Qeios

Peer Review

Review of: "Calibrating DRAMPower Model: A Runtime Perspective from Real-System HPC Measurements"

Christian Weis[1]

1. Universität Kaiserslautern, Germany

The paper targets an important topic since the modeling of DRAM power based on datasheet currents is indeed very inaccurate. However, it has several major flaws, which I will point out in the following.

- First, the novelty of the paper is very limited because similar work was already carried out in the VAMPIRE paper.
- In the introduction, several factors that impact power consumption are mentioned (temperature, aging, and noise). These factors are not taken up again in the rest of the paper.
- Was the word “session” intentionally used instead of “section”?
- It is not clear how you derive the standard IDD currents from the benchmarks shown in Table 1. The IDD currents are measured for specific DRAM command traces. How can you ensure that these command traces are issued on the real system?
- For simulation, you use Intel Pin, Ramulator, and DRAMPower. But you cannot simply feed the extracted memory accesses from Pin into Ramulator. What about out-of-order execution, cache hierarchies, cache replacement policies, hardware prefetching, etc.?
- How do you know that the memory controller modeled inside Ramulator behaves like the real memory controller on the Intel CPU?
- You state the following: “A single memory command’s simulated energy in DRAMPower Model can be simplified as equation(1), where Top stands for related timing spec of the command; Nbank stands for the number of banks in the DIMM, Iop stands for the currents related to the command, and Ncmd is the total numbers of this particular command.” This is simply wrong. DRAMPower models DRAM energy as a combination of background energy and command energy. The

 doi.org/10.32388/4TJ5TM

background energy depends on the state of the DRAM device (number of active banks, active/power down/self refresh), while only the command energy directly depends on the issued commands.

- In Section 4, you state: “To address this discrepancy, our calibration only focuses on the IDD currents that have the major contributors to overall DRAM power consumption, since the operating voltage of the DRAM chip is fixed to a standard value (1.2V for DDR4 in our case).” Again, this is wrong. DDR4 uses three operating voltages, namely VDD, VPP, and VDDQ. The wordline boost voltage VPP is 2.5V.
- It is also unclear how you consider the interface power consumption because current is drawn both over the VDDQ power supply of the DRAM devices and the memory controller.
- You state: “Figure 4(b) compares the original IDD currents with the calibrated values, where five key IDD currents—corresponding to activation, pre-charge, act-standby, read, and write operations of DRAM—are successfully calibrated.” According to Figure 4b, the corresponding currents are IDD0, IDD2N, IDD3N, IDD4R, and IDD4W. This is again wrong. Please carefully read the standard and understand how the currents are measured. IDD0 is measured with a repeating ACT-PRE-ACT-PRE... pattern. This is not identical to “activation”. IDD2N is measured with all banks precharged and no commands at all are issued, i.e., it is different from “precharge”.
- In Figure 4b, the calibrated value of IDD2N is higher than the calibrated value of IDD3N. This is impossible. IDD3N is measured with all banks active, which definitely consumes more power than all banks precharged. This shows that your calibration approach simply does not work. For a real evaluation, you should calibrate with a benchmark set A and then evaluate with a different benchmark set B. In this case, I am sure that the measured and simulated energy values would differ.
- I would not call DRAMPower a “statistical tool”.

Declarations

Potential competing interests: No potential competing interests to declare.