

Checkpointing in a real time OS

David Garriou

February 1, 2020





- 1 Introduction
- 2 Context of the project
- 3 Energy estimation and prediction
- 4 Trampoline RTOS
- 5 Transcient computing platform

- Smile
 - 1800 geeks
 - European expert in open source solutions
 - www.smile.eu
 - Smile did fund a part of the work
- LS2N – UMR CNRS 6004 (French National Centre for Scientific Research)
 - École Centrale de Nantes, IMT Atlantique, Université de Nantes
 - www.ls2n.fr
 - Special thanks to my friends and former colleagues, Jean-Luc Béchenec, Mikaël Briday, Sébastien Faucou



- 1 Introduction
- 2 Context of the project
- 3 Energy estimation and prediction
- 4 Trampoline RTOS
- 5 Transcient computing platform

Some facts

- So many connected objects in the coming years
- So much raw data sent to the cloud
- Too much data will be stored in the cloud

Our motivation


- Decreasing the environmental impact of IoT
- Decreasing the maintenance cost of IoT

Our Goal

- Avoid transmission of raw data
- Provide a platform for usual sensing, transmission, "heavy" computing
- Have fun of course !



- Constrained embedded devices
- Application over an operating system
- Batteryless devices, so we need hypothesis :
 - Some non volatile RAM (NVRAM)
 - A supercapacitor





- 1 Introduction
- 2 Context of the project
- 3 Energy estimation and prediction**
- 4 Trampoline RTOS
- 5 Transcient computing platform



- System equipped with a supercapacitor
- Available energy in the capacitor stated with a relation on voltage
- Available energy depends on operating mode, clocked peripherals, microcontroller operations, etc.
- Estimate the operating time (the time before a checkpoint shall be done)
- May be predicted using a linear model
- Estimate a global slope (from different modules' slope)
- Adjust the slopes dynamically at runtime



- By using the current voltage, V_c and the voltage drop slope over time, it is possible to predict when a checkpoint execution is required (When the voltage will cross V_{cp}).
- A timer is programmed so that an interruption (cp_{int}) occurs at this time.
- In addition, the exact voltage is read periodically to adjust the prediction. This can lead to a reprogramming of the timer.


Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done




Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done




Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done




Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done




Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done




Operating time prediction



V_{min} Minimum voltage at which the microcontroller runs

V_{max} Maximum voltage of the supercapacitor

V_{cp} Voltage at which a checkpointing shall be done





- 1 Introduction
- 2 Context of the project
- 3 Energy estimation and prediction
- 4 Trampoline RTOS**
- 5 Transcient computing platform

- Real-time Operating System, developed by the LS2N Real-time Systems Group
- Modular and portable
- Free Software License (GPLv2) + industrialized version
- Conforms to OSEK/VDX OS and AUTOSAR OS (fixed priority scheduling)
- Many targets including AVR, ARM-Cortex, PowerPC, Posix, MSP430
- Support for multicore architectures

Contributors 13 identified, but **4** in reality; Jean-Luc Béchenec, Mikaël Briday, Sébastien Faucou, David Garriou

Commits > 2600

Organization No organization ! But I hope things will evolve

- First Trampoline sprint (à la Buildroot) on march.
- A lot of subjects to deal with; new features, build system, code generation, legacy code, documentation
- As far as we know, used in industrial environment by two major undertaking

■ Version 1

- 2005: Start of development
- 2007: AUTOSAR
- 2008: Pass OSEK certification, AUTOSAR SC1&2
- 2009: Integrated trace system
- 2010: AUTOSAR SC3&4, MODISTARC & AUTOSAR test suite


■ Version 2

- 2013: Runtime Verification
- 2014: Multicore

■ Version 3


- 2016: Hypervisor support (XVisor)
- 2017: Runtime enforcement
- 2018: Runtime verification on FPGA

- Fixed priority scheduling
- Basic task : may not synchronize
- Extended task : may synchronize
- Critical sections use the ICPP protocol (no deadlock no priority inversion)
- Each task can be configured as preemptable or non-preemptable



The configuration of Trampoline uses a dedicated language

- OIL or arXML
- extensible (based on a template engine)






```
ALARM one_second {  
    COUNTER = SystemCounter;  
    ACTION = ACTIVATETASK { TASK = t0; };  
    AUTOSTART = TRUE { APPMODE = std; ALARMTIME = 100; CYCLETIME = 100; };  
};
```

```
TASK t0 {  
    AUTOSTART = FALSE;  
    PRIORITY = 3;  
    ACTIVATION = 1;  
    SCHEDULE = FULL;  
    MESSAGE = s00;  
};
```

```
MESSAGE s00 {  
    MESSAGEPROPERTY = SEND_STATIC_INTERNAL {  
        CDATATYPE = "uint8";  
    };  
    NOTIFICATION = NONE;  
};
```





Service	Prototype	Description
OS services		
StartOS	<i>app_mode</i>	Starts the operating system
ShutdownOS	<i>app_mode</i>	Stops the operating system
Task services		
ActivateTask	<i>task_id</i>	Activates task <i>task_id</i>
TerminateTask		Terminates the caller
Alarm services		
SetRelAlarm	<i>alarm_id, offset, cycle)</i>	Starts alarm <i>alarm_id</i>
SetAbsAlarm	<i>alarm_id, date, cycle)</i>	Starts alarm <i>alarm_id</i>
CancelAlarm	<i>alarm_id</i>	tops alarm <i>alarm_id</i>
GetAlarm	<i>alarm_id, &remaining)</i>	Gets alarm state
Resource services		
GetResource	<i>rez_id</i>	The caller enters critical section
ReleaseResource	<i>rez_id</i>	The caller leaves critical section
Event services		
WaitEvent	<i>ev</i>	The caller waits for event <i>ev</i>
SetEvent	<i>task_id, ev</i>	Set event <i>ev</i> to task <i>task_id</i>
ClearEvent	<i>ev</i>	Clear event <i>ev</i> of the caller
GetEvent	<i>task_id, &ev</i>	Puts a copy of events received by task <i>task_id</i> in <i>ev</i>



- 1 Introduction
- 2 Context of the project
- 3 Energy estimation and prediction
- 4 Trampoline RTOS
- 5 **Transcient computing platform**

Port of MSP430 to Trampoline


- Two boards; MSP430FR5969, MSP430FR5994 Launchpads
- Still in progress, not completely tested
- We need to better integrate it in Trampoline (e.g. templates)

New services


Service	Prototype	Description
OS services		
RestartOS		Starts the operating system from a checkpoint
Hibernate		Terminates the current task and records a checkpoint in an available FRAM slot



- 16 bits micro-controller
- 16 registers 20 bits wide (CPUX)
- 16MHz
- 4kB + 4kB SRAM
- 512B + 256kB FRAM
- $118\mu\text{A}/\text{MHz}$ active mode, Peripheral low-frequency (LPM3)
- 0.22F supercapacitor. Works x minutes with blink.



MSP430FR5994 memory mapping



Energy

- Remaining energy prediction strategy not implemented
- We periodically read the voltage
- Commit a checkpoint if voltage value is below a threshold

Checkpoints

- Two checkpoint slots, a mark selects the available checkpoint slot
- Copy a part of RAM to a slot in FRAM
- Exclude the system stack

Restart

- When energy comes back
- Get back from FRAM the last committed checkpoint



- Implement the strategy for operating time estimation, dynamically adjust the slopes
- Do not store unnecessary data, higher level RAM management
- Restore the state of some peripherals
- Thinking with the "normally off, instantly on" paradigm
- Integrate this paradigm in the configuration of Trampoline.
 - Checkpointing some tasks, not all
 - A task may specify the energy it would need to surely complete, a kind of worst case execution energy, the Worst Case Voltage Drop.



Trampoline `https://github.com/TrampolineRTOS/trampoline`

Work Work in progress on branch `checkpointing`

Documentation Available on `checkpointing` branch,
`documentation/manual/msp430/`