

Non-Volatile Computing for Embedded Cyber-Physical Systems

G. Edward Suh and Edwin Kan, Cornell University

Today's computer systems are mostly built with volatile memory components such as flip-flops, SRAM, and SDRAM except for the secondary non-volatile storage. As a result, traditional computer systems quickly lose information stored in its memory if the power supply is interrupted even for a short period of time. Unfortunately, such volatile devices present significant limitations for computations on embedded cyber-physical systems that have to rely on unreliable power sources.

This project aims to develop a new computing device where non-volatile elements based on flash (floating gate) transistors are pervasively used in all levels of the memory hierarchy to enable almost instantaneous checkpointing and recovery of program state not subject to the data bus bandwidth limit. Effectively, this new system allows its power source to be cut off at any time, and yet resumes regular operation without loss of information when the power comes back. This non-volatile capability will enable continuous computations across power failures, save static power consumption when a device is idle, and enhance the security and reliability.

The initial design and evaluation of non-volatile state elements showed that the tight integration of metal nanocrystal flash memory transistors into traditional volatile memory elements such as SRAMs and flip-flops can indeed achieve almost instantaneous non-volatile checkpointing with minimal impacts on normal operations. Our non-volatile SRAM design is shown to be almost as fast as normal SRAMs (~6% slowdown) and can checkpoint the entire processor state in parallel within 10 μ s. The flash-based design also has a significant advantage on its energy requirements over other non-volatile options. To store one bit in a non-volatile fashion, the flash SRAM cell only consumes 0.075 μ W including the overhead of a charge pump compared to tens of μ W that is required for other options such as ReRAM, PCRAM, and MRAM. The wakeup process takes less than 10ns and consumes even less energy than checkpointing. We are currently working on implementing a prototype non-volatile microcontroller for fabrication.

The project has also led to new discoveries in using non-volatile and hybrid memory technologies. We found that there exist enough variations in the analog characteristics of flash memory so that one can uniquely identify and authenticate each chip that contains flash memory. Therefore, the non-volatile processors can be authenticated for security. The idea of tightly integrating multiple memory types within a cell also turned out to be more widely applicable beyond handling power failures. As an example, we studied an integration of SRAM and DRAM and showed that such a hybrid memory structure can enable a near-instant recovery from an error or significantly improve the area and energy efficiency of multi-threaded registers.