SoftMC Tutorial

Hasan Hassan

01.10.2020

Summary

- **SoftMC:** an FPGA-based DRAM testing infrastructure
- ***Characterize***, ***analyze***, and ***understand*** DRAM cell behavior
- Flexible and Easy to Use (C++ API)
- Open-source (*github.com/CMU-SAFARI/SoftMC*)
- **SoftMC enables a wide range of studies**

DRAM Organization

Memory Bus

Memory Controller

CPU

Accessing DRAM

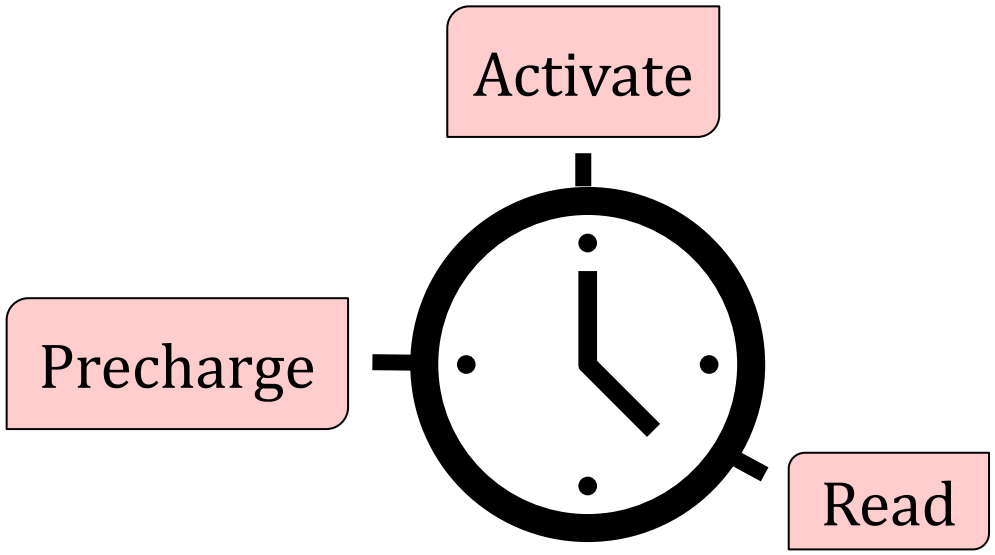

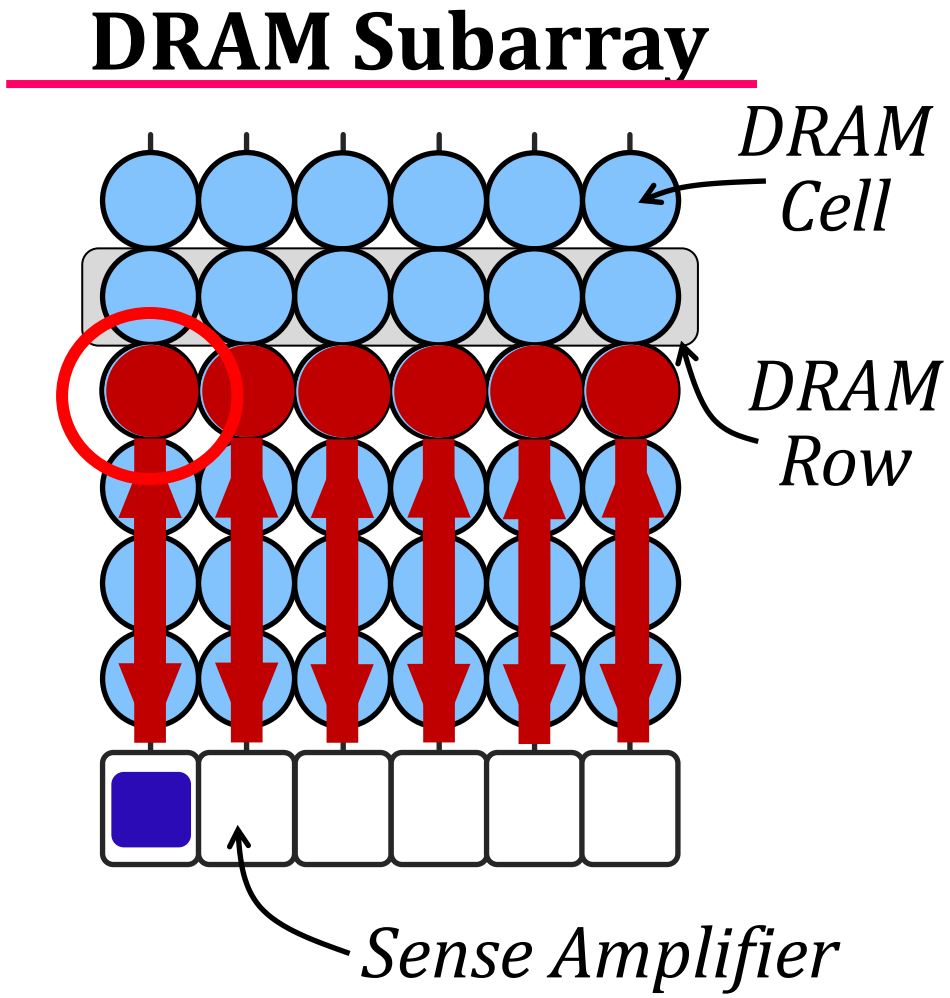

DRAM Latency

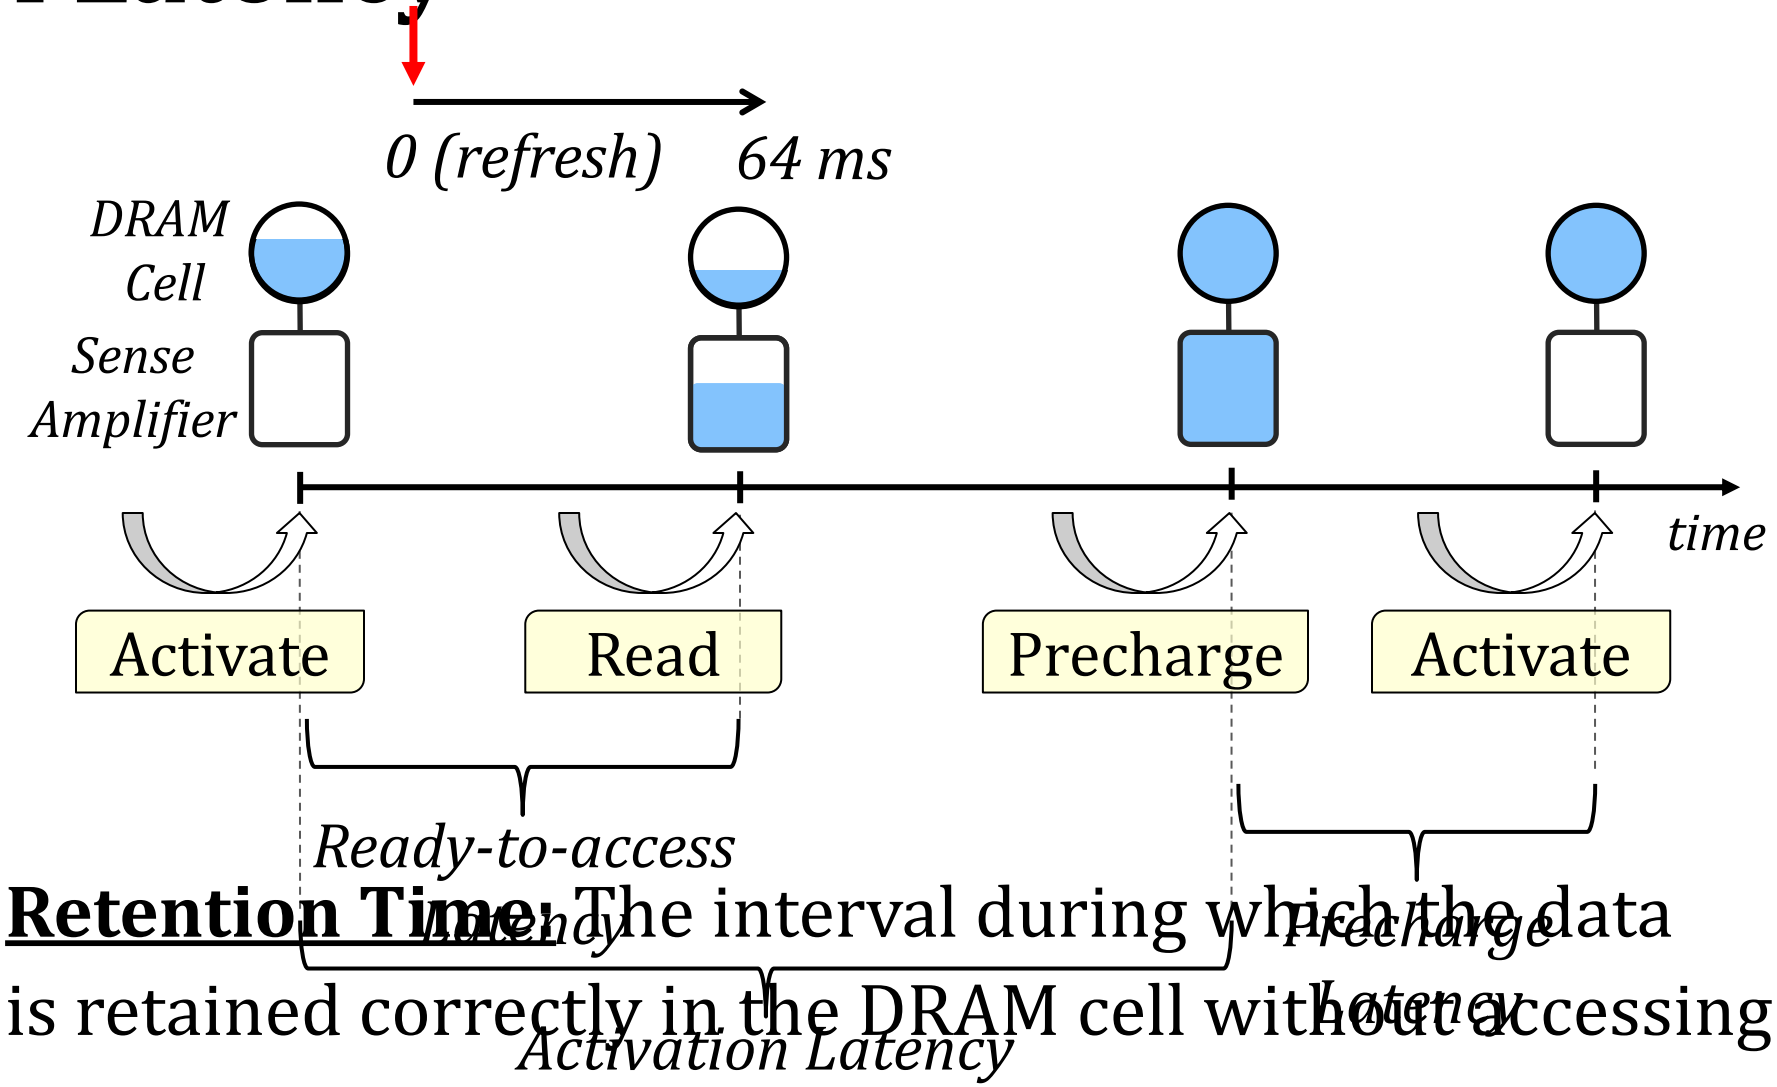

Retention Time: The interval during which the data is retained correctly in the DRAM cell without accessing it

Latency vs. Reliability

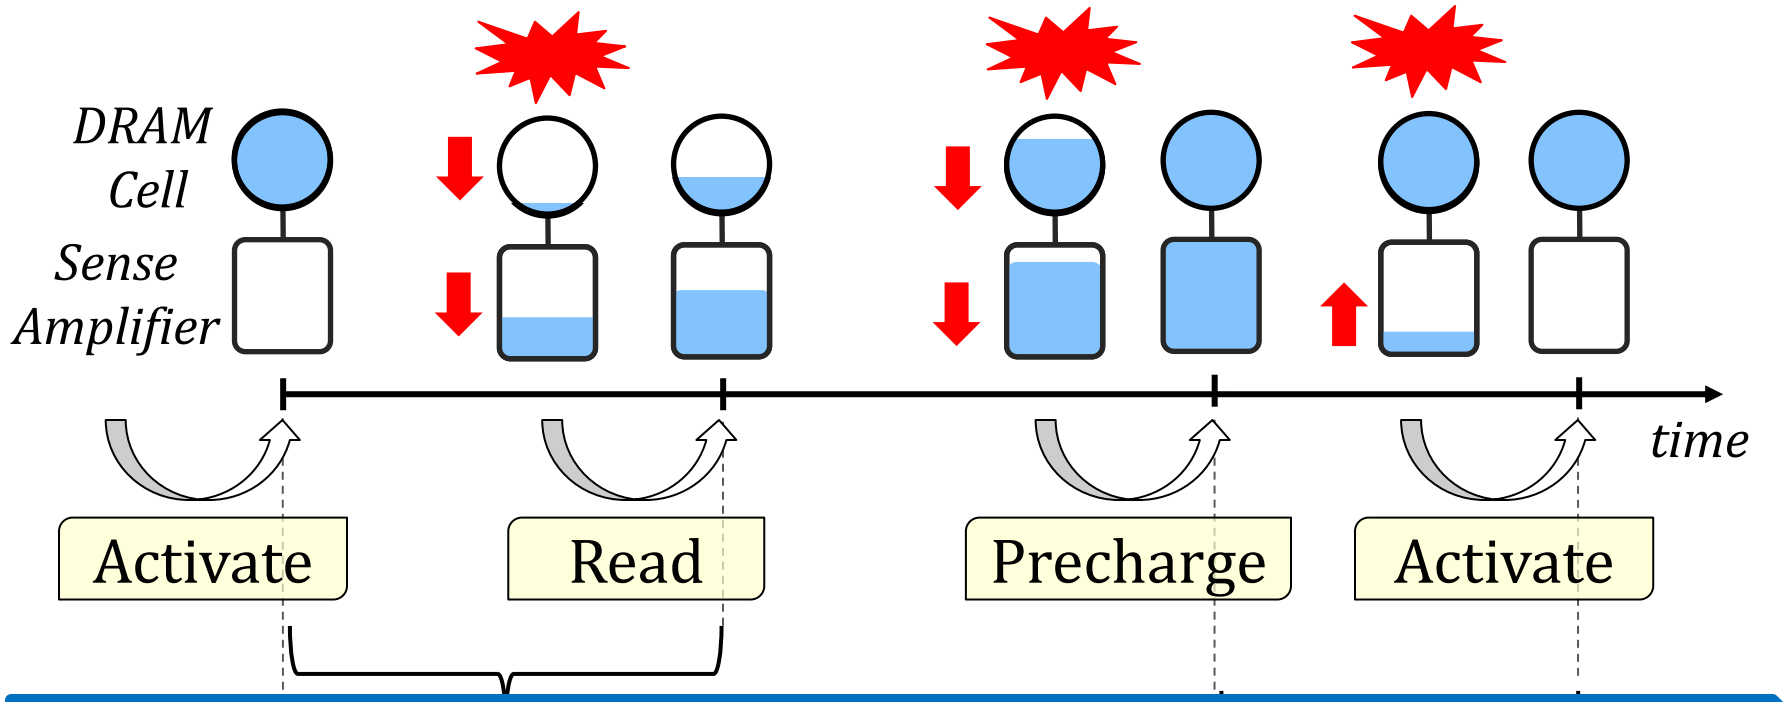

Violating latencies negatively affects DRAM reliability

Other Factors Affecting Reliability and Latency

- Temperature
- Voltage
- Inter-cell Interference
- Manufacturing Process
- Reten
- …

To develop new mechanisms improving reliability and latency, we need to better understand the effects of these factors

Characterizing DRAM

Many of the factors
affecting DRAM reliability and latency
cannot be properly modeled

We need to perform
experimental studies
of *real* DRAM chips

Prior Research Enabled by SoftMC

EDEN: Enabling Energy-Efficient, High-Performance Deep Neural Network Inference Using Approximate DRAM

Skanda Koppula Lois Orosa A. Giray Yağlıkçı
Roknoddin Azizi Taha Shahroodi Konstantinos Kanellopoulos Onur Mutlu
ETH Zürich

Understanding Reduced-Voltage Operation in Modern DRAM Chips: Characterization, Analysis, and Mechanisms

Kevin K. Chang[†] Abdullah Giray Yağlıkçı[†] Saugata Ghose[†] Aditya Agrawal[¶] Niladrish Chatterjee[¶]
Abhijith Kashyap[†] Donghyuk Lee[¶] Mike O'Connor[¶,‡] Hasan Hassan[§] Onur Mutlu[§,†]
[†]Carnegie Mellon University [¶]NVIDIA [‡]The University of Texas at Austin [§]ETH Zürich

TRRespass: Exploiting the Many Sides of Target Row Refresh

Pietro Frigo*[†] Emanuele Vannacci*[†] Hasan Hassan[§] Victor van der Veen[¶]
Onur Mutlu[§] Cristiano Giuffrida* Herbert Bos* Kaveh Razavi*
*Vrije Universiteit Amsterdam [§]ETH Zürich [¶]Qualcomm Technologies Inc.
[†]Equal contribution joint first authors

ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs

Fei Gao
feig@princeton.edu
Department of Electrical Engineering
Princeton University

Georgios Tziantzioulis
georgios.tziantzioulis@princeton.edu
Department of Electrical Engineering
Princeton University

David Wentzlaff
wentzlaf@princeton.edu
Department of Electrical Engineering
Princeton University

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim[‡§] Minesh Patel[§] Hasan Hassan[§] Lois Orosa[§] Onur Mutlu[§‡]
[‡]Carnegie Mellon University [§]ETH Zürich

Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content

Samira Khan* Chris Wilkerson[†] Zhe Wang[†] Alaa R. Alameldeen[†] Donghyuk Lee[‡] Onur Mutlu*
*University of Virginia [†]Intel Labs [‡]Nvidia Research *ETH Zürich

What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study

Saugata Ghose[†] Abdullah Giray Yağlıkçı*[†] Raghav Gupta[†] Donghyuk Lee[‡]
Kais Kudrolli[†] William X. Liu[†] Hasan Hassan* Kevin K. Chang[†]
Niladrish Chatterjee[‡] Aditya Agrawal[‡] Mike O'Connor[‡◇] Onur Mutlu*[†]
[†]Carnegie Mellon University *ETH Zürich [‡]NVIDIA Research [◇]The University of Texas at Austin

Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization

Kevin K. Chang[1] Abhijith Kashyap[1] Hasan Hassan[1,2]
Saugata Ghose[1] Kevin Hsieh[1] Donghyuk Lee[1] Tianshi Li[1,3]
Gennady Pekhimenko[1] Samira Khan[4] Onur Mutlu[5,1]
[1]Carnegie Mellon University [2]TOBB ETÜ [3]Peking University [4]University of Virginia [5]ETH Zürich

SoftMC

- FPGA-based DRAM characterization infrastructure
- Ability to test *any* DRAM operation
- Ability to test *any* combination of DRAM operations and *custom* timing parameters
- Simple programming interface (C++)

Current SoftMC Setups

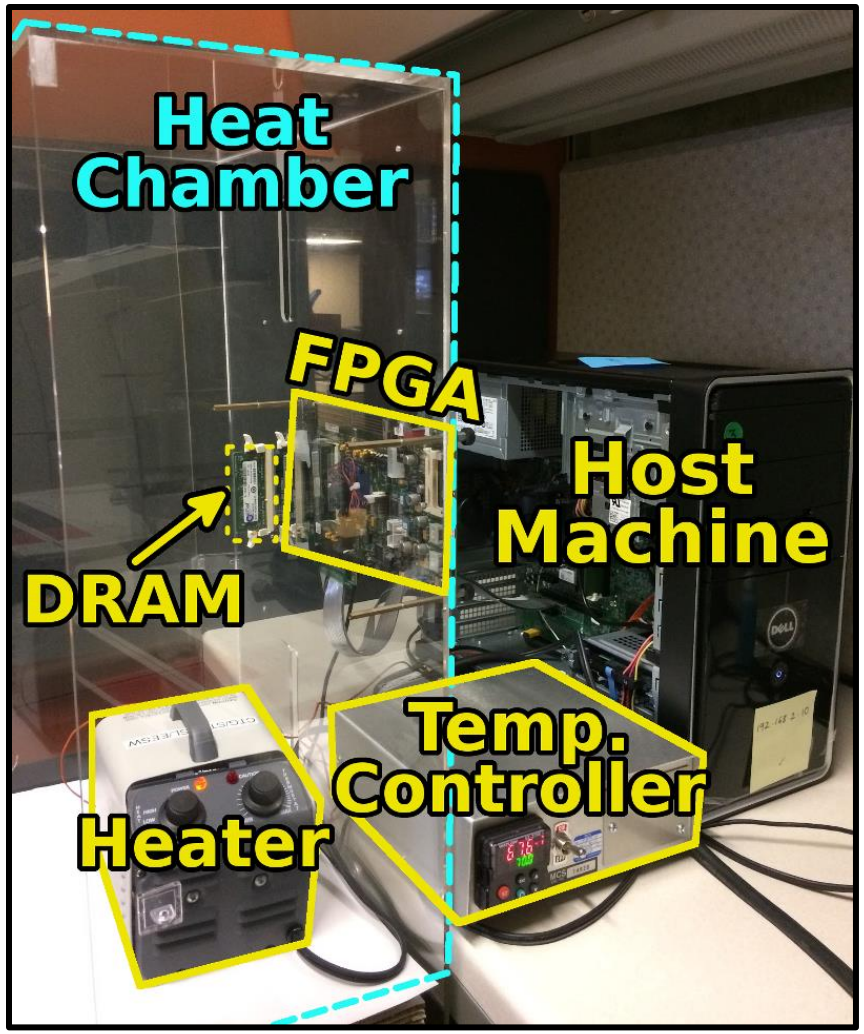

DDR3

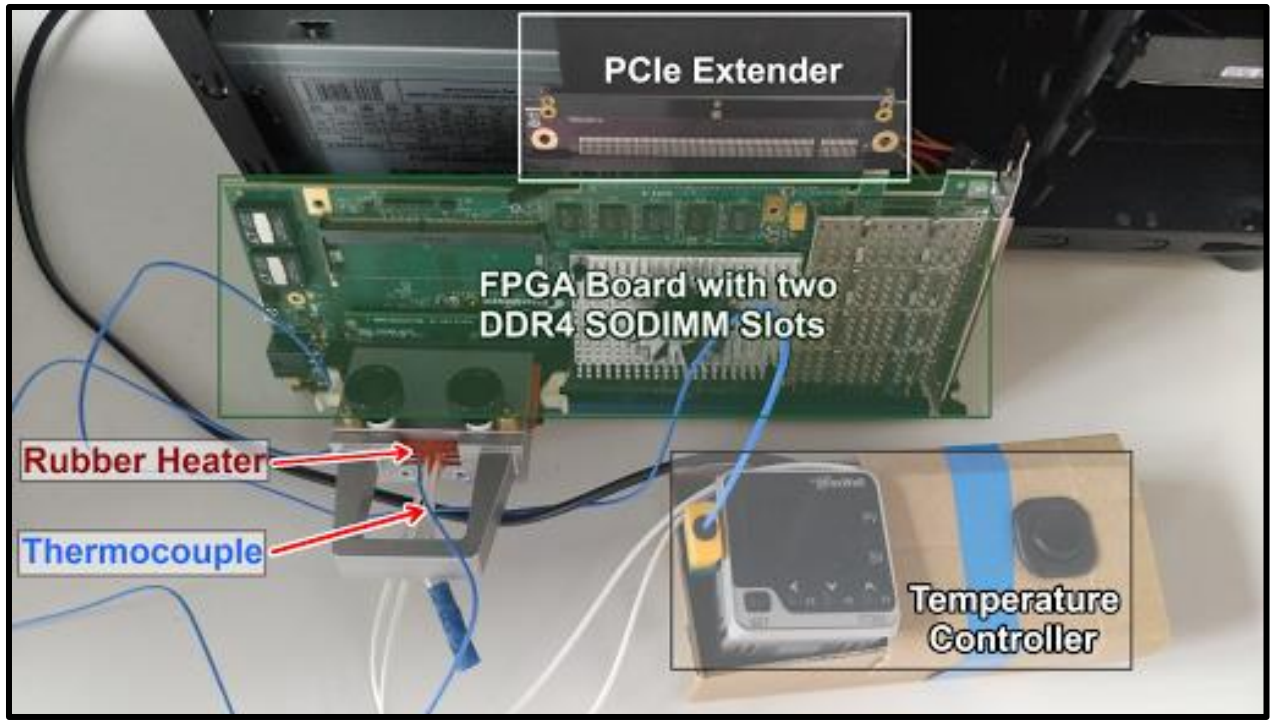

DDR4

SoftMC Programming DEMO