technische universität dortmund

computer science 12

Bachelor/Master Thesis

Normally-off Real-Time Operating System with Non-Volatile Main Memory

Dr.-Ing Kuan-Hsun Chen
Prof. Dr. Jian-Jia Chen
Otto-Hahn Str. 16
Technische Universität Dortmund
Email: kuan-hsun.chen@tu-dortmund.de
January 10, 2021

Emerging byte-addressable non-volatile memories (NVM) have been considered to serve as main memory, since they feature low leakage power, high density, and low unit costs. However, coins have two sides. Such NVMs often face a lot of different challenges, e.g., significantly lower write endurance compared to the classic DRAM. However, one feature is of key interest in embedded system designs, that they allow a computing system to turn off the main memory without the need of storing and restoring any data in the main memory. This feature allows the system to apply advanced dynamic power management (DPM) to switch to low power states (hibernation) almost immediately to achieve energy saving.

For certain normally-off embedded systems which spend most of their lifetime idle, it is crucial to reduce the energy consumption and thus to extend, for instance, battery lifetimes in a resource constrained environment. However, such embedded systems may still serve real-time sensitive applications, where the execution of the system and the internal tasks has to meet deadlines. When the main memory is shared by all processors, the memory cannot be simply turned off as long as at least one processor still executes. Thus, to enter the hibernation state with disabled memories, all processors have to idle at the same time, i.e., no task is executing on the processors. Recently our group developed an approach to guide the scheduler when to force all processors to be idle so that the NVM can be turned off while satisfying the hard real-time constraints.

As most of the NVM relevant research is suffering, it is still lacking a suitable real platform equipped with NVM to conduct a case study. Motivated by one recent work [1] in 2019, a Japanese group develops a FPGA-based emulation environment which features two DRAM slots to emulate the system with hybrid memory models (NVM and DRAM). However, using such an environment to evaluate a real-time system is not a trivial task. A RTOS must be deployed on top of it to support new scheduling policies for evaluations.

In this thesis, students should first get familiar with the targeted RTOS, e.g. RTEMS [2], and platform, e.g. Xilinx Zynq-7000 ZC706 [3]. Afterwards, the studied RTOS should be deployed on top of the targeted platform. With the successful deployment, the students are expected to implement the proposed scheduling algorithm with reasonable mock-up examples to demonstrate the success of the integration. In case for master students, the RTOS should be enhanced to manage two different memory slots, e.g., emulate a NVM memory model with two DRAM devices. In addition, the features of the aforementioned FPGA-based emulation environment should be considered as well.

Other suggestions and related topics are also welcome. Please do not hesitate to make an appointment.

Required Skills:

- Knowledge of real-time operating systems
- Real time scheduling knowledge is beneficial
- Comfortable to program in C/C++ with Linux
- FPGA-related knowledge is beneficial

Acquired Skills after the thesis:

- Knowledge of real-time operating systems
- Experience of open source development
- Experience of system programming
- Experience of scientific writing and

[1] Yu Omori and Keiji Kimura, Performance Evaluation on NVMM Emulator Employing Fine-Grain Delay Injection, Preprint
[2] RTEMS, Github repository
[3] Xilinx Zynq-7000 ZC706, Official Website