

FlexWatts: A Power- and Workload-Aware Hybrid Power Delivery Network for Energy-Efficient Microprocessors

*Jawad Haj-Yahya*¹

*Mohammed Alser*¹ *Jeremie S. Kim*¹ *Lois Orosa*¹ *Efraim Rotem*²
Avi Mendelson^{3,4} *Anupam Chattopadhyay*⁴ *Onur Mutlu*¹



The Power Delivery Network (PDN) Debate



2013 Intel Haswell Uses **IVR** PDN (FIVR)

The Haswell Review: Intel Core i7-4770K & i5-4670K Tested

With FIVR, it's easy to implement tons of voltage rails

on-package and efficiently distribute power to all areas of the chip. Voltage ramps are 5 - 10x quicker with FIVR than with a traditional on-board voltage regulator implementation.

FIVR also comes with a reduction in board area and component cost.

2015 Intel Skylake Uses **MBVR** PDN

The Intel Skylake Mobile and Desktop Launch, with Architecture Analysis

CPU Power Delivery – Moving the FIVR

For Skylake, the voltage regulation is moved back into the hands of the motherboard manufacturers. This should allow for cooler processors depending on how the silicon works, but it will result in slightly more expensive motherboards.

2017 AMD Ryzen Uses **LDO** PDN

Ryzen Mobile is Launched: AMD APUs for Laptops, with Vega and Updated Zen

When Intel introduced their FIVR implementation, they said that they found better efficiency using their big inductors and decided against the linear LDO regulators because they were inefficient at low power. We put that to Sam Naffziger, AMD's top guy on power, and he responded that yes, as a percentage, the power efficiency at idle might be lower than expected – but the power consumption of an idle core while another is loaded is still a very tiny proportion. Sam stated that when the LDO is in complete power gate mode, it can be

they still worked hard on the LDO implementation for power efficiency anyway, to make sure everything still worked. Overall, total current requirements were down 36%, which reduces the motherboard-side power regulation, leading to smaller, lighter, and potentially cooler designs.

2019 Intel Icelake Uses **IVR** PDN (FIVR)

Examining Intel's Ice Lake Processors: Taking a Bite of the Sunny Cove Microarchitecture

Intel is keen to promote that one of the new features of Ice Lake is its Thin Magnetic Inductor Array, which helps the FIVR achieve better power conversion efficiencies and waste less power. The main issue with a FIVR is at low power consumption states that have a lot of inefficiency – some other processor designs have



Power Delivery Affecting Performance At 7nm

One particularly troublesome area involves the power delivery network (PDN). To distill it to its simplest form, resistance is going up because of decreasing dimensions. That causes more IR drop, which in turn affects timing, sometimes in unexpected ways. Chips are coming back that are not able to run at intended clock speed.

<https://semiengineering.com/power-delivery-affecting-performance-at-7nm/>

Executive Summary

Problem: Client processors typically use one of three commonly-used power delivery network (PDN) architectures: 1) **MBVR**, 2) **IVR**, and 3) **LDO**. The energy-efficiency of each of these PDNs varies with the **processor power** and **workload characteristics**.

Goal: Provide high energy-efficiency and high performance PDN architecture by leveraging the advantages of each one of the three PDN architectures **across the processor's wide range of power consumption and workloads**.

Mechanism: **FlexWatts**, a new hybrid adaptive PDN for client processors that introduces

- **New hybrid voltage regulators (VRs)** that are allocated for processor domains **with a wide power consumption range** (e.g., CPU cores and graphics engines) that dynamically switch between two modes: **IVR-Mode** and **LDO-Mode**.
- **Static allocation of off-chip VRs**, which have high energy-efficiency for **low and narrow power ranges**, for the rest of the domains (e.g., IO domain).
- A **novel prediction algorithm** that switches the hybrid PDN to the mode (IVR-Mode/LDO-Mode) that is the most beneficial based on **processor power** and **workload characteristics**.

Evaluation: We evaluate FlexWatts using **our new architectural PDN model (PDNspot)**:

- **Improves the average performance** of the SPEC CPU2006 and 3DMark06 workloads by 22% and 25%, respectively, for 4W thermal design power (TDP) system.
- **Reduces the average power consumption** of video playback workload by 11% across all tested TDPs (4W-50W).

Overview of Modern Client PDN Architectures

- The **Power Delivery Network (PDN)** is the electrical system that provides **supply voltage** to the processor's **domains**

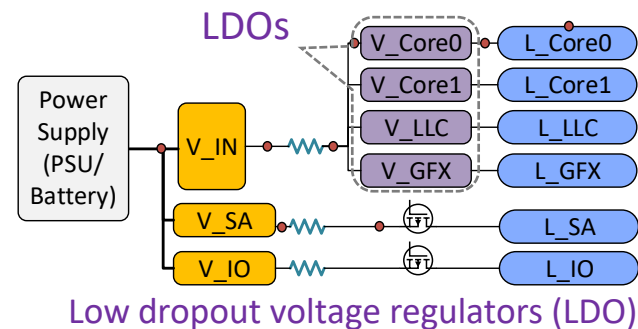
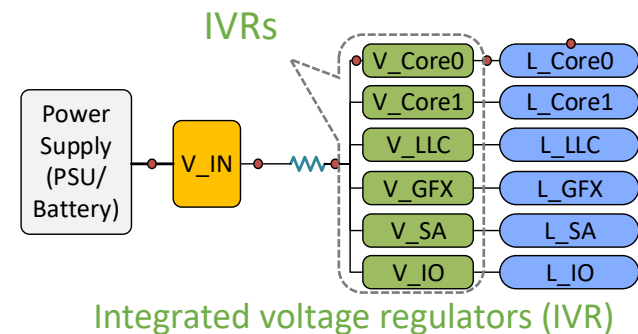
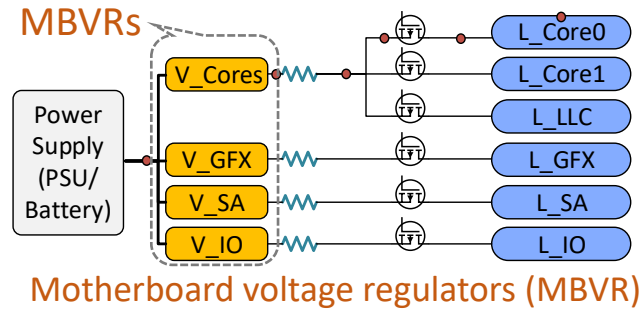
- A PDN consists of 1) a power supply, 2) **voltage regulators** (VRs), 3) network of interconnections, 4) decoupling capacitors (not graphed), and 5) **power-gates**

- There are **3 different commonly-used PDNs**

- Use Switching VRs (**SVRs**) and/or Low dropout VRs (**LDO VRs**)
- An SVR can be placed into the motherboard (**MBVR**) or integrated on-chip (**IVR**)

- The PDNs perform differently based on the **processor power** and **workload characteristics**

- MBVR** and **LDO** PDNs are more energy-efficient at **low thermal design power** (TDP) and **light workloads** compared to **IVR**
- IVR** PDN is more energy-efficient at **high TDP** and **heavy workloads** than **LDO** and **MBVR**
- MBVR** is more energy-efficient than **LDO** in **graphics workloads**



Overview of Modern Client PDN Architectures

- The **Power Delivery Network (PDN)** is the electrical system that provides **supply voltage** to the processor's **domains**

- A PDN consists of 1) a power supply, 2) **voltage regulators** (VRs), 3) network of interconnections, 4) decoupling capacitors (not graphed), and 5) **power-gates**

- There are **3 different commonly-used PDNs**

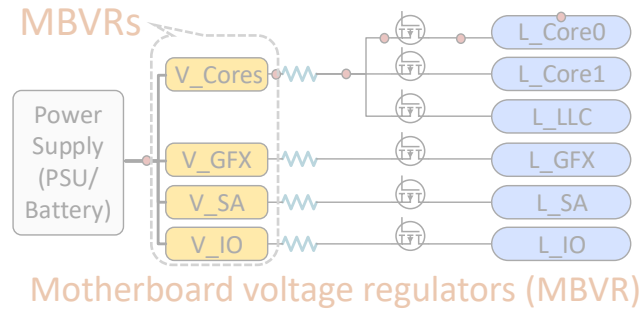
- Use Switching VRs (**SVRs**) and/or Low dropout VRs (**LDO VRs**)
- An SVR can be placed into the motherboard (**MBVR**) or integrated on-chip (**IVR**).

- The PDNs perform differently based on the

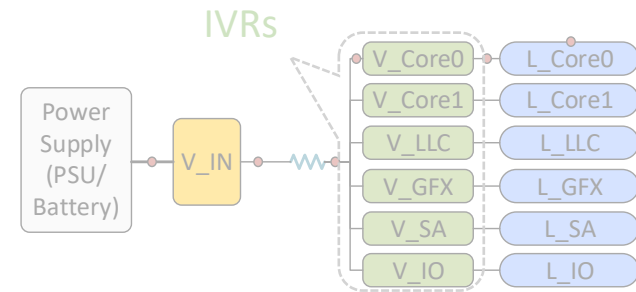
A **single PDN architecture** cannot provide **high energy-efficiency** across a wide **power consumption range** and wide **variety of workloads**

workloads than LDO and MBVR


- MBVR is **more energy-efficient** than LDO in **graphics workloads**



Motherboard voltage regulators (MBVR)



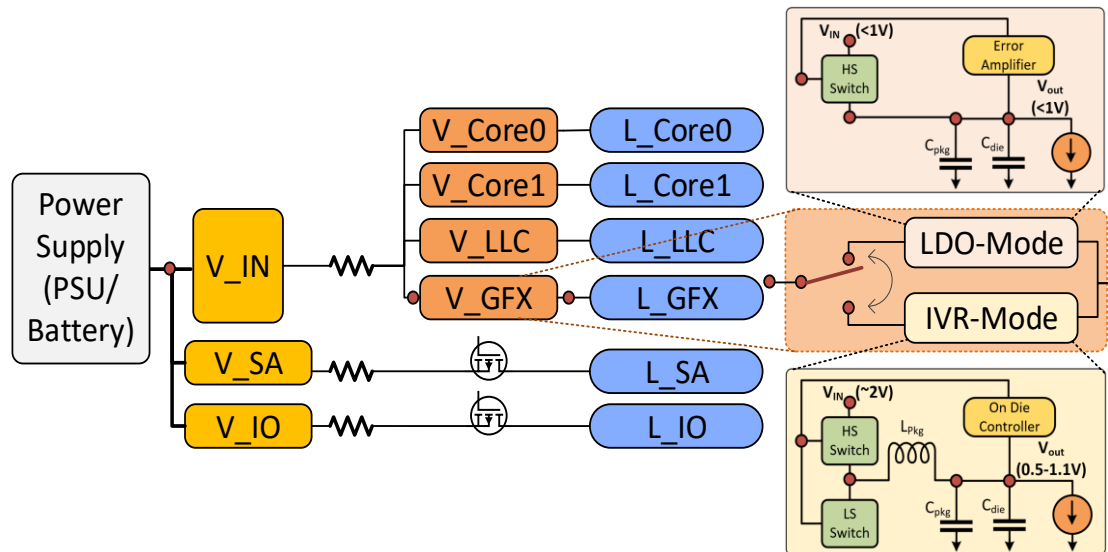
Integrated voltage regulators (IVR)



Low dropout voltage regulators (LDO)

Our Goal: A Hybrid and Adaptive PDN

- Our goal is to provide a PDN architecture that provides high energy-efficiency across the wide range of power consumption and variety of workloads
- To this end, we propose a **hybrid and adaptive PDN** that provides the **advantages** of each one of the **three commonly-used PDNs**
 - by dynamically adapting the hybrid PDN based on **processor power consumption** and **workload characteristics**



FlexWatts: Key Results

- FlexWatts is the **first** hybrid PDN to use two types of on-chip voltage regulators (**IVR** and **LDO**) to leverage the **advantages** of both
- FlexWatts efficiently **chooses the processor PDN** based on the **power demands** and **workload characteristics**
- We evaluate FlexWatts using our **new open-sourced PDNspot model**
 - FlexWatts **improves the performance** of CPU and graphics workloads (by up to 22% and 25%, respectively, for 4W thermal design power (TDP))
 - FlexWatts **reduces the average power consumption** of battery life workloads (by up to 11%) across all TDPs
- We **show** that FlexWatts is an effective approach to **maintain high efficiency and high performance** in metrics of interest in client processors across a wide spectrum of TDPs and workloads with **minimal overhead**

FlexWatts: A Power- and Workload-Aware Hybrid Power Delivery Network for Energy-Efficient Microprocessors

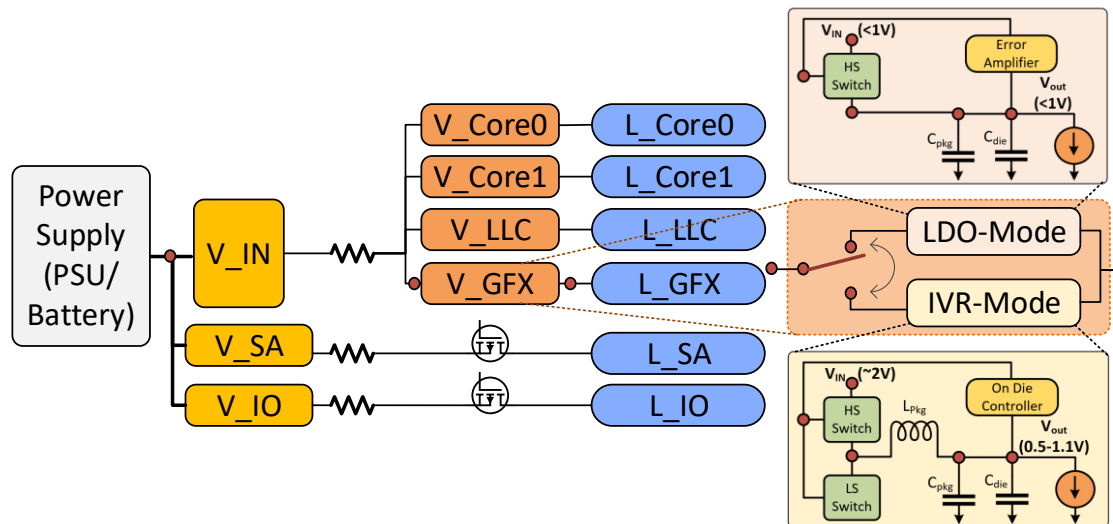
*Jawad Haj-Yahya*¹

*Mohammed Alser*¹ *Jeremie S. Kim*¹ *Lois Orosa*¹ *Efraim Rotem*²
Avi Mendelson^{3,4} *Anupam Chattopadhyay*⁴ *Onur Mutlu*¹



Component 1: Hybrid Adaptive PDN

1. Off-chip VRs allocated to each system domain with a **low and narrow power consumption** range (i.e., SA and IO domains).
2. **Hybrid VRs** allocated to system domains with a **wide power consumption** range (e.g., CPU cores and graphics engines)
 - This hybrid PDN can dynamically **switch between two modes**, **IVR-Mode** and **LDO-Mode**, based on the expected ETEE benefits of each mode for the current workload and power consumption
 - The hybrid PDN **shares multiple** die, package, and board **resources**



Component 2: Voltage Noise-Free Mode-Switching

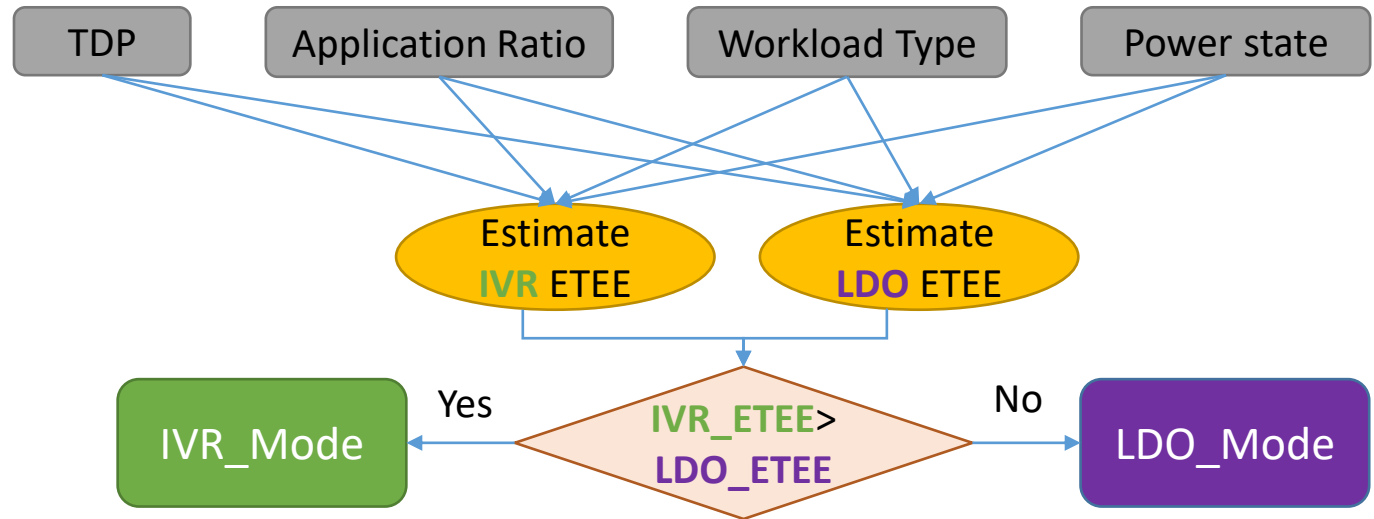
FlexWatts performs mode-switching using 3 steps, while the compute domains are idle to prevent any voltage noise

1. Place the processor in an idle power-state for a short period of time
2. Configure the hybrid PDN and update the on-chip and off-chip VR levels
3. Exit the idle power-state and resume the processor with the new PDN mode

Component 3: Runtime Mode Prediction Algorithm

A new runtime mode-prediction algorithm predicts which PDN mode, among the two modes (IVR-Mode and LDO-Mode), provides the best end-to-end power conversion efficiency (ETEE). ETEE is a function of:

- The AR and the workload type (i.e., single-thread, multi-thread, and graphics)
 - The TDP and the power-state of the system
- We store two sets of ETEE curves inside the power management unit (PMU) firmware, one set for the IVR PDN and the other set for the LDO PDN
 - The PMU firmware estimates each of the input parameters at runtime



Motivational Experiment

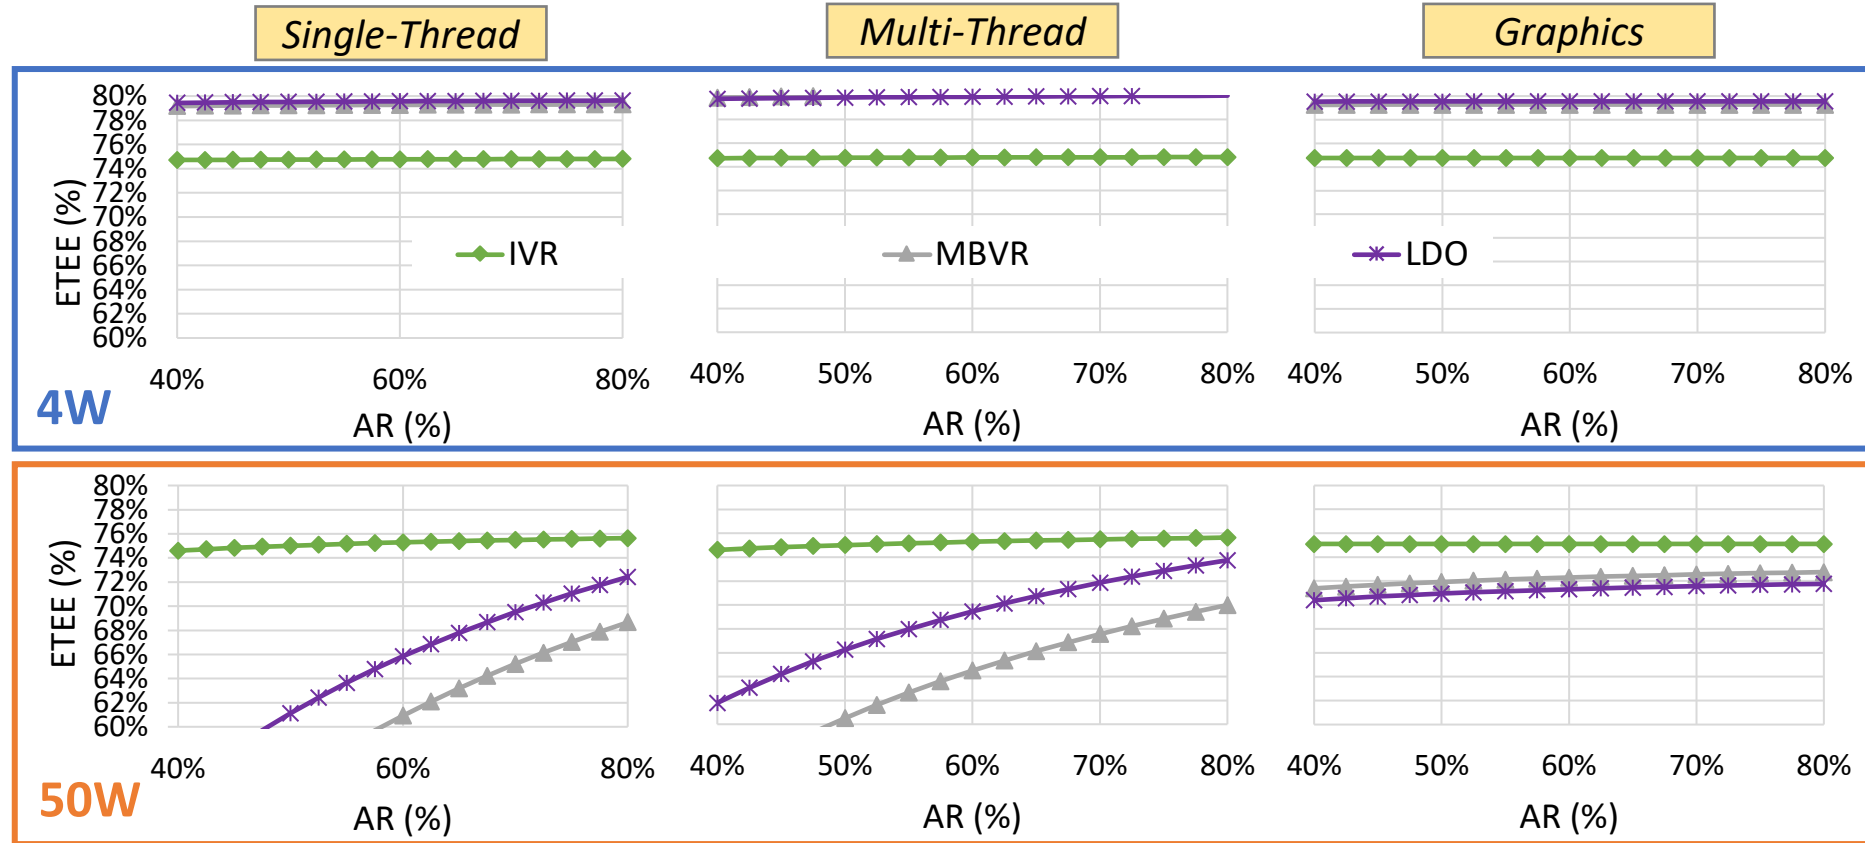
- We evaluate the potential benefits of the 3 PDNs (IVR, MBVR, LDO) across different **TDPs**, **application ratios (ARs)** and **workloads**
 - **AR** is the **switching rate** of a component (e.g., CPU core) for a workload when compared to the highest possible power (**power-virus**)
- We use our validated model, **PDNspot**, to evaluate the efficiency of the three PDNs
- We use the Metric **end-to-end power conversion efficiency (ETEE)**
 - **ETEE** of a PDN is the ratio between total **output power** consumption to total **input power** consumption of the PDN



- We use multiple **workload traces** from **SPEC CPU2006**, **3DMark06**, and **video playback** workloads

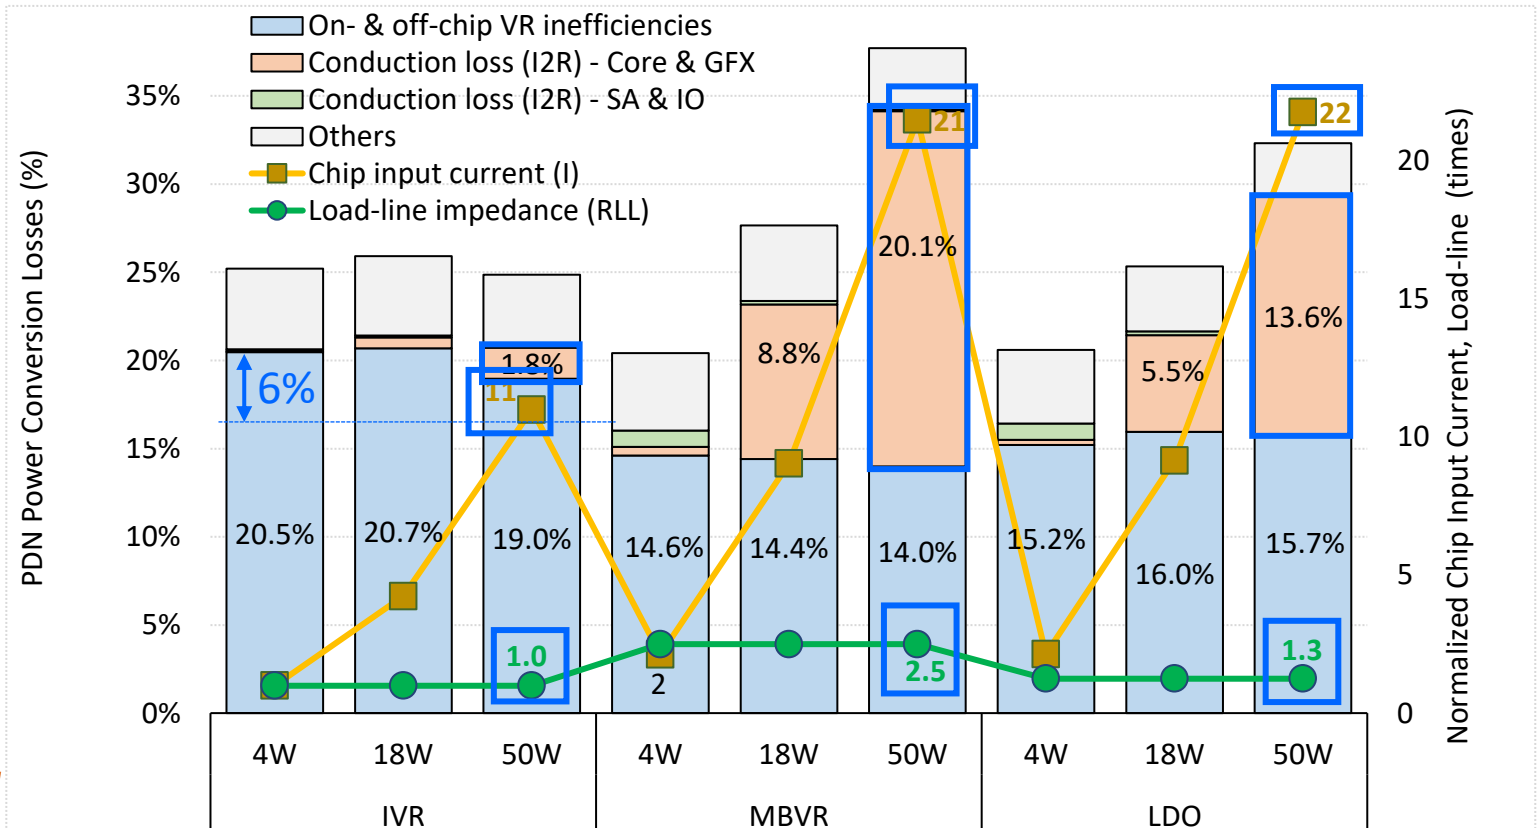
Observation 1: TDP Effect on PDN's Efficiency (1)

- Executing CPU- and graphics-intensive workloads, with different ARs
 - IVR PDN has a lower ETEE at the 4W TDP compared to MBVR and LDO PDNs
 - IVR PDN has a higher ETEE at the 50W TDP compared to MBVR and LDO PDNs
- The ETEE crossover point, at which the IVR ETEE becomes higher than the MBVR/LDO ETEE, exists at some TDP between 4W and 50W.



Observation 1: TDP Effect on PDN Efficiency (2)

- A **breakdown** of PDN power conversion **loss**:
 - At **low-TDP** (e.g., 4W), the **on-chip** and **off-chip VRs' power conversion inefficiencies** are **higher** for the **IVR** PDN than that of the **MBVR** and **LDO**.
 - At **high-TDP** (e.g., 50W), the **I²R loss** in **core and graphics** domains is **higher** for the **MBVR** and **LDO** PDNs than for **IVR**. The high I²R loss is due to:
 - A **2x** higher **chip input current** in the MBVR and LDO PDNs compared to the **IVR** PDN, and
 - A **2.5x/1.3x** higher **load-line impedance (RLL)** in the **MBVR/LDO** PDNs compared to the **IVR** PDN



Observation 1: TDP Effect on PDN's Efficiency

- Executing CPU- and graphics-intensive workloads, with different ARs
 - IVR PDN has a lower ETEE at the 4W TDP compared to MBVR and LDO PDNs
 - IVR PDN has a higher ETEE at the 50W TDP compared to MBVR and LDO PDNs
- The ETEE crossover point, at which the IVR ETEE becomes higher than the MBVR/LDO ETEE, exists at some TDP between 4W and 50W


TDP affects the **ETEE** of a PDN.

At **low-TDP**, the **MBVR** and **LDO** PDNs are more efficient.

At **high-TDP**, the **IVR** PDN is more efficient.

Observation 2: Workload Effect on PDN's Efficiency

- Workload's AR and the workload type (e.g., single-threaded, multi-threaded, graphics) affects PDN ETEE
- The ETEEs of the MBVR & LDO PDNs increase with AR due to load-line effect
- Different workload types have different ETEE curves
 - The LDO ETEE is higher than the MBVR ETEE for CPU-intensive (single- and multi-threaded) workloads, but is lower than the MBVR ETEE for graphics-intensive workloads
 - LDO inefficiency is more dominant in graphics workloads, due to the high voltage difference between the core and graphics domains (graphics' voltage is significantly higher than cores')




Observation 2: Workload Effect on PDN's Efficiency

- Workload's AR and the workload type (e.g., single-threaded, multi-threaded, graphics) affects PDN ETEE

Application ratio and **workload type** have significant effects on each PDN's ETEE.


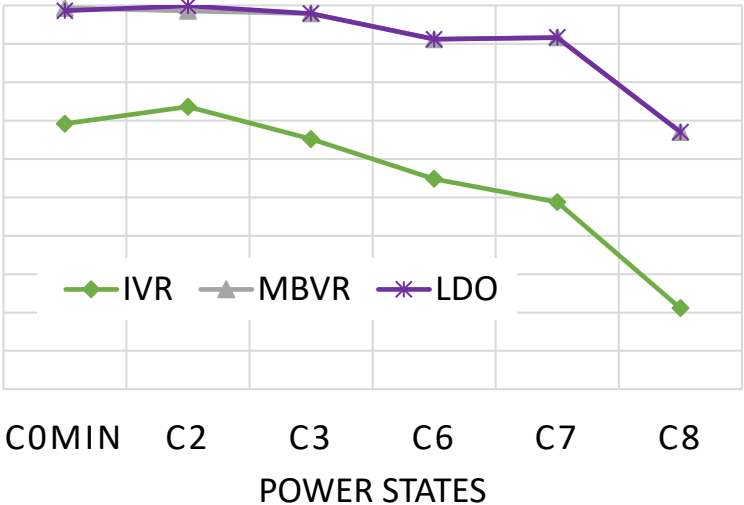
MBVR and **LDO** PDNs have **lower ETEE** with **low application ratios**.

LDO PDN has **lower ETEE** on **graphics workloads**.



Observation 3: Power-State Effect on PDN's Efficiency


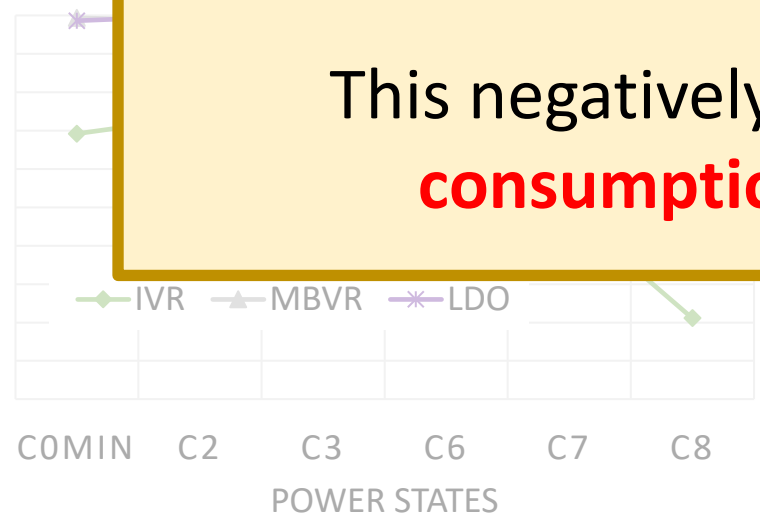
- The ETEE of the IVR PDN is lower than that of MBVR and LDO PDNs for computationally light workloads and low power states across all TDPs
- A video playback workload spends 10%, and 90% of its execution time in light workload (C0_{MIN}) and low power states (C2 and C8) respectively
- MBVR and LDO PDNs have 12% and 11% lower average power consumption, respectively, than the IVR PDN due to the higher ETEE of MBVR and LDO PDNs



Observation 3: Power-State Effect on PDN's Efficiency

- The ETEE of the IVR PDN is lower than that of MBVR and LDO PDNs for computationally light workloads and low power states across all TDPs
- A video playback workload spends 10%, and 90% of its execution time in light workload (C0_{MIN}) and low power states (C2 and C8) respectively

The IVR PDN is energy-inefficient for computationally-light workloads. This negatively impacts both energy consumption and battery life.



Our Goal: A Hybrid and Adaptive PDN

- We conclude that there is **no single PDN** for modern client processors that **provides a high ETEE** across all **TDPs**, **workload types** and **application ratios (ARs)**

Our Goal: A Hybrid and Adaptive PDN

- We conclude that there is **no single PDN** for modern client processors that **provides a high ETEE** across all **TDPs**, **workload types** and **application ratios (ARs)**

Our goal is to provide a **PDN architecture** that provides **high energy-efficiency** across the wide range of **power consumption** and **variety of workloads**.

Presentation Outline

1. Overview of Client Processor PDN Architectures

2. Motivation and Goal

3. FlexWatts

- I. Hybrid Adaptive PDN
- II. Voltage Noise-Free Mode-Switching
- III. Runtime Mode Prediction Algorithm

4. Evaluation

5. Conclusion

FlexWatts

- **FlexWatts** is a new hybrid adaptive PDN that provides a high ETEE for the wide power consumption range and workload diversity of client processors
- FlexWatts is based on three key ideas:
 1. New hybrid voltage regulators (VRs) that are allocated for processor domains with a wide power consumption range that dynamically switch between two modes (IVR-Mode and LDO-Mode) to maintain high energy-efficiency across the wide range
 2. Static allocation of off-chip VRs, which have high energy-efficiency for low and narrow power ranges, for the rest of the domains (e.g., IO domain)
 3. A novel prediction algorithm that switches the hybrid PDN to the mode (IVR-Mode/LDO-Mode) that is the most beneficial based on processor power and workload characteristics

FlexWatts Architecture: 3 Components

1. **Hybrid adaptive PDN:** includes **hybrid VRs** and **off-chip VRs**
2. **Voltage Noise-Free Mode-Switching:** transitions the hybrid PDN between two modes (**IVR-Mode** and **LDO-Mode**)
3. **A new mode prediction algorithm:** automatically determines, which PDN mode would be the **most beneficial** based on **system** and **workload** characteristics

Presentation Outline

1. Overview of Client Processor PDN Architectures

2. Motivation and Goal

3. FlexWatts

I. Hybrid Adaptive PDN

II. Voltage Noise-Free Mode-Switching

III. Runtime Mode Prediction Algorithm


4. Evaluation

5. Conclusion

Methodology


- **Framework**: We evaluate the PDNs using our **new open-sourced PDNspot model**
- **Workloads**: We evaluate **FlexWatts** with three classes of workloads
 - **CPU**: SPEC CPU2006 benchmarks
 - **Graphics**: 3DMARK06 benchmarks
 - **Battery life**: web browsing, light gaming, video conferencing, and video playback benchmarks
- **Comparison Points**: We compare **FlexWatts** to the three commonly-used (**IVR**, **MBVR**, **LDO**) PDNs of client processors

Results – CPU Workloads (TDPs 4-50W)



- At TDPs **lower than 18W**, FlexWatts provides up to **22%** higher performance over the IVR PDN
 - by **operating mainly in LDO-Mode**
- At TDPs **higher than 18W**, FlexWatts provides up to **7%/4%** higher performance over the **MBVR/LDO** PDNs
 - by **operating in IVR-Mode**


Results – CPU Workloads (TDPs 4-50W)



FlexWatts significantly improves **CPU performance** compared to **IVR** at **low TDPs** by operating in **LDO-Mode**.


FlexWatts provides higher **CPU performance** compared to **MBVR/LDO** at **high TDPs** by operating in **IVR-Mode**.

Results – Graphics Workloads



- At TDPs **lower than 25W**, **FlexWatts** provides up to **25%** higher performance over the **IVR** PDN
 - by **operating mainly in LDO-Mode**
- At TDPs **higher than 25W**, FlexWatts provides up to **3%/6%** higher performance over the **MBVR/LDO** PDNs
 - by **operating mainly in IVR-Mode**
- At TDPs **lower than 25W**, FlexWatts performs **slightly worse** (up to **2%**) than **MBVR/LDO** PDNs due to
 - The **higher load-line** of FlexWatts
 - The large difference in operating voltages across the cores/LLC/graphics domains

Results – Graphics Workloads



- At TDPs lower than 25W, FlexWatts provides up to 25% higher performance over the IVR PDN


FlexWatts significantly improves graphics performance compared to IVR at low TDPs by operating in LDO-Mode.

FlexWatts provides higher graphics performance compared to MBVR/LDO at high TDPs by operating in IVR-Mode.

- The higher load-line of FlexWatts
- The large difference in operating voltages across the cores/LLC/graphics domains

Results – Battery Life Workloads


- Battery life workloads have **fixed** performance requirements
 - they consume the same power at all TDPs



- FlexWatts reduces average power consumption (**8%-11%**) on **battery life** workloads compared to **IVR** PDN
- FlexWatts consumes up to **1%** more power than **MBVR** PDN

Results – Battery Life Workloads

- Battery life workloads have **fixed** performance requirements
 - they consume the same power at all TDPs



FlexWatts is almost **as energy-efficient** as both **MBVR** and **LDO**
and up to **11% more energy-efficient** than **IVR**.

Other Results in the Paper

- FlexWatts **board area** and **bill of materials (BOM)** compared to other PDNs:
 - FlexWatts PDN has **comparable BOM** cost and **board area** to IVR PDN
 - FlexWatts PDN has **significantly lower BOM** cost and **board area** compared to **MBVR** and **LDO** PDNs
- Comparison to the **PDN** used in **Intel Skylake-X** processors that **combines IVR** and **MBVR** PDNs
 - **Intel Skylake-X PDN** provides **higher** performance and **lower** energy consumption than the **IVR** PDN
 - **FlexWatts** provides significantly **higher** performance and **lower** energy consumption than the **Intel Skylake-X PDN**