ThermoGater: Thermally-Aware On-Chip Voltage Regulation

S. Karen Khatamifard, Longfei Wang†, Weize Yu†
Selçuk Köse†, Ulya R. Karpuzcu

University of Minnesota
{**khatami**, ukarpuzc}**@umn.edu**

†University of South Florida
{longfei, weizeyu}@mail.usf.edu kose@usf.edu

26/6/2017



UNIVERSITY OF MINNESOTA

UNIVERSITY OF SOUTH FLORIDA

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload
- On-chip voltage regulation enables fast, fine-grain voltage control.

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload
- On-chip voltage regulation enables fast, fine-grain voltage control.
- Power-limited computing platforms of today feature many on-chip regulators.

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload
- On-chip voltage regulation enables fast, fine-grain voltage control.
- Power-limited computing platforms of today feature many on-chip regulators.
- On-chip voltage regulators
  - Convert power from an external energy source to the processor
  - Power conversion loss is inevitable and sizable

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload
- On-chip voltage regulation enables fast, fine-grain voltage control.
- Power-limited computing platforms of today feature many on-chip regulators.
- On-chip voltage regulators
  - Convert power from an external energy source to the processor
  - Power conversion loss is inevitable and sizable
    - Lost power gets dissipated as heat
    - Small regulator footprint → potential thermal hotspots

Motivation

- How to maximize power efficiency?
  - Tailor the voltage to spatio-temporal changes in workload
- On-chip voltage regulation enables fast, fine-grain voltage control.
- Power-limited computing platforms of today feature many on-chip regulators.
- On-chip voltage regulators
  - Convert power from an external energy source to the processor
  - Power conversion loss is inevitable and sizable
    - Lost power gets dissipated as heat
    - Small regulator footprint → potential thermal hotspots

ThermoGater

Architectural governor to orchestrate **thermally-aware on-chip regulation.**

Distributed On-Chip Voltage Regulation

- Many regulators dispersed across chip → maximize physical proximity to load
  - Enables **fast** response time in tailoring operating point to load activity
  - Mitigates voltage noise

Distributed On-Chip Voltage Regulation

- Many regulators dispersed across chip → maximize physical proximity to load
  - Enables **fast** response time in tailoring operating point to load activity
  - Mitigates voltage noise

- Regulator Power conversion efficiency: $\eta = P_{out} / P_{in}$
  - Due to inevitable conversion loss, eta typically is less than 100%
  - Lost power is dissipated as heat
  - eta evolves as a function of microarchitectural activity in the load block

Distributed On-Chip Voltage Regulation

- Many regulators dispersed across chip → maximize physical proximity to load
  - Enables **fast** response time in tailoring operating point to load activity
  - Mitigates voltage noise

- Regulator Power conversion efficiency: $\eta = P_{out} / P_{in}$
  - Due to inevitable conversion loss, eta typically is less than 100%
  - Lost power is dissipated as heat
  - eta evolves as a function of microarchitectural activity in the load block
    - Regulators are calibrated to reach peak eta at specific activity
    - Deviation in activity from the calibrated point degrades eta
    - Even the peak eta barely exceeds 90%

Distributed On-Chip Voltage Regulation

- Many regulators dispersed across chip → maximize physical proximity to load
  - Enables **fast** response time in tailoring operating point to load activity
  - Mitigates voltage noise

- Regulator Power conversion efficiency: $\eta = P_{out} / P_{in}$
  - Due to inevitable conversion loss, eta typically is less than 100%
  - Lost power is dissipated as heat
  - eta evolves as a function of microarchitectural activity in the load block
    - Regulators are calibrated to reach peak eta at specific activity
    - Deviation in activity from the calibrated point degrades eta
    - Even the peak eta barely exceeds 90%

How to sustain operation at peak eta?

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

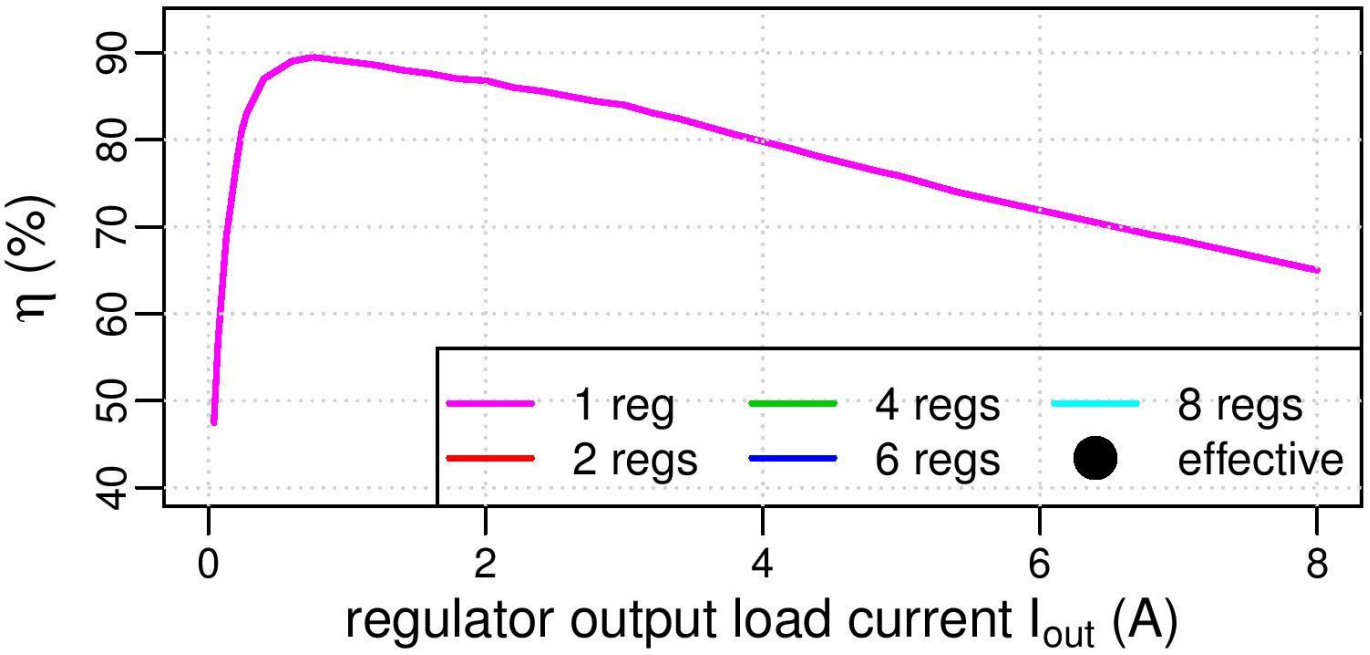

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

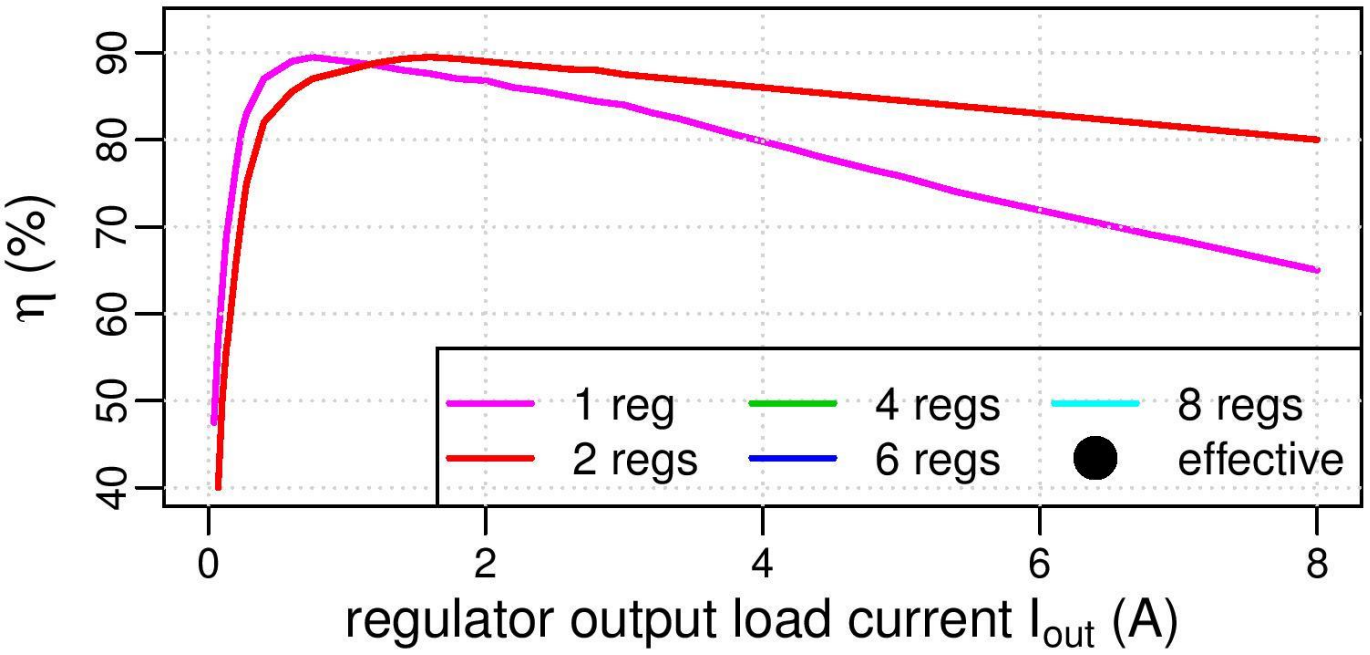

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

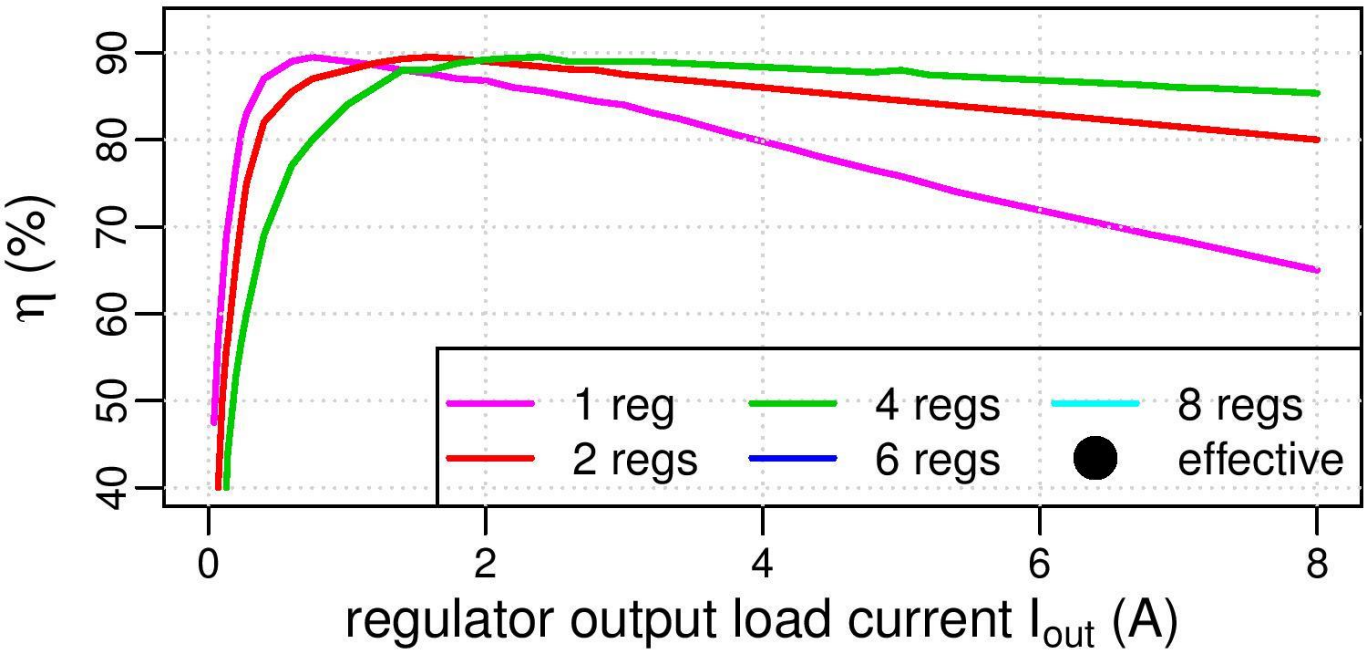

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

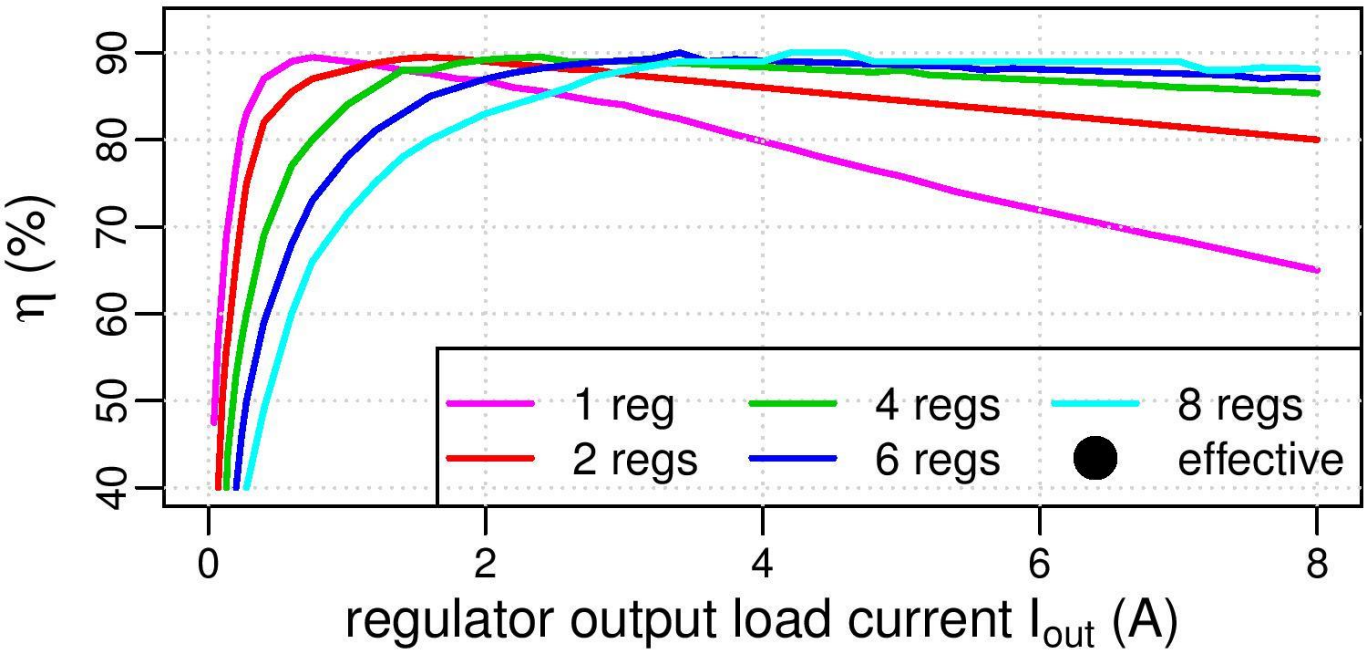

(Temperature-oblivious) Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

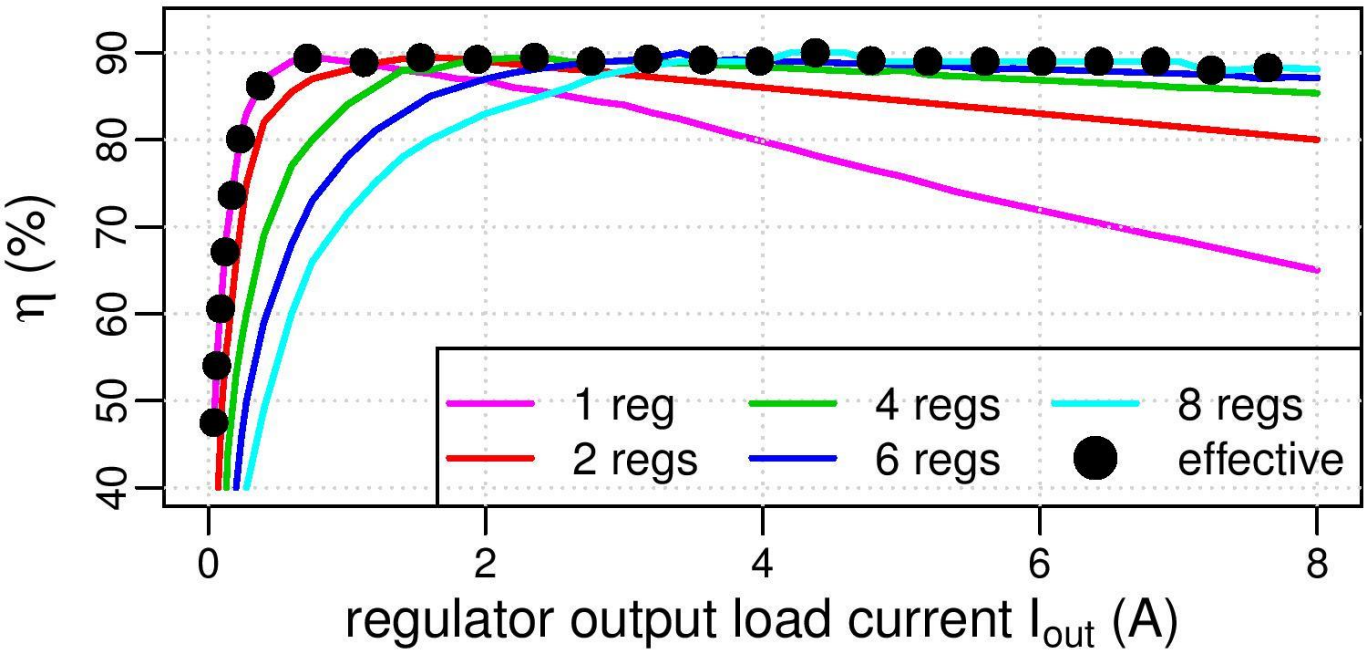

The Case for Temperature-Aware Regulator Gating

- Sustain operation at peak eta
  - By selective shut-down, i.e., gating of component regulators
  - As a function of changes in microarchitectural activity

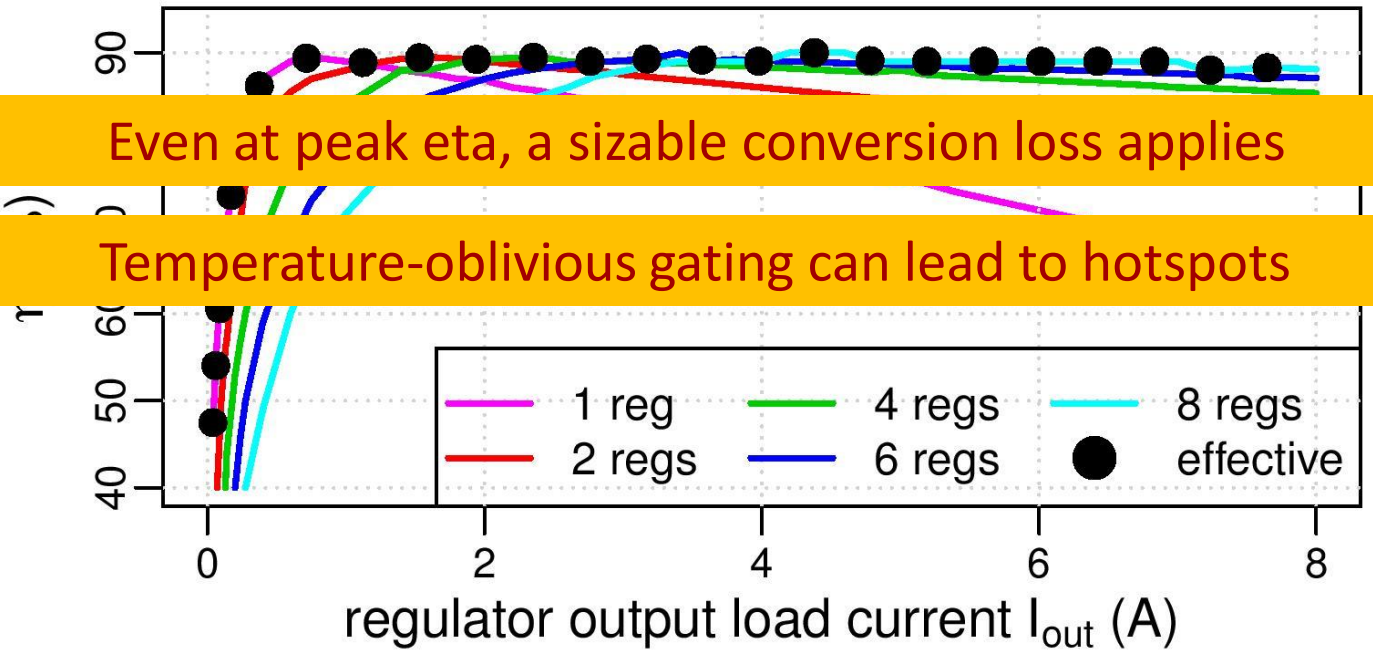

Even at peak eta, a sizable conversion loss applies

Temperature-oblivious gating can lead to hotspots

Gating Policy Design Space

	Operate at peak eta?	Minimize V-noise?	Optimize Temperature?
Current policies			
ThermoGater			

The Case for ThermoGater

Voltage noise ~ 16.8%

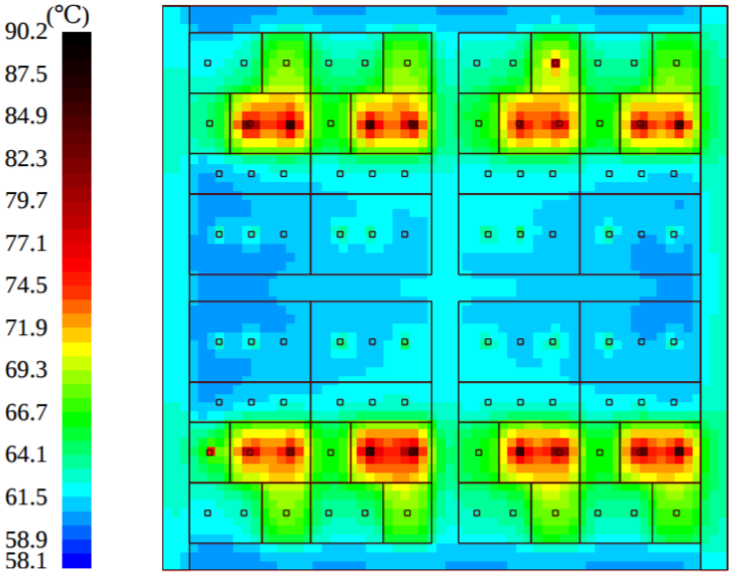

Voltage noise only

The Case for ThermoGater

Voltage noise ~ 16.8%

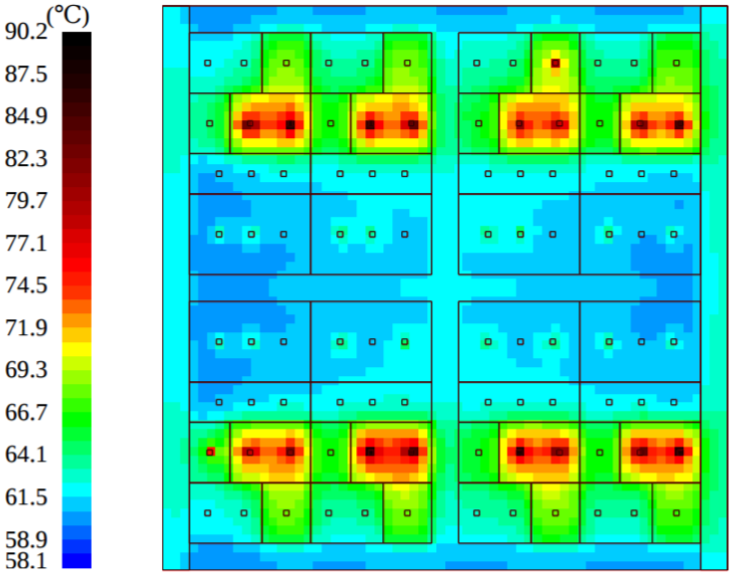

Voltage noise only

Poor **thermal** profile!

The Case for ThermoGater

Voltage noise ~ 16.8%

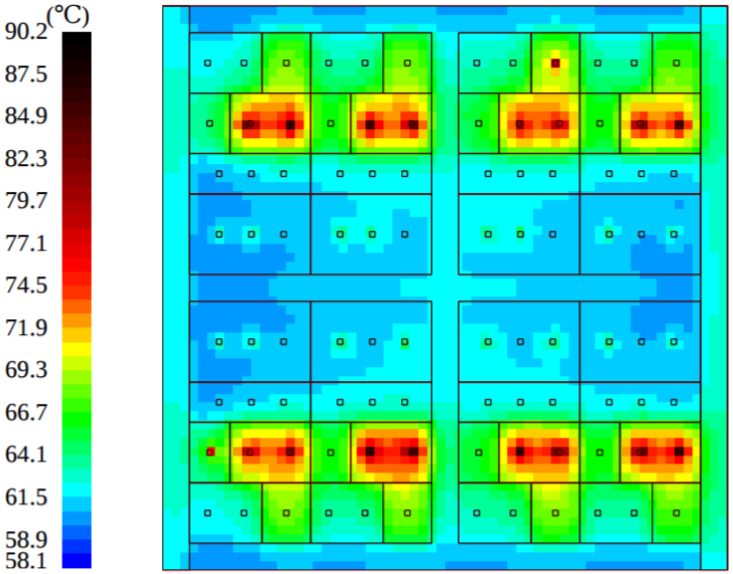

Voltage noise only

Voltage noise ~ 23.4%

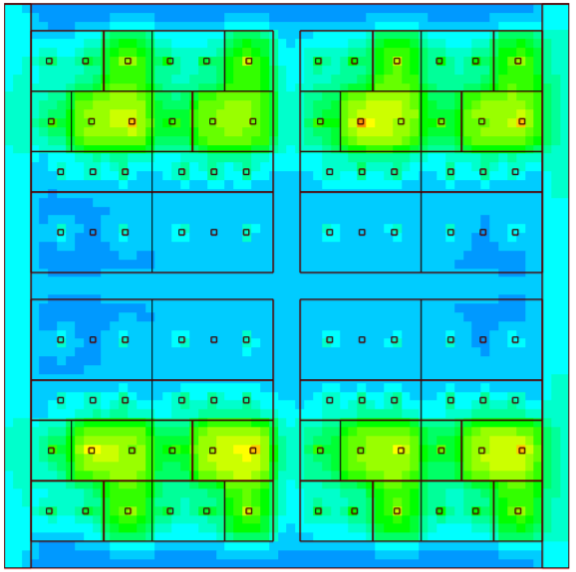

Temperature only

Poor **thermal** profile!

The Case for ThermoGater

Voltage noise ~ 16.8%

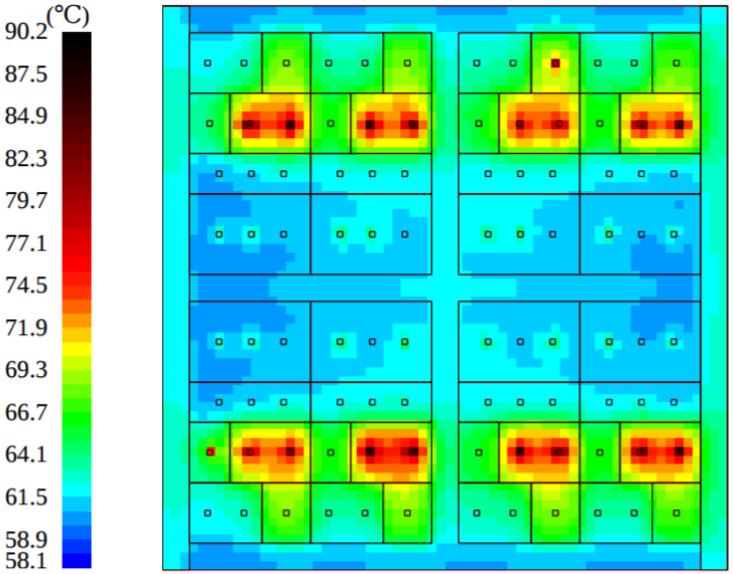

Voltage noise only

Poor **thermal** profile!

Voltage noise ~ 23.4%

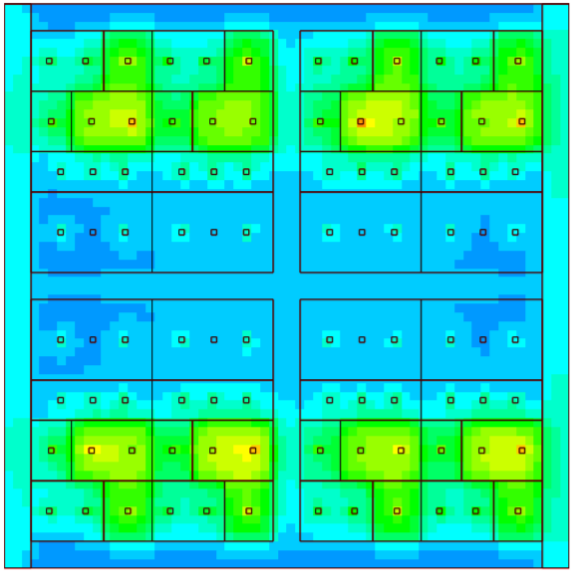

Temperature only

Poor **Voltage noise** profile!

The Case for ThermoGater

Voltage noise ~ 16.8%

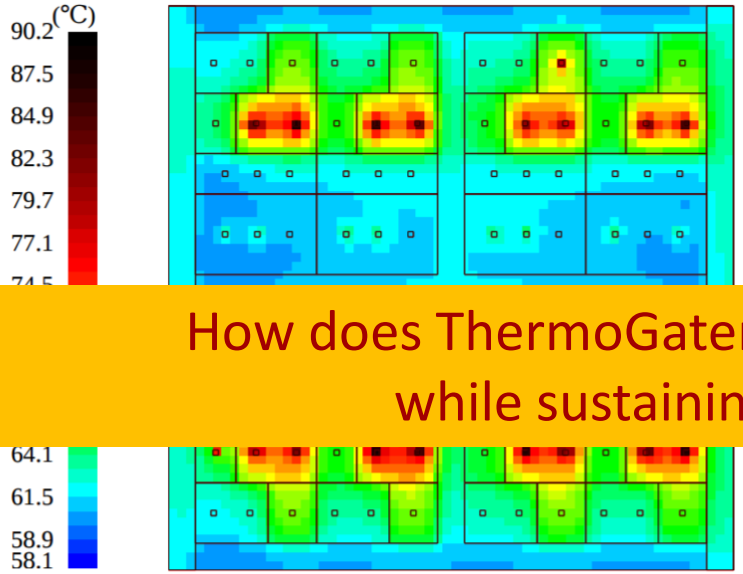

Voltage noise ~ 23.4%

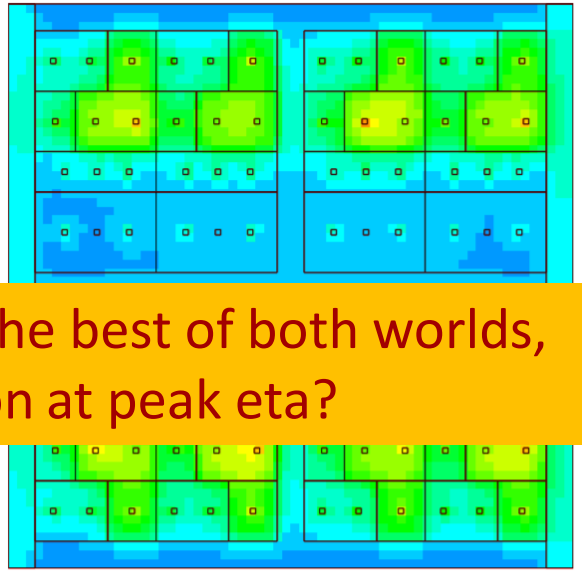

How does ThermoGater achieve the best of both worlds, while sustaining operation at peak eta?

Voltage noise only

Poor **thermal** profile!

Temperature only

Poor **Voltage noise** profile!

Experimental Setup

- IBM POWER8 like 8-core processor
- 96 on-chip regulators, in 16 domains.
- Architectural simulator: SniperSim
- Power simulator: McPAT (MR2 version)
- Thermal simulator: HotSpot
- Voltage noise simulator: VoltSpot
- Benchmarks: Splash2X

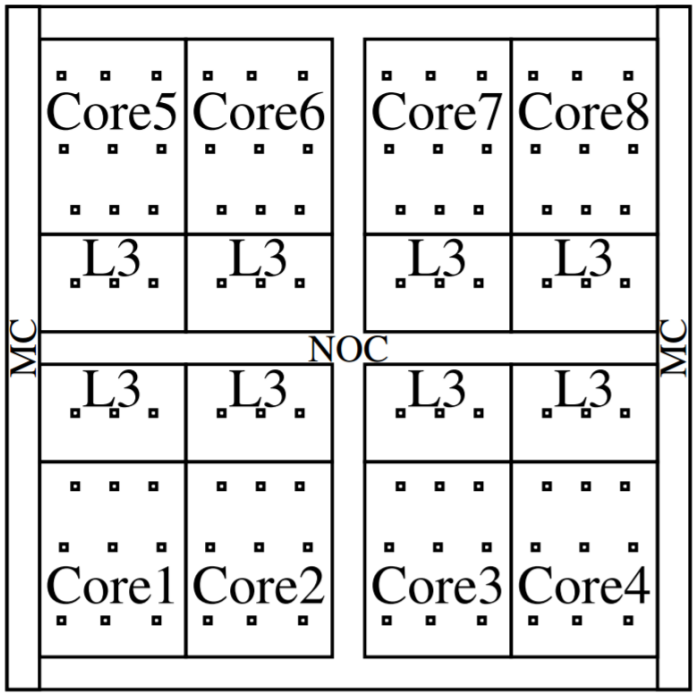

Which Regulators to Turn On?

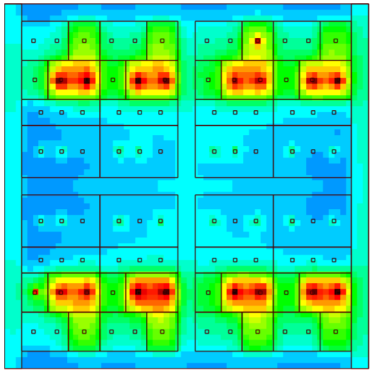

Voltage noise only

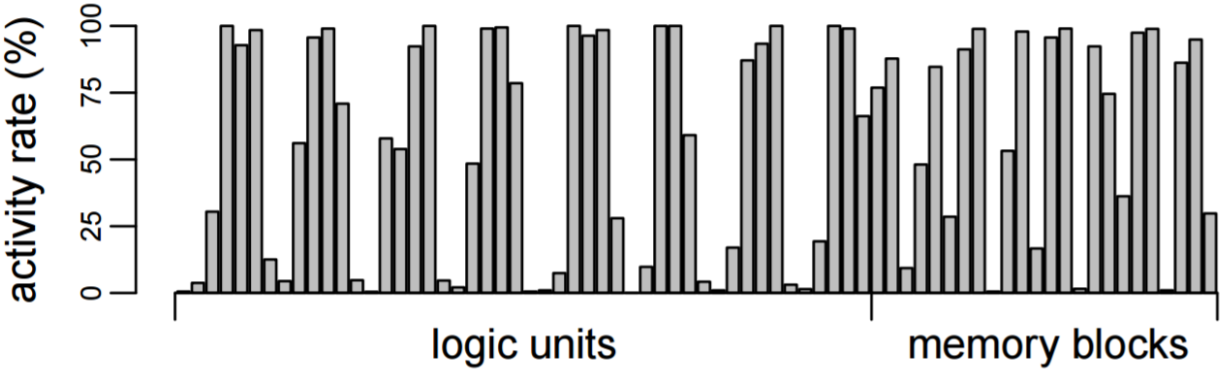

Which Regulators to Turn On?

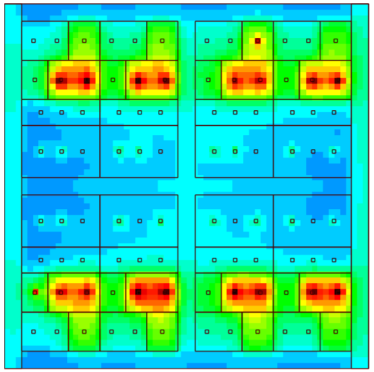

Voltage noise only

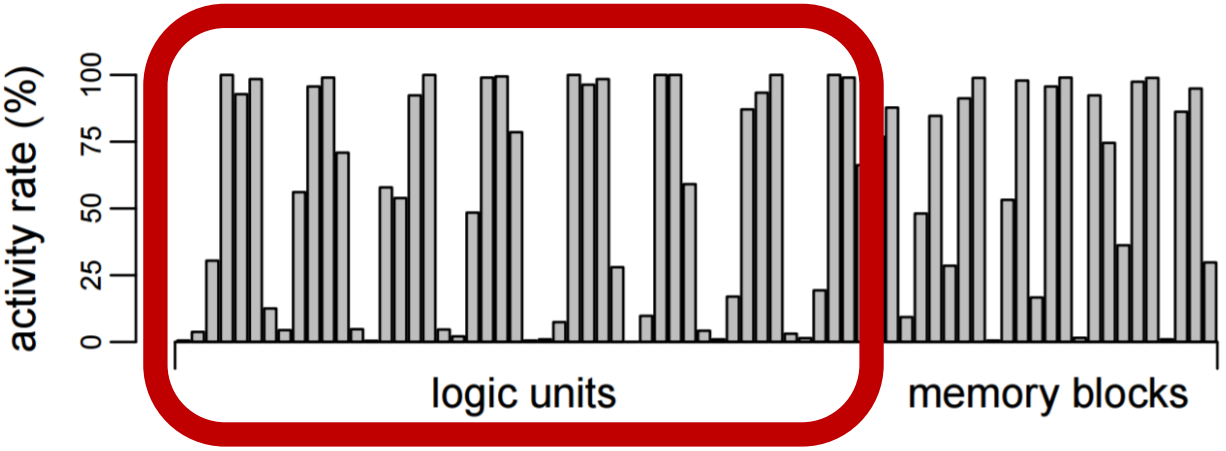

Which Regulators to Turn On?

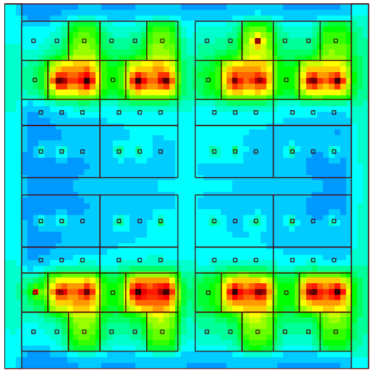

Voltage noise only

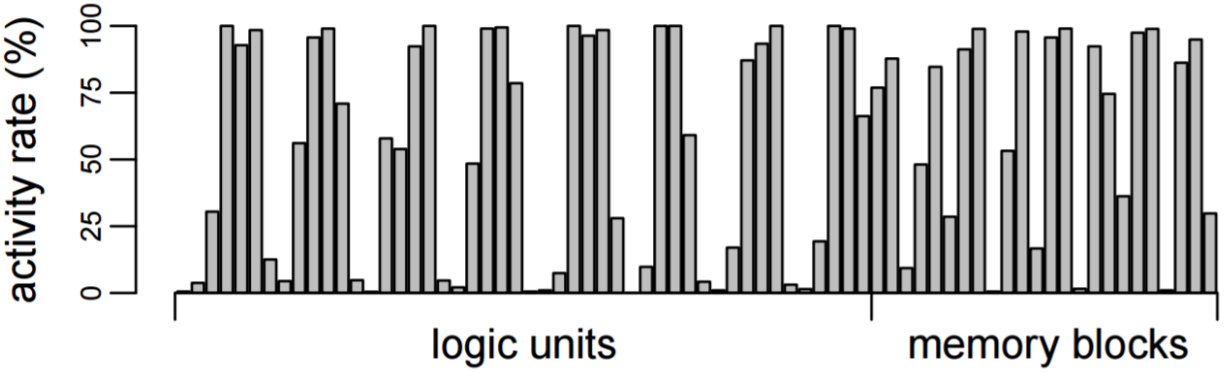

Which Regulators to Turn On?

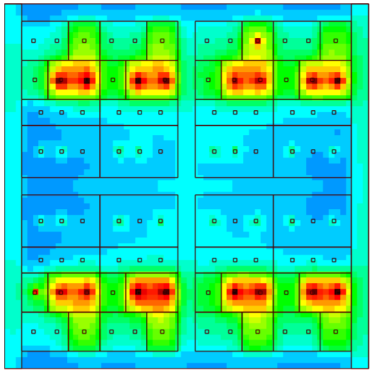

Voltage noise only

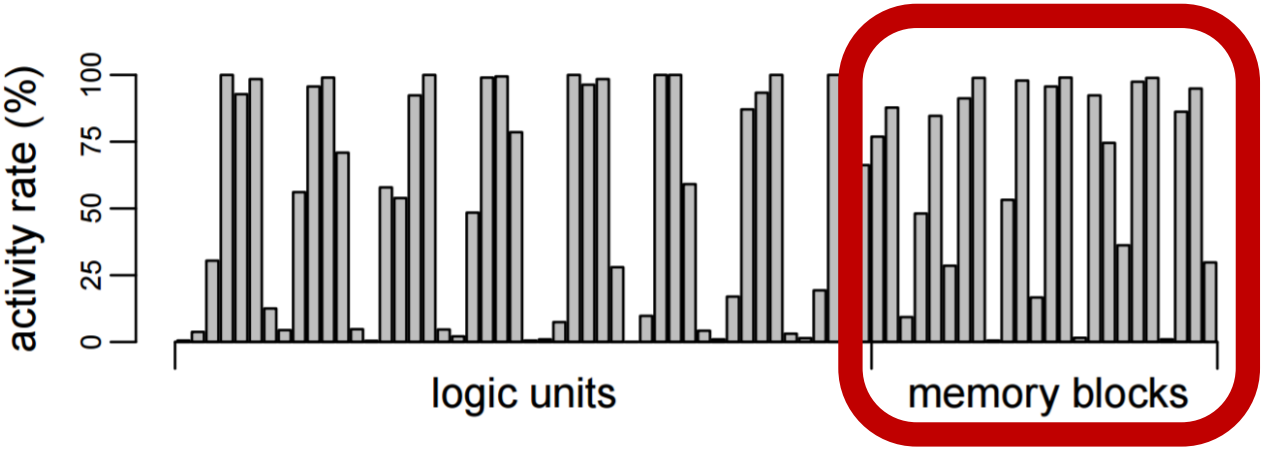

Which Regulators to Turn On?

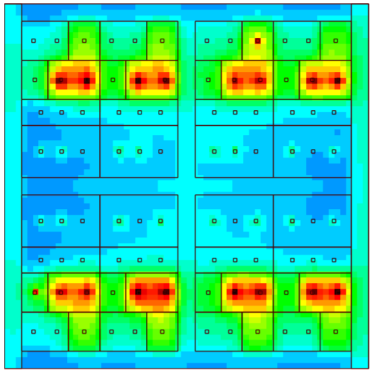

Voltage noise only

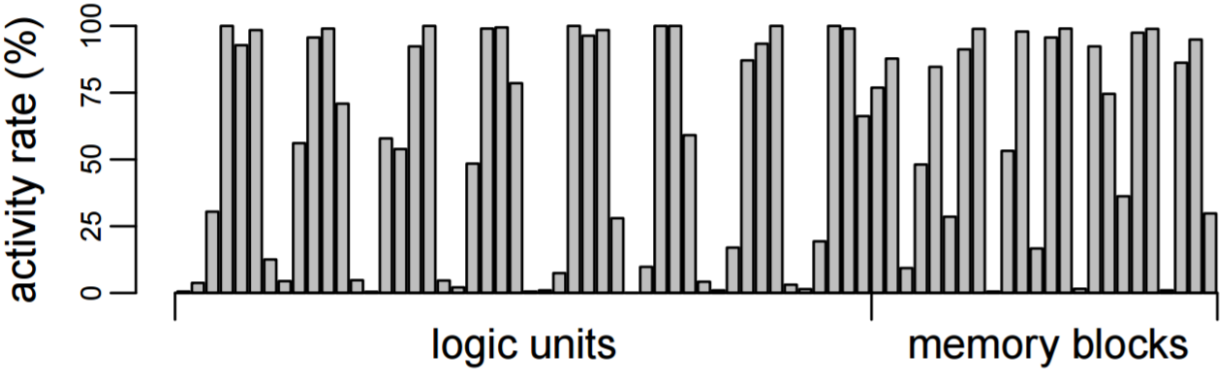

Which Regulators to Turn On?

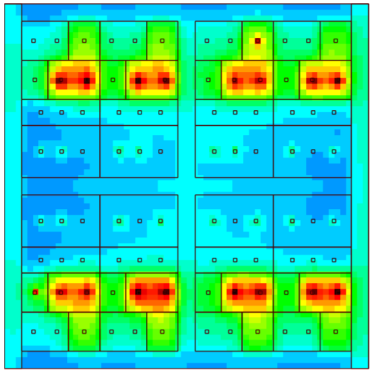

Voltage noise only

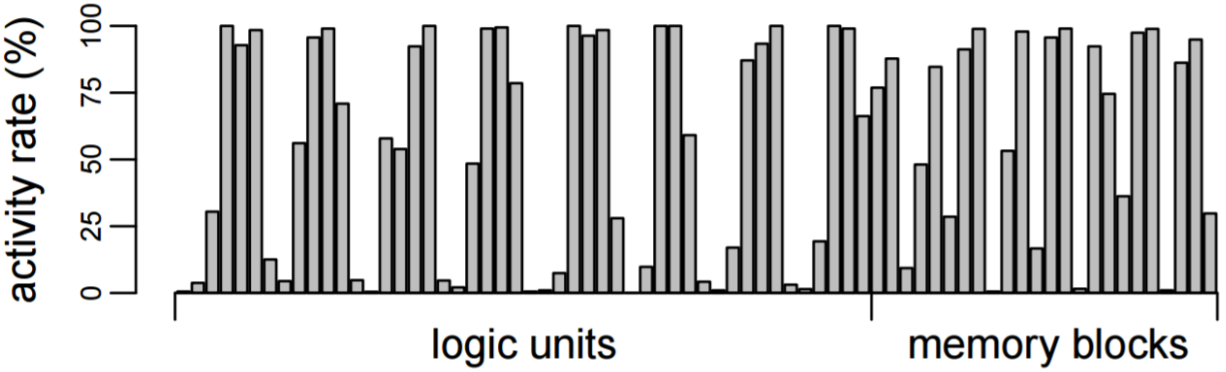

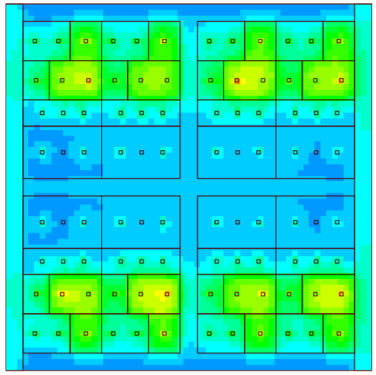

Temperature only

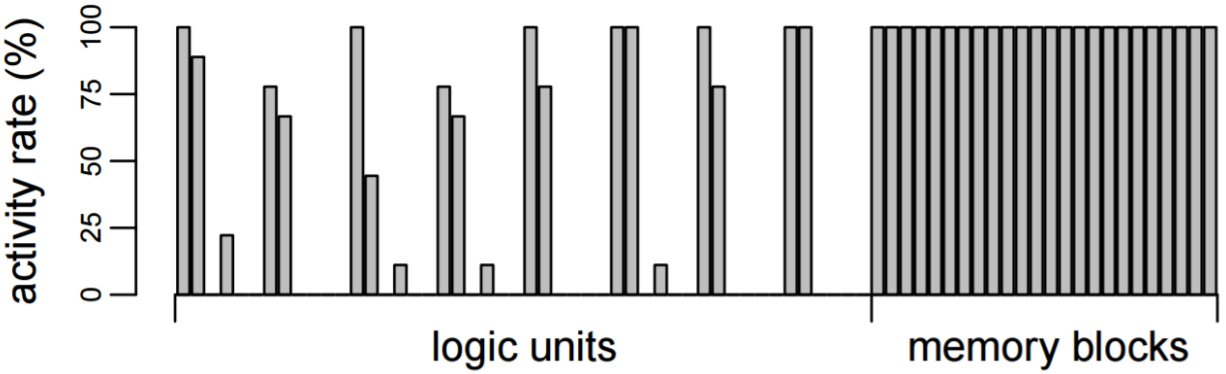

ThermoGater (TG) Policies

- Keep always as many active regulators as required at peak eta: N
  - Track microarchitectural activity
  - Turn more regulators on (off) under high (low) activity

ThermoGater (TG) Policies

- Keep always as many active regulators as required at peak eta: N
  - Track microarchitectural activity
  - Turn more regulators on (off) under high (low) activity
- For a given N, which regulators to select for turning on/off?
  - Constraint: prevent both hotspots **and** voltage emergencies
  - Different ways to enforce this constraint leads to different TG policies

Oracular ThemoGater (TG) Policy

- Assumption: oracular knowledge about
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies

Oracular ThemoGater (TG) Policy

- Assumption: oracular knowledge about
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies

- Observations
  - Voltage emergencies are short (~ ns).
  - Thermal emergencies are long (~ ms).
  - Voltage emergencies are rare

Oracular ThemoGater (TG) Policy

- Assumption: oracular knowledge about
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies

- Observations
  - Voltage emergencies are short (~ ns).
  - Thermal emergencies are long (~ ms).
  - Voltage emergencies are rare

- **Oracular TG Policy**
  - (I) Always mimics temperature-only
  - (II) On a voltage emergency, switches all regulators on

Oracular ThemoGater (TG) Policy

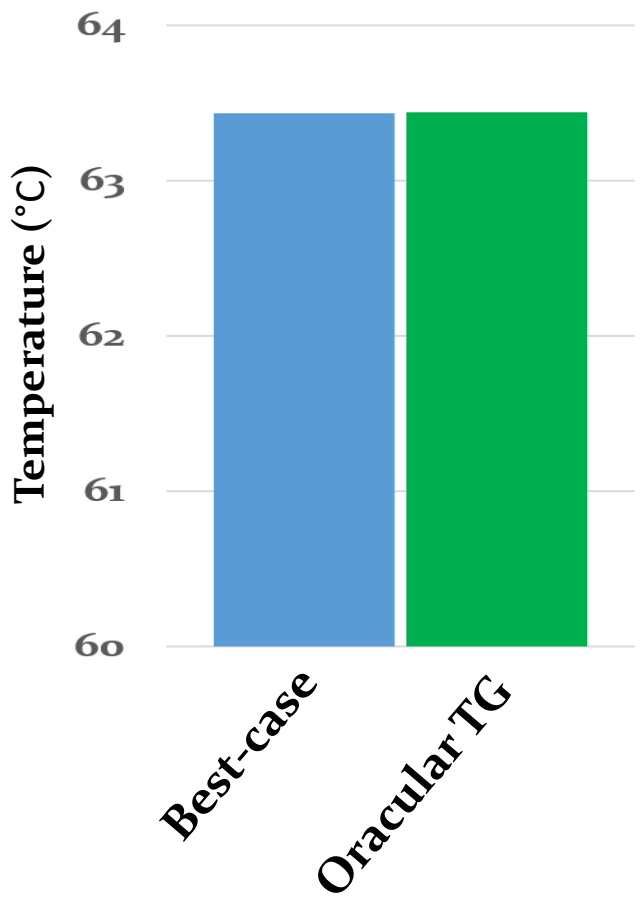

Oracular ThemoGater (TG) Policy

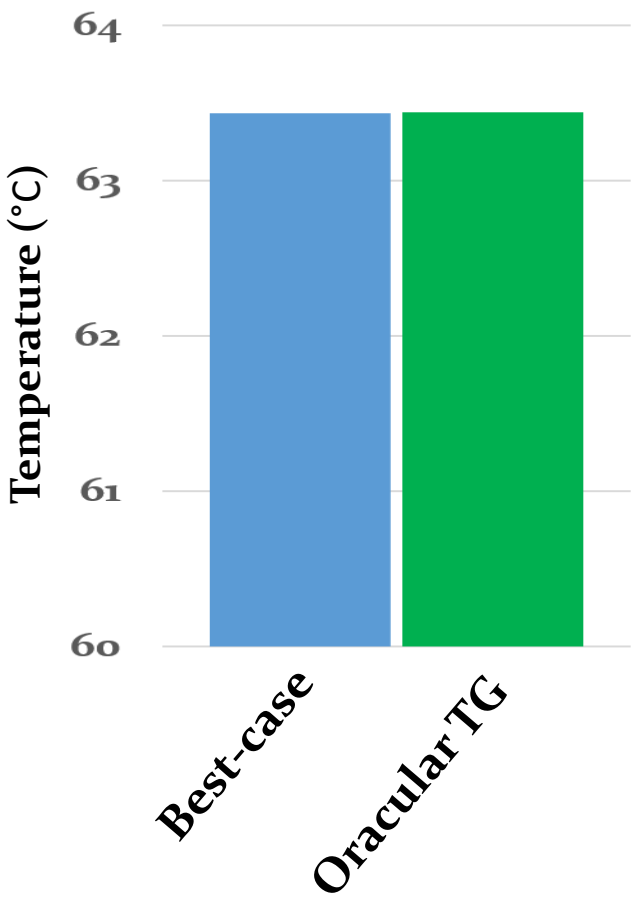

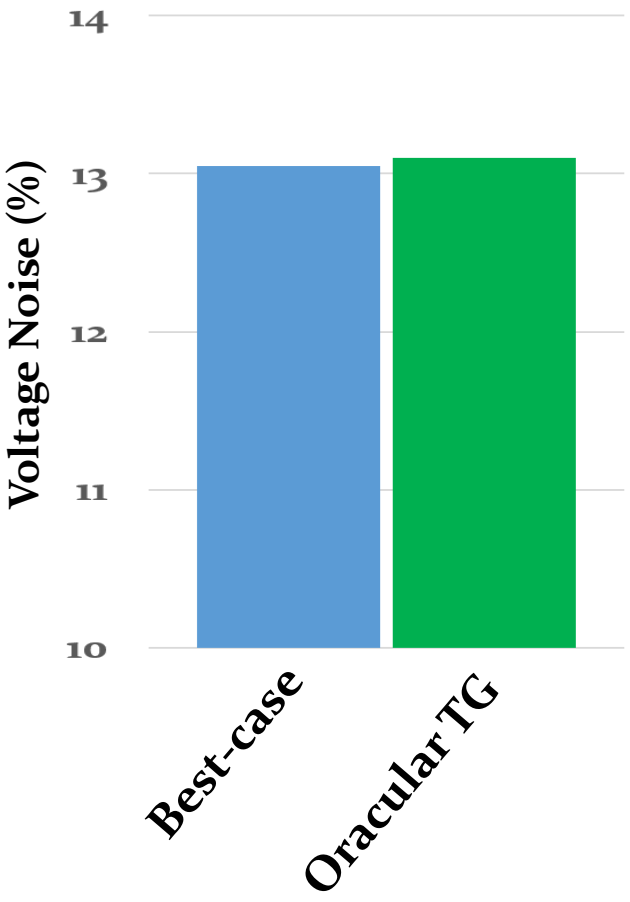

Oracular ThemoGater (TG) Policy

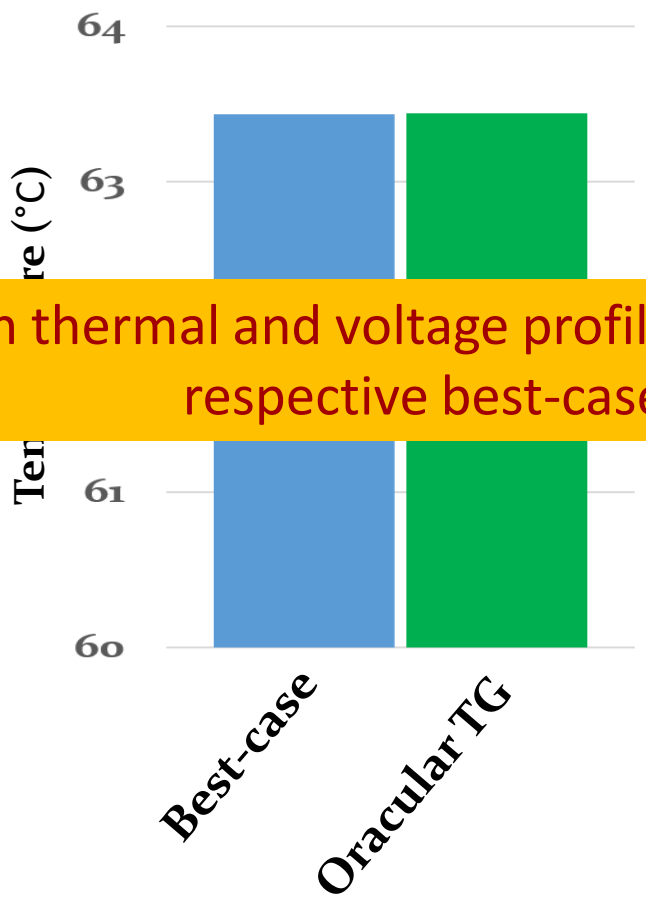

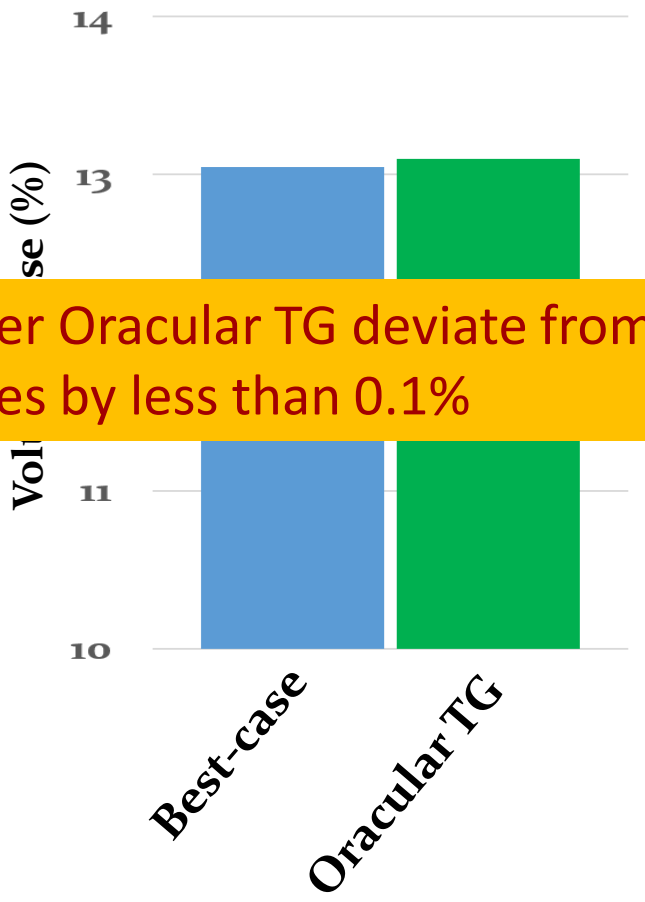

Both thermal and voltage profiles under Oracular TG deviate from the respective best-case profiles by less than 0.1%

Practical ThemoGater (TG) Policy

- Challenge: How to predict
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies

Practical ThemoGater (TG) Policy

- Challenge: How to predict
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies
- Power demand prediction
  - Keep a short history of power demand (at a few previous decision points)
  - Take Weighted Moving Average (of power demand history)

Practical ThemoGater (TG) Policy

- Challenge: How to predict
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies
- Power demand prediction
  - Keep a short history of power demand (at a few previous decision points)
  - Take Weighted Moving Average (of power demand history)
- Temperature prediction
  - Read current temperature from on-chip sensors
  - Use a simple linear model
    - Rank anticipated temperatures at the next decision point

Practical ThemoGater (TG) Policy

- Challenge: How to predict
  - Output power demand
  - Temperature of all regulators under all possible gating decisions
  - Potential voltage emergencies
- Power demand prediction
  - Keep a short history of power demand (at a few previous decision points)
  - Take Weighted Moving Average (of power demand history)
- Temperature prediction
  - Read current temperature from on-chip sensors
  - Use a simple linear model
    - Rank anticipated temperatures at the next decision point
- Voltage emergency detection
  - Deploy a predictive per-core voltage emergency detector

Practical ThemoGater (TG) Policy

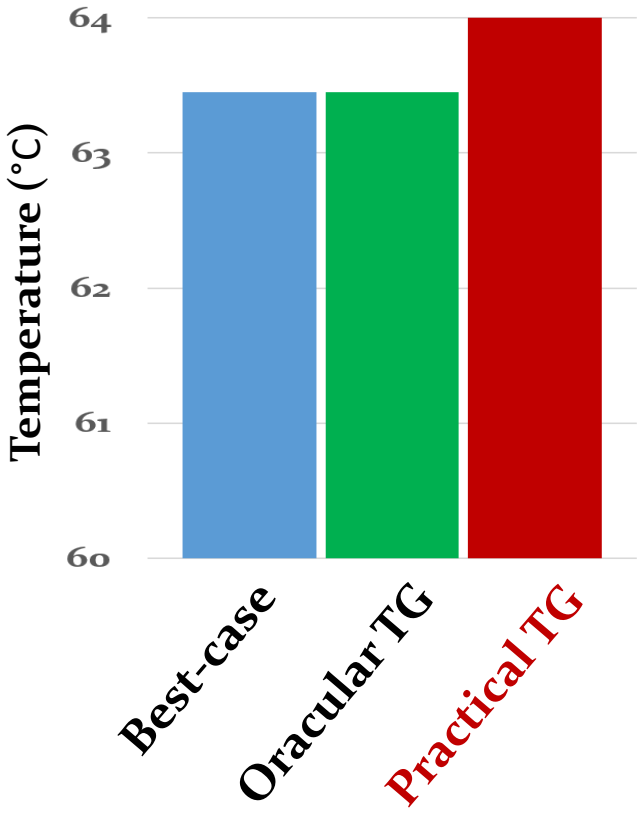

Practical ThemoGater (TG) Policy

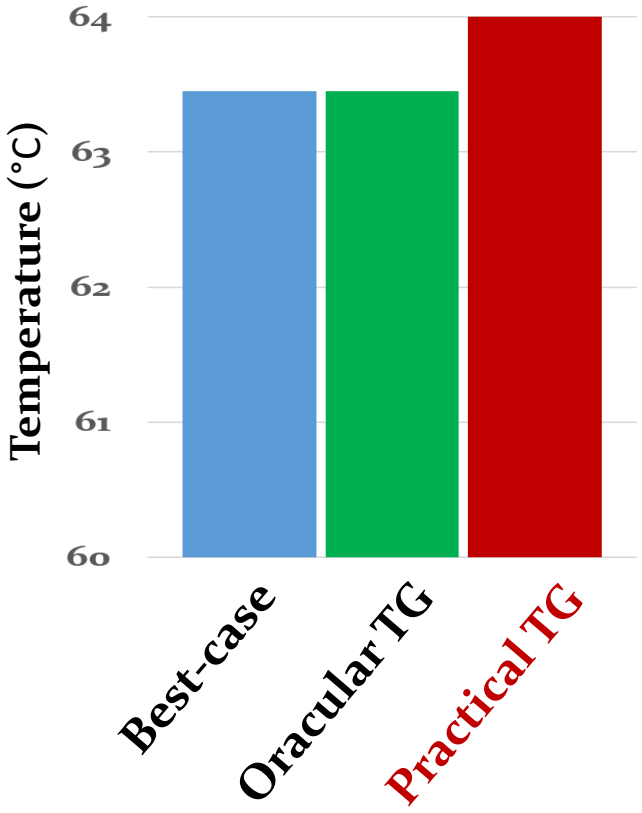

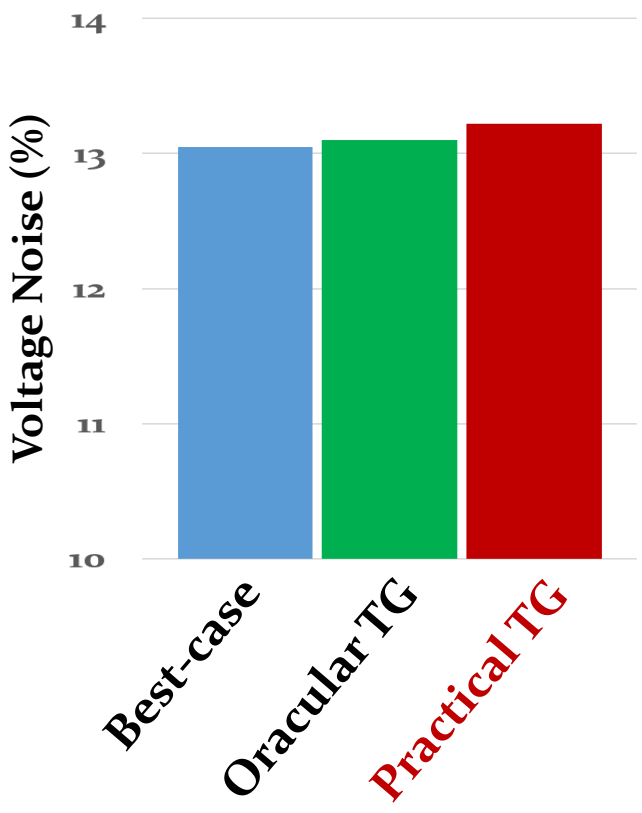

Practical ThemoGater (TG) Policy

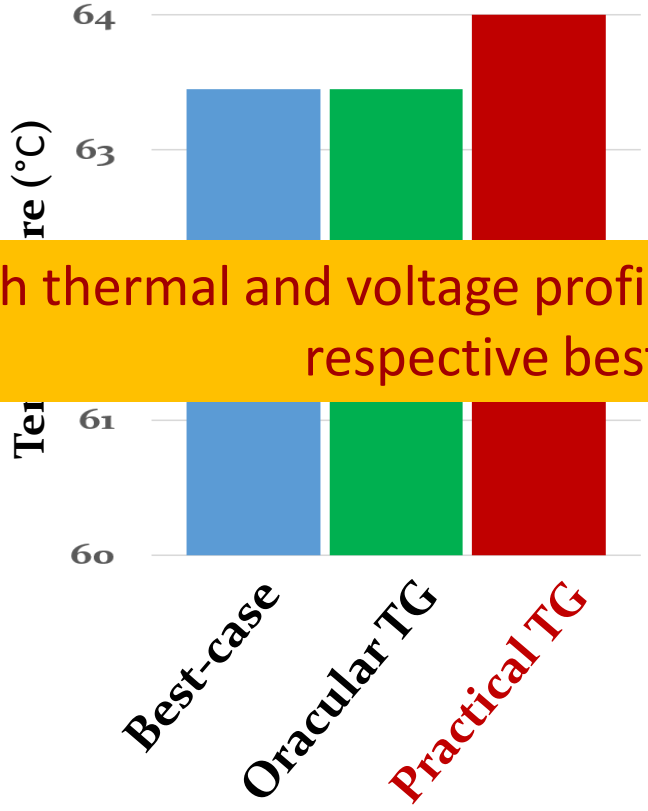

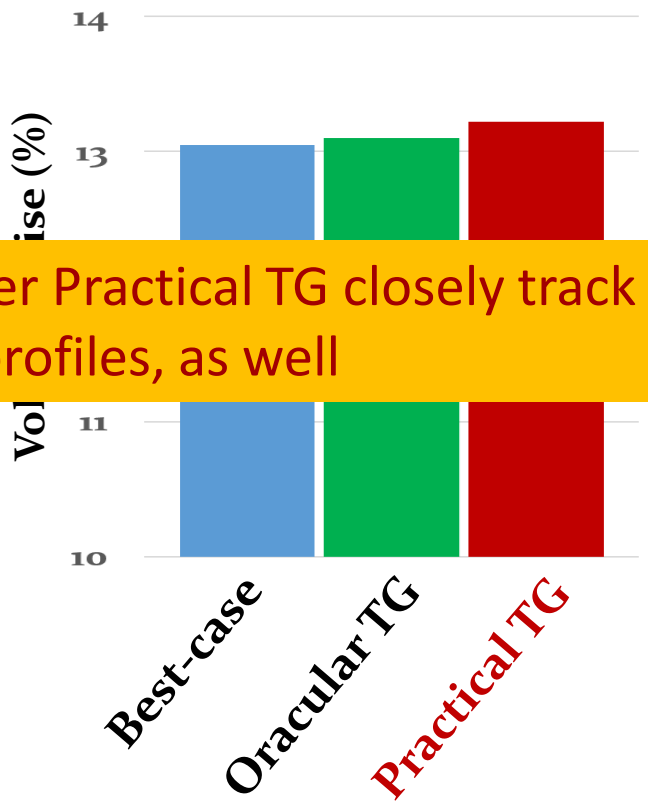

Both thermal and voltage profiles under Practical TG closely track the respective best-case profiles, as well

Impact on Aging

- Utilization per regulator is not uniform throughout execution
- Higher regulator utilization near cooler regions such as memory
  - TG mimics temperature-only policy by default
  - Periodic gating decision interval is based on temperature
  - Gating based on voltage is event-driven
- Aging rate increases with both utilization and temperature
  - Higher utilization near cooler regions likely to balance out aging

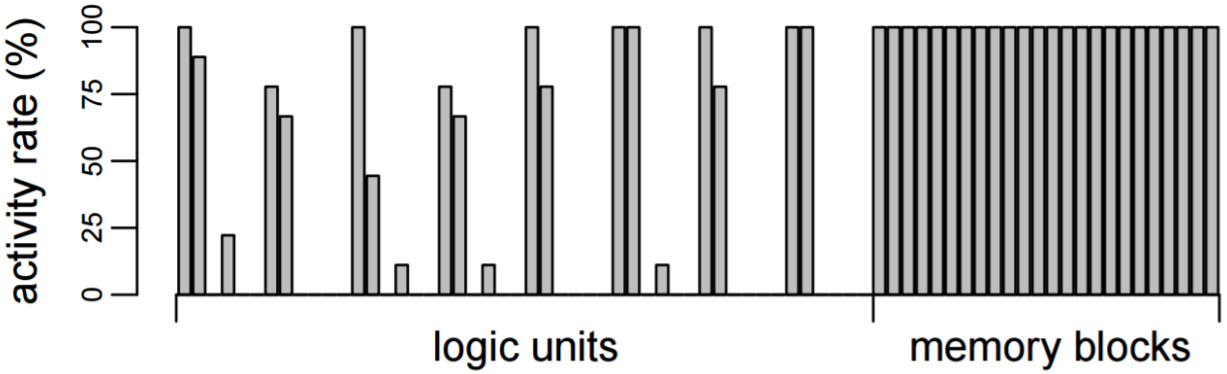

Conclusion

- ThermoGater
  - An architectural governor for practical, thermally-aware regulator gating
  - Sustains operation at peak power conversion efficiency
  - Mitigates regulator-induced thermal emergencies
  - Considers the impact on voltage noise

Conclusion

- ThermoGater
  - An architectural governor for practical, thermally-aware regulator gating
  - Sustains operation at peak power conversion efficiency
  - Mitigates regulator-induced thermal emergencies
  - Considers the impact on voltage noise
- Practical ThermoGater policies can
  - Sustain operation at 1% of the peak power conversion efficiency
  - Keep the temperature only 0.6°C higher than the best-case thermal profile
  - Keep the voltage noise only 0.2% higher than the best-case voltage profile

Thanks!

For questions or feedback, please contact

khatami@umn.edu

ukarpuzc@umn.edu

ThermoGater: Thermally-Aware On-Chip Voltage Regulation

S. Karen Khatamifard, Longfei Wang†, Weize Yu†, Selçuk Köse†, Ulya R. Karpuzcu

University of Minnesota
{khatami, ukarpuzc}@umn.edu

†University of South Florida
{longfei, weizeyu}@mail.usf.edu kose@usf.edu

26/6/2017



UNIVERSITY OF MINNESOTA

UNIVERSITY OF SOUTH FLORIDA