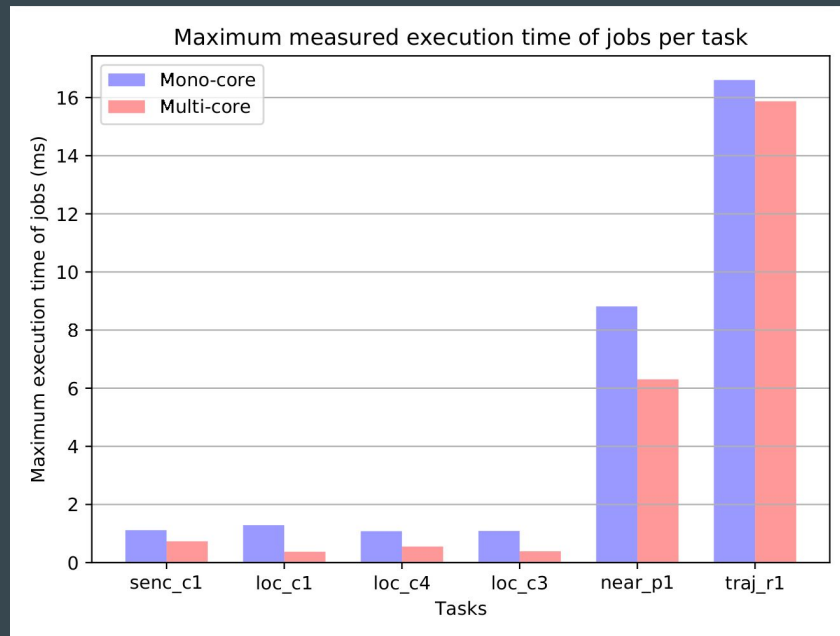
- Thales FMS application: 18 tasks
- 12 tasks have negligible execution time and/or are not periodic
- We plot the maximum observed execution time of all the jobs of each task
- Repeated 10 times, very stable measurements (small stdev)

Job execution time measurements (max observed)

T2080RDB



SABRE Lite



Observation

**Use case uses less than 10%
of the platform CPU**

Summary: evaluation scheme

- A. Measure job execution time to bound WCET

- B. Run the HI use case with a LO application

How to set WCET LO?

How to set WCET LO?




How to set WCET LO?




How to set WCET LO?



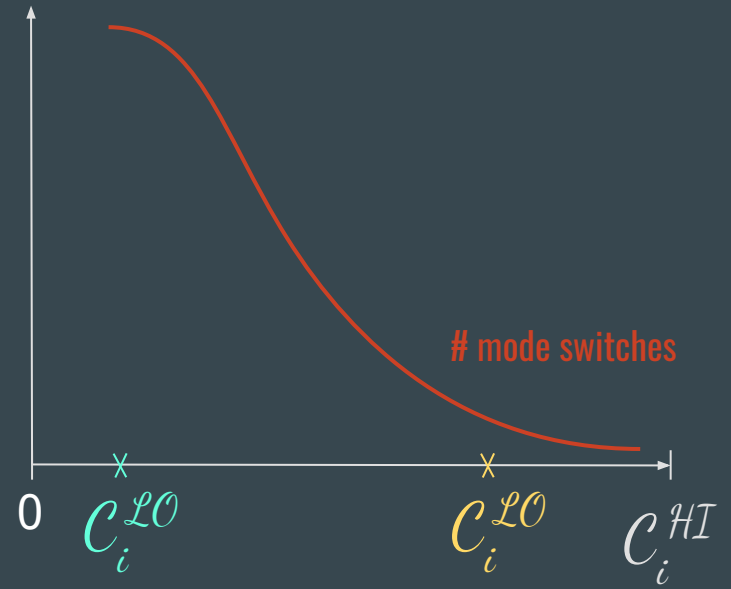
How to set WCET LO?



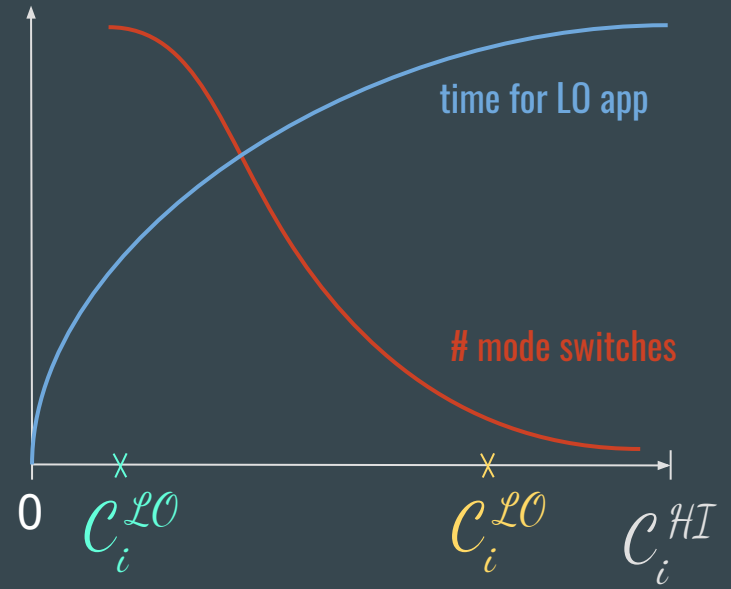
How to set WCET LO?



How to set WCET LO?

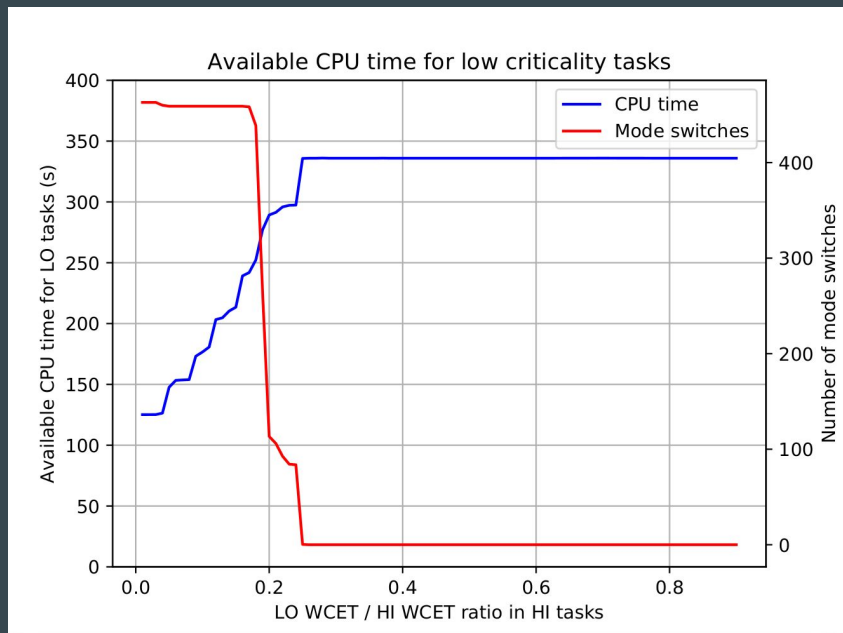


How to set WCET LO?




Evolution of available LO CPU time and # mode switches

T2080RDB





SABRE Lite



Evolution of available LO CPU time, per core

T2080RDB

SABRE Lite




Conclusions

Conclusions

- About the porting : no line of the original core application were modified
- Industrial use case takes less than 10% on the considered platform
- We can use the remaining 90% for LO app without compromising HI app
- Straightforward MC mode switch support at OS level
- Future work involves:
 - Fine grained control over interferences
 - Implementation of MC scheduling algorithms
 - Evaluations: non trivial LO app (memory, file systems, I/O) with compressible/incompressible
 - Typical industrial use cases

Where is the balance?



Where is the balance?

