ESA ITT 6185 - System Impact of Distributed Multicore systems

XtratuM porting to LEON4

Alfons Crespo, Miguel Masmano
UPV
Mathieu Patte, Vincent Leftz
Astrium

UNIVERSITAT POLITÈCNICA DE VALÈNCIA

ASTRIUM
AN EADS COMPANY

Index

- Approach
- Main issues
- Scheduling policies
- Configuration file

Approach

- XtratuM was initially designed for monocore architectures: LEON2 and LEON3
- It offers virtual machines (vCPU) to execute partitions

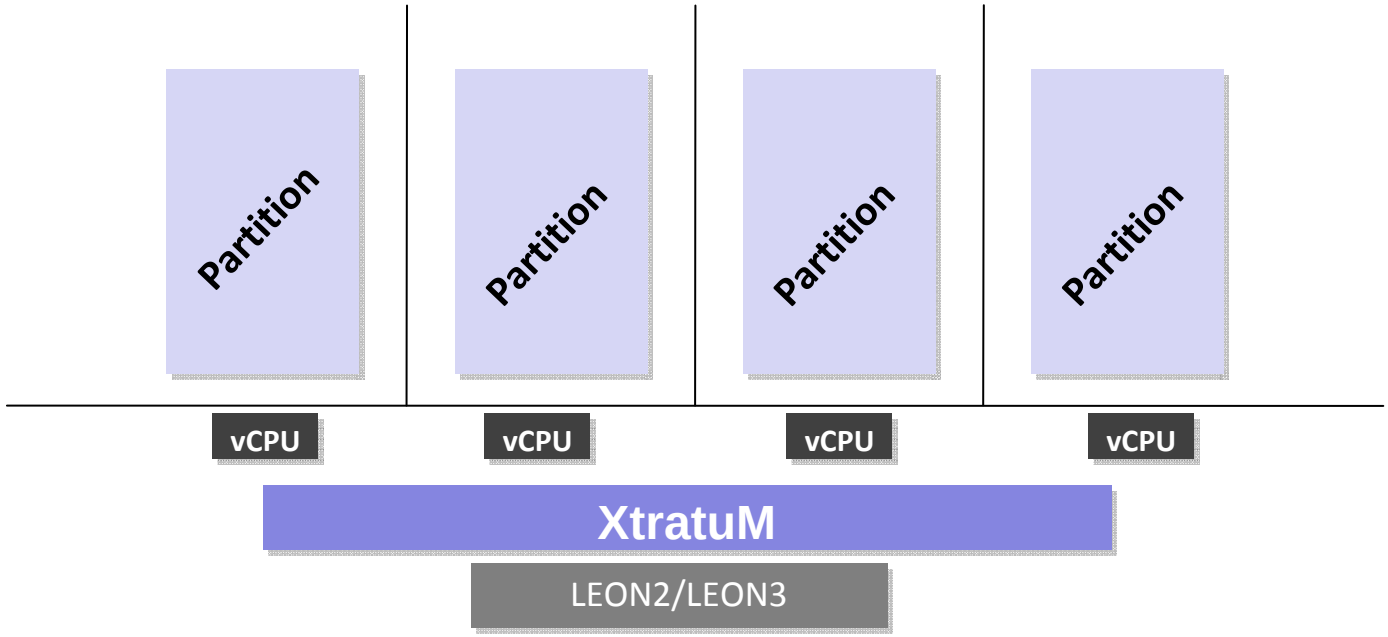

Approach

- In a multicore architecture the scheme can be:
  - Implicit
  - Explicit
- Partitions can be:
  - Monocore 1 vCPU
  - Multicore N vCPUs

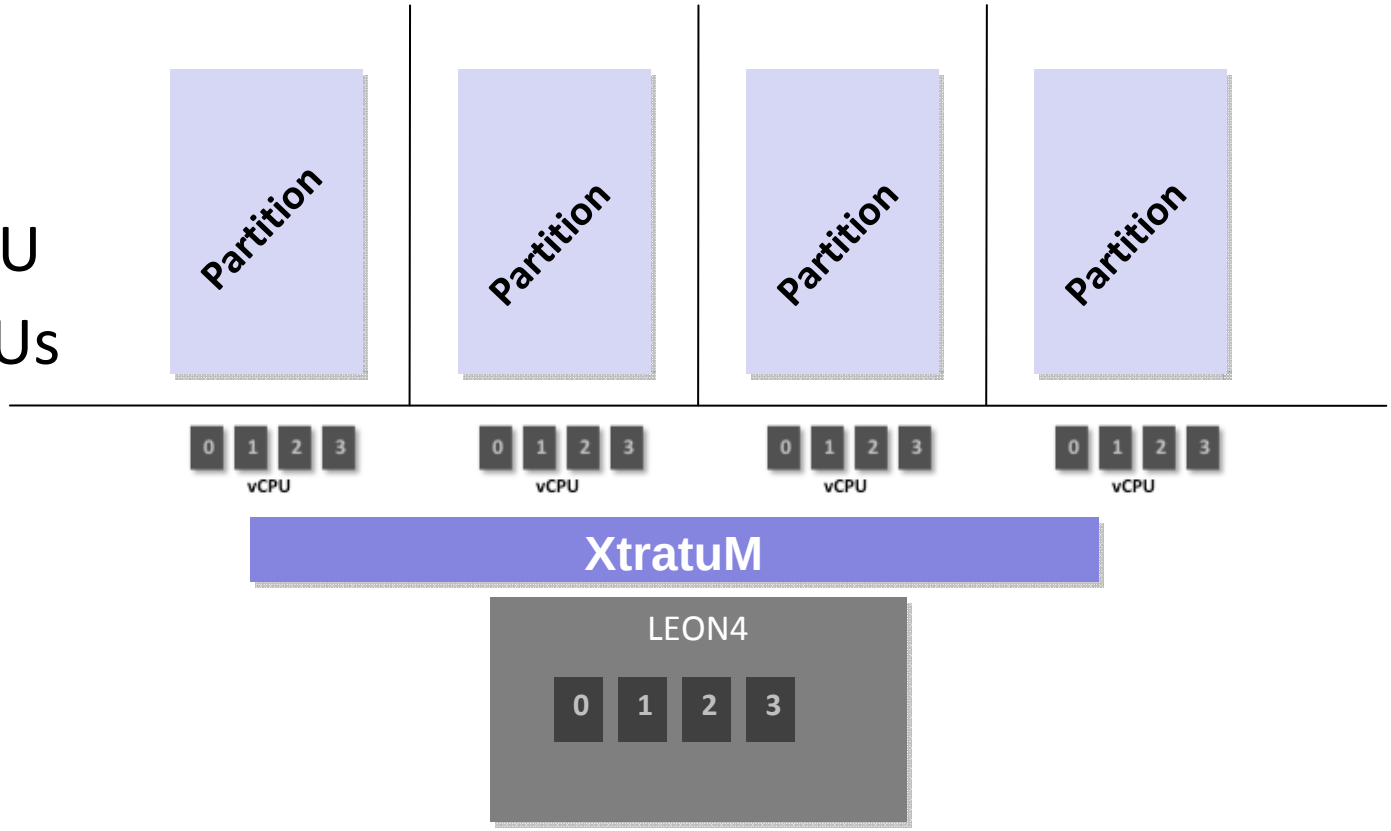

Issues related to Multicore

- Impact of Multicore on the services provided by the hypervisor:
  - Interrupts; Partition management; Health Monitor;
- Virtualised resources
  - Clock and Timers
  - Interrupt management (Set up/use the multiprocessor interrupt controller with extended ASMP)
  - IPI's management (Emulate IPIS through interrupts)
  - Memory management
- Scheduling
  - Main aspect to be analysed

Issues related to Multicore

- Virtual CPUs:
  - Inclusion of the virtual CPU (VCPU) concept
  - Each partition has one or more VCPUs (multi-core)
  - Each VCPU has a local partition control table
  - The clock is shared among the VCPUs
- New hypercalls are required:
  - get_VCPUID_self
  - Start-up/resume/suspend/halt VCPU
- XML extension
  - Each partition shall define the number of VCPUs supported (omission means 1)
  - Each slot shall indicate the VCPUID (omission means VCPUID=0)

Scheduling

- Several scheduling policies
  - Basic scheduling policy: **Cyclic scheduling**
  - Alternative scheduling policies (IO activities)
    - Fixed Priority Scheduling
    - Limited preemptive Priority Scheduling
    - Deferrable/Sporadic Server
- Each core can have different scheduling policy
  - i.e.:
    - 3 cores under a cyclic scheduling
    - 1 core other policy

Plan management in Multicore

- **Definition of Plans**
  - **MAF definition**
  - **Multiple schedules**

Single Core Scheduling Plan

In the single core approach, ARINC-653 defines the scheduling policy as a **cyclic scheduling for partitions**.

In ARINC-653 extended services, it proposes a **Multiple schedule** scheme to deal with modes of operation.

Example of a schedule with 3 plans.

- Each plan is the system architect response to a mode of operation (i.e. initialisation, normal, maintenance, etc.)
- Each plan has the appropriated MAF (least common multiple of all partition/task periods)
- A plan change can be requested only by *system* partitions
- Plan changes are effective at the end of the MAF

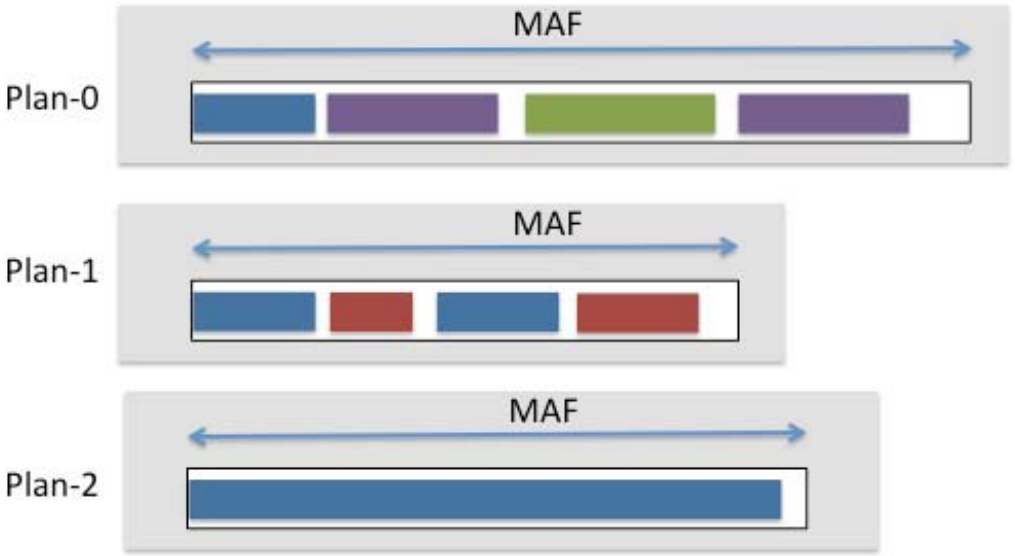

Multicore Scheduling Plan

- Each Schedule Plan defines the set of partitions to be executed in each **core**
- Each **core** defines a policy to be used to schedule partitions

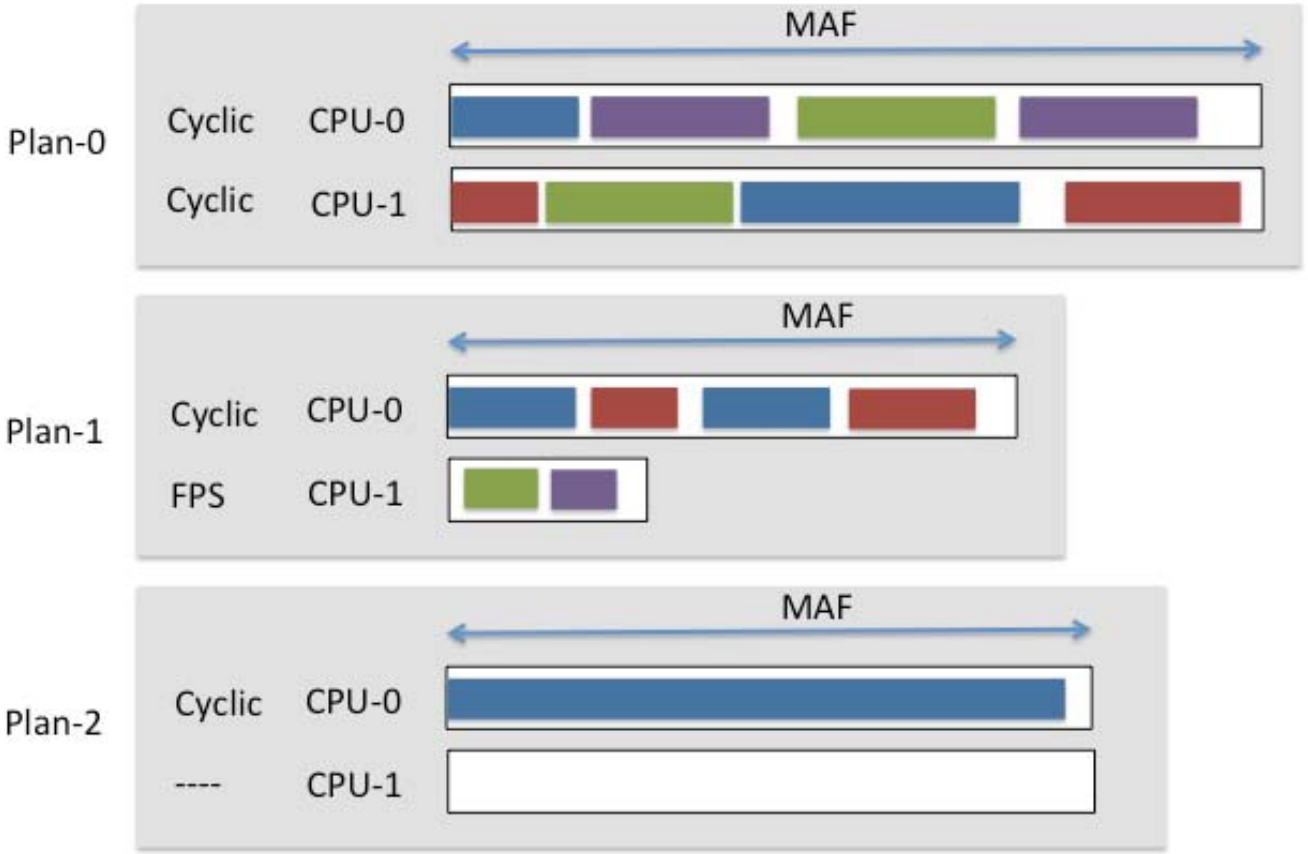

Multicore Scheduling Plan

- Plan specification: configuration file

Plan view

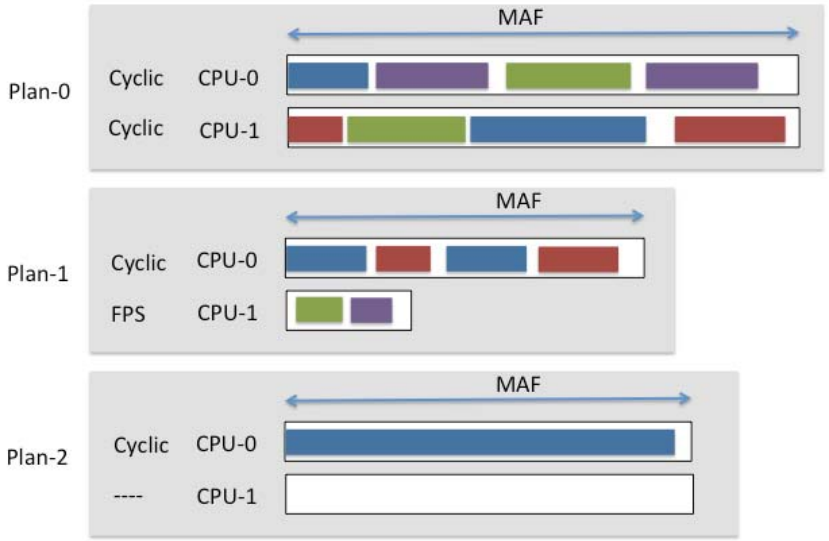

Core view

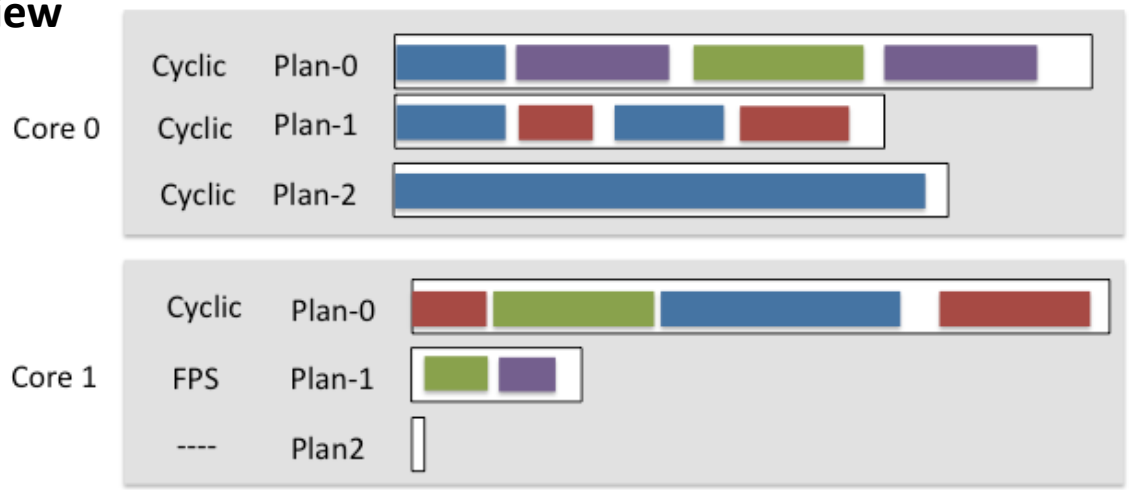

Multicore Scheduling Plan

- Plan specification: configuration file

```
<Processor id="0" frequency="50Mhz">
    <CyclicPlanTable>
     <Plan id="0" majorFrame="400ms">
      <Slot id="0" start="0ms" duration="200ms" partitionId="0" vCpuId="0"/>
            <Slot id="1" start="200ms" duration="200ms" partitionId="0" vCpuId="1"/>
     </Plan>
     ..............
    </CyclicPlanTable>
 </Processor>

   <Processor id="1" frequency="50Mhz">
     --------
     <FixedPriority>
            <Partition id="0" vCpuId="1" priority="10"/>
            <Partition id="2" vCpuId="0" priority="5"/>
     </FixedPriority>
    </Processor>

<PartitionTable>
        <Partition id="0" name="Partition1" flags="system" console="Uart" noVCpus="4">
...
```