UF | UNIVERSITY of FLORIDA

Characterization of Workload-induced Voltage Noises on CPU Power Delivery Networks

Hybrid IBIS Summit at IEEE EMC+SIPI
Grand Rapids, MI
August 4, 2023

Shuo Wang[1], Zihao Zhan[1], Dean Sullivan[2]

1. University of Florida
2. University of New Hampshire
Email: shuowang@ieee.org

UF UNIVERSITY of FLORIDA

Outline

- The security impact of di/dt noise induced by workloads on CPU
- Task 1: Use microbenchmark programs to assist in measuring the PDN impedance profile
- Task 2: Use microbenchmark programs to assist profile the power consumption contributed by different CPU microarchitectural activities

Motivation

The workload in a CPU will cause voltage fluctuations on the power delivery network (PDN).

- Effect 1: Power/EM side-channel leakages that expose computing states.
- Effect 2: di/dt voltage noise can interfere with the other components on the PDN, which can potentially be exploited to implement fault injection attacks.

Can we find an approach that can accurately and efficiently estimate the impact of di/dt noise for any workloads?

Essential information required:

- Impedance profile: Due to the complicated circuit model of the PDN, workload at different frequencies have varied impact on the PDN voltage
- Noise intensities: How large can the interference be when different components in the CPU are activated

Task 1: Measure The Impedance Profile of PDN

Simplified PDN model: Voltage regulator regulates output voltage based on workload, switching frequency ~1MHz. High-frequency noise is harder to handle.

Challenge: Hard to directly measure the impedance inside the CPU because no pins are directly accessible to apply excitation signals

Solution: Measure the EM and voltage trace while running frequency-sweeping microbenchmark programs

Two measurements:

- Measure EM emanations on CPU
- Measure Vcc/Vss sense pins

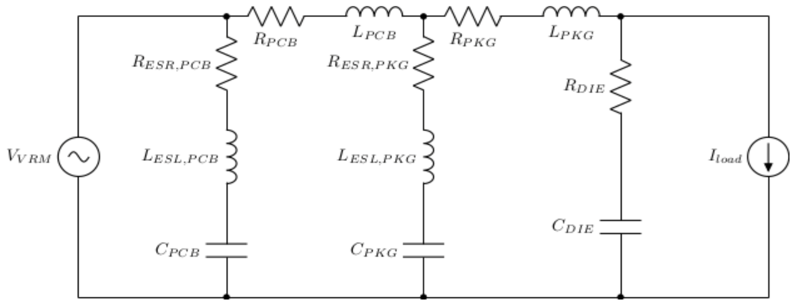

Simplified PDN model

UF UNIVERSITY of FLORIDA

EM and Voltage Signal Measurement

EM emanation: Placing a magnetic probe, under the CPU mounted on a motherboard

Core voltage: Vcc_sense and Vss_sense pins available near the voltage regulator module on the motherboard

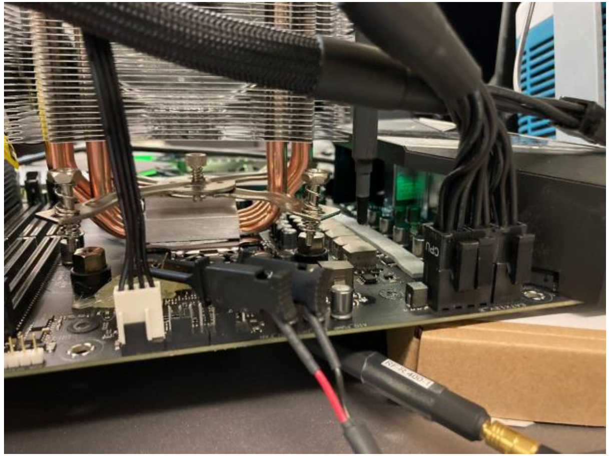

Voltage and EM signal measurement setup

Frequency-sweeping Microbenchmark Construction

Requirements:

- Capable of switching CPU activity intensity between distinct levels (fixed amplitude)
- Frequency can be tuned at a fine resolution

Implementation:

- SUB/NOP instruction pair, use 4 different destination registers for SUB instructions to avoid dependencies between instructions so that SUB instructions can be executed at their theoretical maximum instruction per cycle (IPC)
- On average, each instruction takes 0.25 cycles
- Use run-time compilation to increase the speed of sweeping the frequency

Measured Impedance Profile

- Frequency-sweeping benchmark program switching between SUB/NOP instructions
- Run-time compilation to gradually increase the loop length to reduce the switching frequency
- Collect the spectra using Spectrum analyzer and plot all spectra together.
- As the microbenchmark program sweeps the frequency, the frequencies of peaks at the spectra sweeps from high frequency to low frequency.
- EM and Voltage measurements give spectra with the same shape.Two highest peaks correspond to the first and second order resonance frequency can be observed.

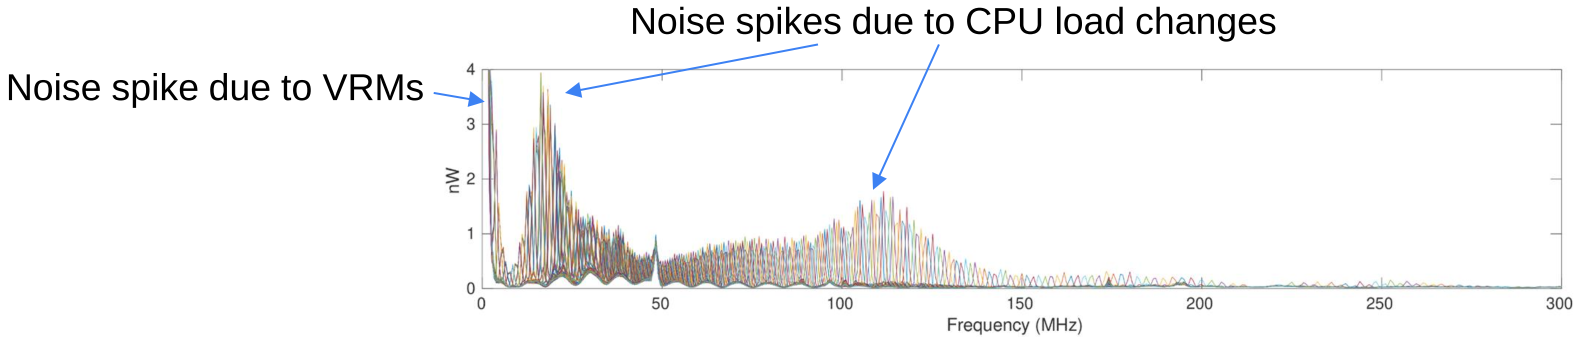

Measured EM signal spectra

Task 2: Measure Noise Intensities Caused by Workloads

The load current inside the CPU die I_{load} can not be directly measured

- The die current need to be estimated using core voltage and power consumption, which are measurable.

Challenges caused by the complex CPU designs

- Multiple cores
- Memory hierarchy – multi-level cache
- Out-of-order execution
- Branch predictor

Power consumption is affected by dynamic voltage and frequency scaling (DVFS), which switches the CPU to different performance levels, each level specifies a clock frequency f_{clk} and core voltage V_{core}.

- At different performance levels, the same workload can induce significantly varied power consumptions.

Model The Impact of Workloads

In CMOS devices, switching behaviors of transistors are the main cause of power consumptions.

- All microarchitectural activities in a CPU contribute to the total power consumption
- The impact of different activities can be described by the following equation
  - For a certain activity pi in CPU, its state O_{pi} represents whether it is activating in a single cycle.
  - The equivalent capacitance C_{pi} reflect the amount of capacitance changes caused by the switching transistors of this activity.
  - Sum up capacitances of all considered subparts to obtain the total load capacitance C_{load}
  - The uncategorized factors not counted in are included in the C_{noise}

$$C_{load} = C_{noise} + \sum_i O_{pi} C_{pi}$$

Power consumption P_{core} can be estimated by V_{core},f_{clk}, and C_{load} by equation:

$$P_{core} = V^2_{core} f_{clk} C_{load}$$

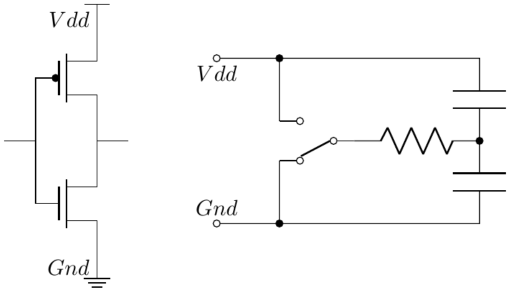

CMOS Inverter Example

Microbenchmarks Construction

- Select an instruction [INSTR_EVAL] to evaluate in an infinite loop
- For the the same instruction, REP[N] duplicate the evaluated instruction [N] times, NOP_REP[N] do the same duplication but insert NOP instructions between evaluated instructions
- In the experiments, we select 9 different N in the range from 1 to 5000 to create loops with varied lengths. This variation will induce differences on some CPU activities such as instruction cache hit/misses and branching instruction rate

```
asm(
        ".intel_syntax noprefix\n\t"
        "CPUID\n\t"
        "lea %%rsi, %0\n\t"
        "lea %%rdi, %1\n\t"
        "mov %%r13, 0\n\t"
        "lea %%r14, %2\n\t"
        "loop:\n\t"
        REP10(INSTR_EVAL)
        "jmp loop\n\t"
    :: "m" (rsi), "m" (rdi), "m" (r14): "memory");
```

```
asm(
        ".intel_syntax noprefix\n\t"
        "CPUID\n\t"
        "lea %%rsi, %0\n\t"
        "lea %%rdi, %1\n\t"
        "mov %%r13, 0\n\t"
        "lea %%r14, %2\n\t"
        "loop:\n\t"
        NOP_REP10(INSTR_EVAL)
        "jmp loop\n\t"
    :: "m" (rsi), "m" (rdi), "m" (r14): "memory");
```

Evaluated Instructions

Select 16 benchmark instruction sequences

- The selected sequences can cause varied activity levels at different components in a CPU. E.g
  - Instruction per cycle
  - Memory load/store
  - Port utilization
  - Front-end/back-end stalls
- Each instruction sequence are used to construct 2x9 microbenchmarks. 288 microbenchmarks on total are evaluated

```
'nop', 'nop', 'nop', 'nop'
'idiv bx', 'idiv bx', 'idiv bx', 'idiv bx'
'idiv ebx', 'idiv ebx', 'idiv ebx', 'idiv ebx'
'idiv rbx', 'idiv rbx', 'idiv rbx', 'idiv rbx'
'imul r8, r12', 'imul r9, r12', 'imul r10, r12', 'imul r11, r12'
'imul r8, qword ptr [r13 + r14]', 'imul r9, qword ptr [r13 + r14]', 'imul r10, qword ptr [r13 + r14]', 'imul r11, qword ptr [r13 + r14]'
'lzcnt r8, r12', 'lzcnt r9, r12', 'lzcnt r10, r12', 'lzcnt r11, r12'
'mov r8, qword ptr [r13 + r14]', 'mov r9, qword ptr [r13 + r14]', 'mov r10, qword ptr [r13 + r14]', 'mov r11, qword ptr [r13 + r14]'
'mov qword ptr [r13 + r14], r8', 'mov qword ptr [r13 + r14], r9', 'mov qword ptr [r13 + r14], r10', 'mov qword ptr [r13 + r14], r11',
'pcmpeqd mm0, mm4', 'pcmpeqd mm1, mm4', 'pcmpeqd mm2, mm4', 'pcmpeqd mm3, mm4'
'ror r8, 2', 'ror r9, 2', 'ror r10, 2', 'ror r11, 2'
'ror qword ptr [r13 + r14], 2', 'ror qword ptr [r13 + r14], 2', 'ror qword ptr [r13 + r14], 2', 'ror qword ptr [r13 + r14], 2'
'shr qword ptr [r13 + r14], 2', 'shr qword ptr [r13 + r14], 2', 'shr qword ptr [r13 + r14], 2', 'shr qword ptr [r13 + r14], 2'
'sub r8, 0x1', 'sub r9, 0x1', 'sub r10, 0x1', 'sub r11, 0x1'
'sub r8, r9', 'sub r11, r8', 'sub r10, r11', 'sub r9, r10'
'sub r8, qword ptr [r13 + r14]', 'sub r9, 0x1', 'sub r10, qword ptr [r13 + r14]', 'sub r11, 0x1'
```

Power Consumption Measurement

Use running average power limit (RAPL) interface to measure the CPU core power consumption

- Read model-specific registers (MSRs) to estimate the power consumption of CPU cores.

Measurements are conducted at 15 different performance levels, where the clock frequency is set to 800, 1000, … 3600 MHz respectively.

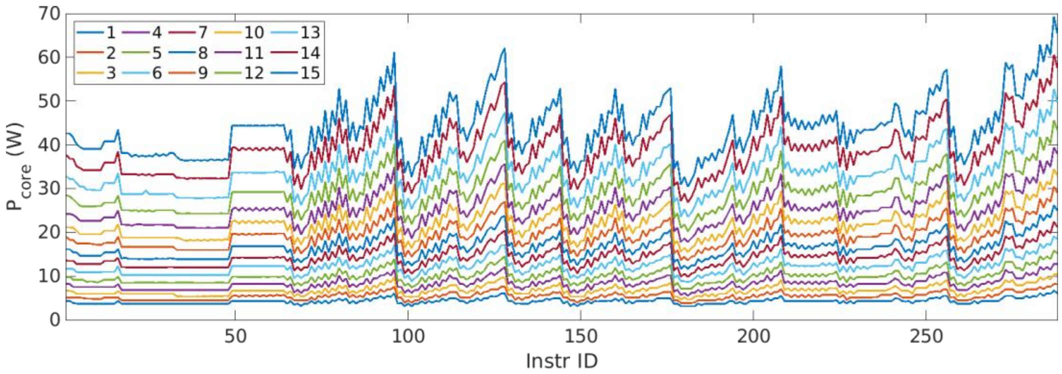

Voltage Measurement

Use Oscilloscope to collect trace from the Vcc_sense and Vss_sense pins

- Preserve the DC component in the voltage trace as the measurement

Voltage is mainly affected by the performance level.

Workload also affect core voltage, likely due to IR drop. High power consumption workload causes more voltage drops.

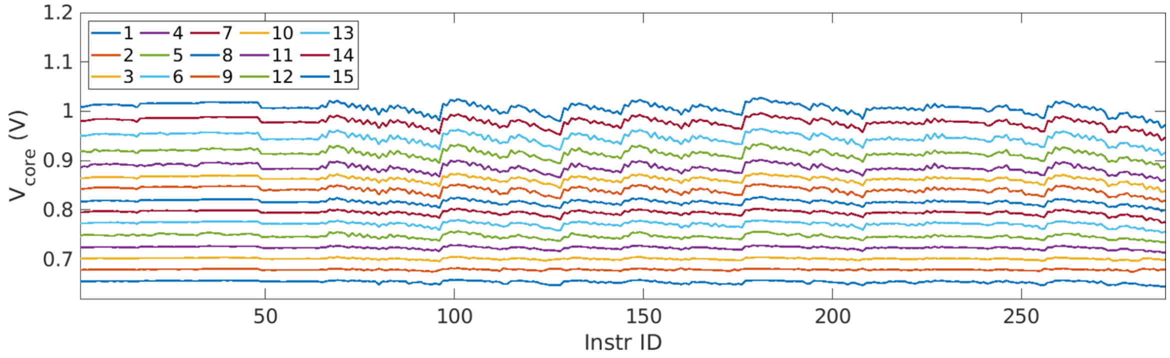

Load Capacitance Estimation

With P_{core}, V_{core} being measured and f_{clk} being documented values.

The load equivalent capacitance can be estimated based on equation:

$$C_{load} = P_{core}/(V^2_{core} f_{clk})$$

Ideally, if the model is correct, even if power consumption and voltage changes significantly when performance level is swiched, C_{load} will keep constant as predicted by our model and verified by the results in the figure below.

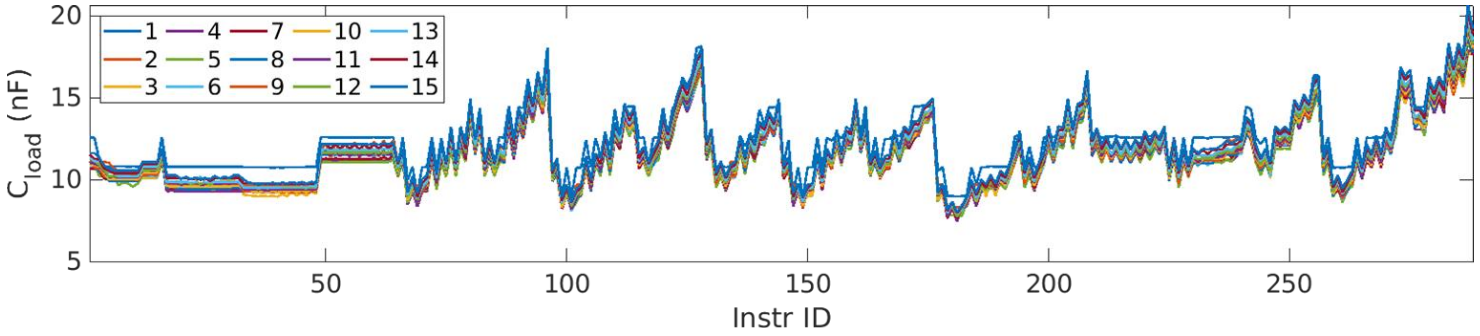

Use Perf program to Measure Microarchitecture Activities

Select 35 metrics related to most important microarchitectural activities

Results indicate that :

- Generated microbenchmark programs can induce various microarchitectural activity patterns

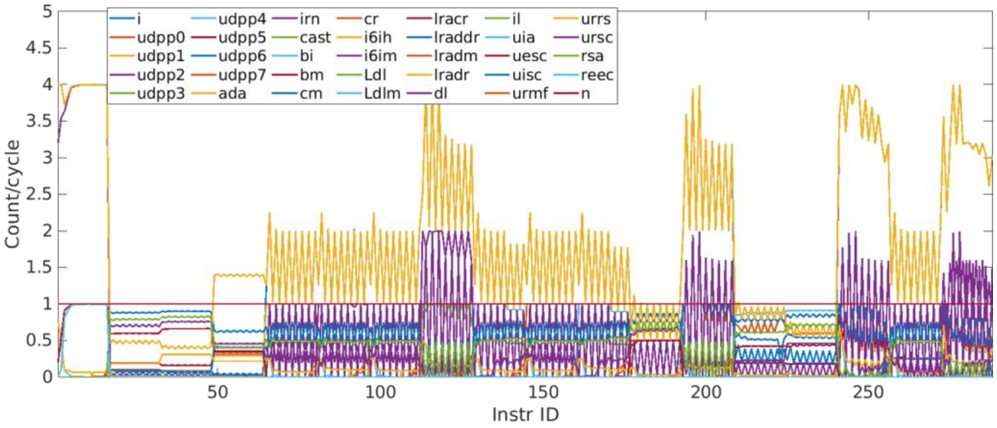

Measure/Calculate Microarchitectural Activity Intensities

The capacitance contribution of different microarchitectural activities can be estimated using previously measured results based on the following equation.

$$\begin{bmatrix} C_{load}(0) \\ C_{load}(1) \\ \dots \\ C_{load}(T) \end{bmatrix} = \begin{bmatrix} O_{p_1}(0) & O_{p_2}(0) & \dots & O_{p_I}(0) \\ O_{p_1}(1) & O_{p_2}(1) & \dots & O_{p_I}(1) \\ \dots & \dots & \dots & \dots \\ O_{p_1}(T) & O_{p_2}(T) & \dots & O_{p_I}(T) \end{bmatrix} \begin{bmatrix} C_{p_1} \\ C_{p_2} \\ \dots \\ C_{p_I} \end{bmatrix} + C_{noise}$$

The results obtained at 15 different performance levels agree well, which verifies the validness of our model.

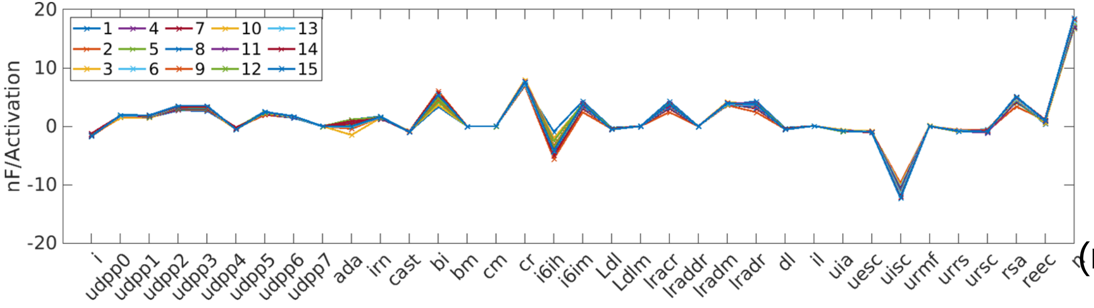

(noise)

Verify The Accuracy of This Model

Using the obtained capacitances corresponding to each activity, predict the power consumption

Compare the results estimated using the model (left) with the measured results (right)

The obtained model can be used to estimate the workload-induced power consumption accurately

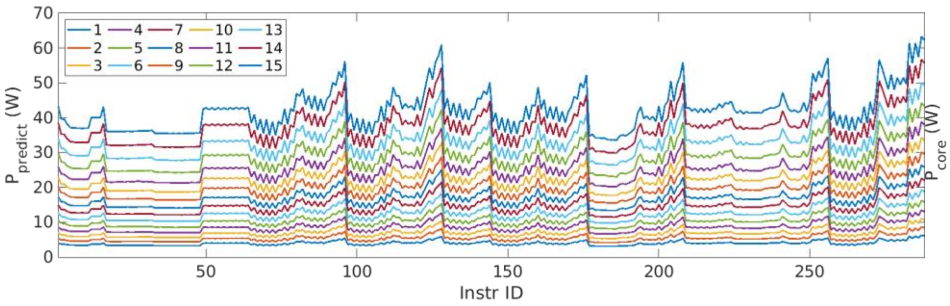

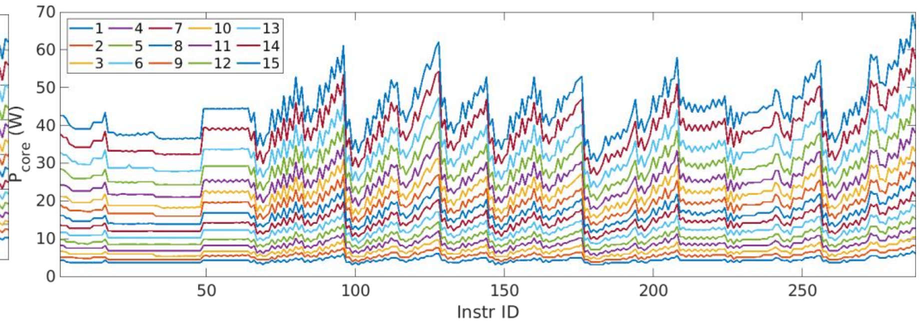

How to Use the Results to Estimate The Voltage Noise

1. Given an instruction sequences, we can estimate the possible micro architecture activity traces.
2. Combining the trace with the obtained results – capacitance of different activities, we can estimate the equivalent load capacitace at different cycles.
3. Estimate the time-series trace of equivalent load capacitance changes at different supply voltages and frequencies defined at different performance levels
4. Use the measured impedance profile to calculate the di/dt voltage noise induced by the workload

UF UNIVERSITY of FLORIDA

Summary

- Through using frequency-sweeping microbenchmarks with spectrum measurement of voltage and EM signals, we can successfully measure the impedance profile inside a CPU.
- Through the microbenchmarks constructed using various instructions to activate different CPU activities, the power consumption contribution of these activities can be estimated.
- On top of this framework, more experiments with more microbenchmarks and more measured microarchitectural activities can provide more accurate parameter estimation which can further increase the effectivness of this model.

Q&A