

A Simple Mechanism to Adapt Leakage-Control Policies to Temperature

Stefanos Kaxiras

Polychronis Xekalakis

Georgios Keramidas

xekalakis@apel.ee.upatras.gr

{kaxiras,keramidas}@ee.upatras.gr

Introduction

- Substantial rise in heat density
- Increasing need for temperature aware micro architectures
- Deep sub-micron/nano era: leakage control necessary
- Leakage greatly depends on temperature →
We propose temperature-adaptive leakage control policies for on-chip caches

Outline

- Why Thermally Adaptive Leakage Control Techniques?
- Cache Decay, Drowsy and Hybrid Schemes
- Our Proposal: Thermally Aware Hybrid Mech.
- Temperature Sensitive Timer: 4T DRAM Cell
- Evaluation of Hybrid vs. Decay and Drowsy
- Summary - Future Work

Why do we Need Thermally Adaptive Leakage-Saving Techniques?

- Static consumption becomes substantial
 - Leakage control
 - Caches primary target: They account for a large fraction of the 'T' budget
- Some mechanisms trade static for dynamic
- Dominant leakage contributor (subthreshold leakage) depends **exponentially** on temperature
- "Stationary" static/dynamic trade-off is **clearly not optimal**

Main Idea

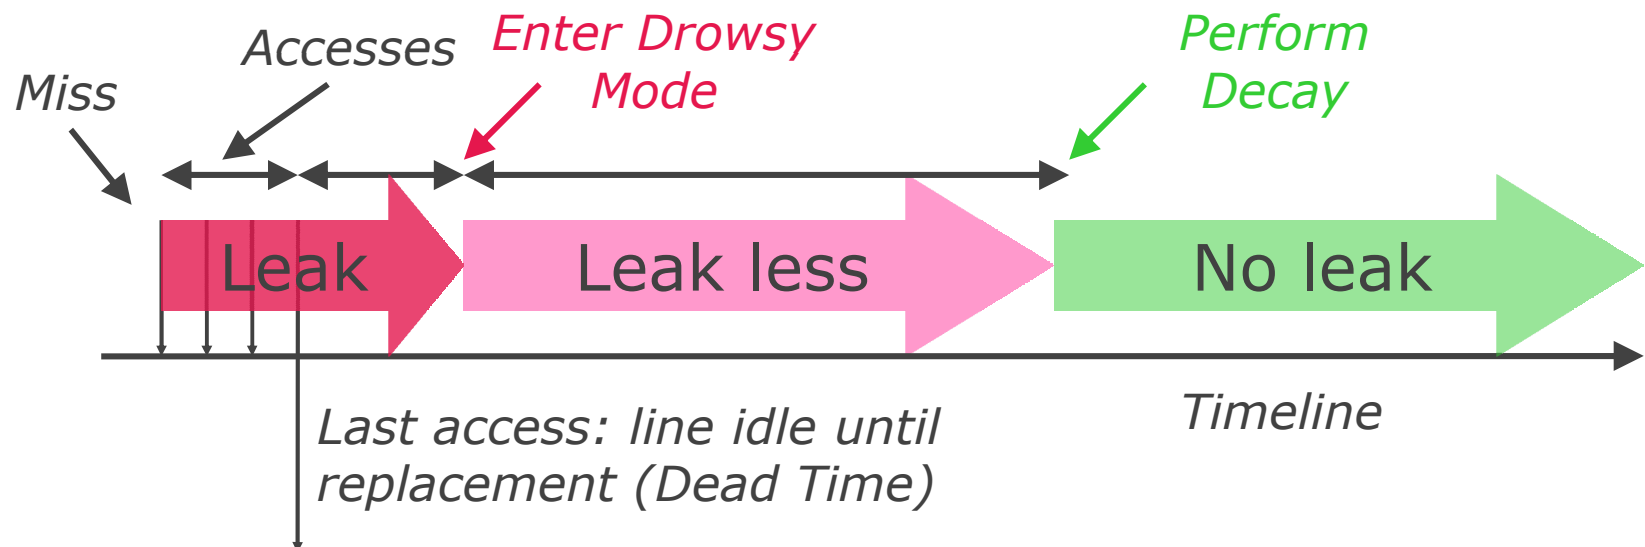
- Gated approach (cache decay) to save leakage:
 - Saves lots of leakage, destroys data
 - Dynamic power penalty if wrong (miss)
- Tradeoff (save leakage but incur misses) changes with temperature!
 - High T (high leakage) →
 - aggressive gating, don't mind misses
 - Low T (low leakage) →
 - very careful gating, misses hurt
 - also use DVS!

Cache Decay and Drowsy Cache

- **Cache Decay**: wait a “decay interval” after the last access and gate the cache line
- Decay interval defines behavior:
 - “decay” the cache lines **soon** → increase in dynamic power
 - “decay” the cache lines **late** → might not save enough leakage
- **Drowsy Cache**: DVS in cache lines
 - leakage savings smaller compared to decay
 - but: no dynamic power penalty
- Hybrid approach captures the benefits of both

Hybrid Mechanism

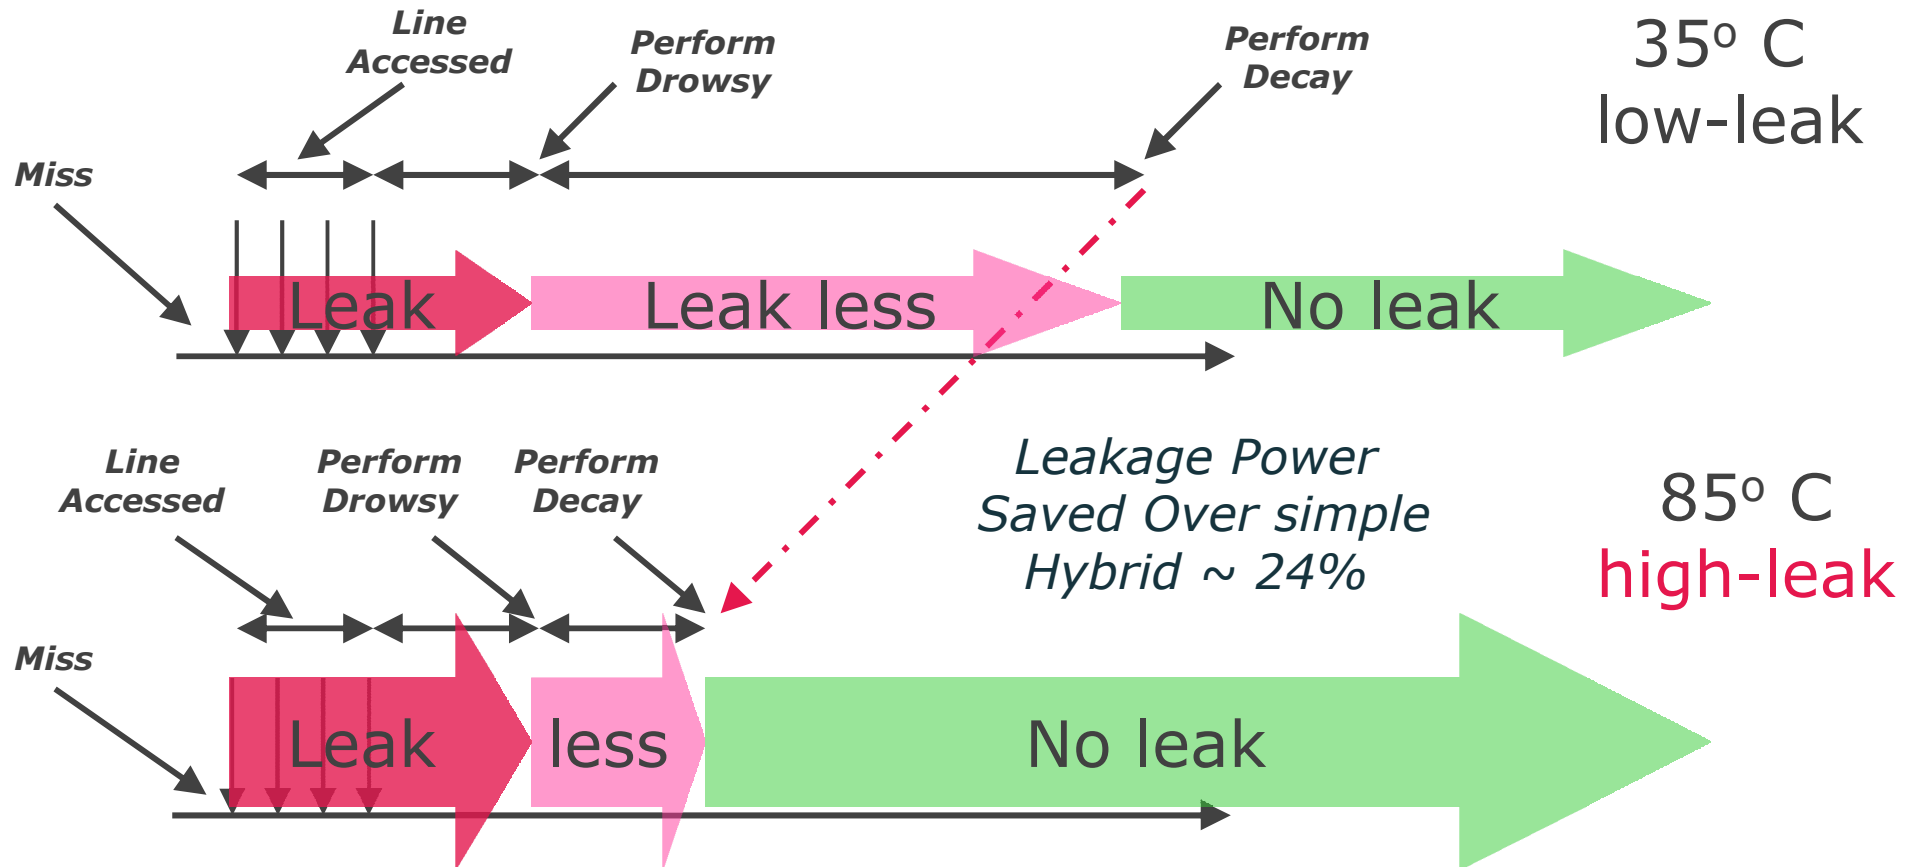
- If line idle for a few cycles we put it in drowsy mode
- If a drowsy line remains idle for a longer period we decay it



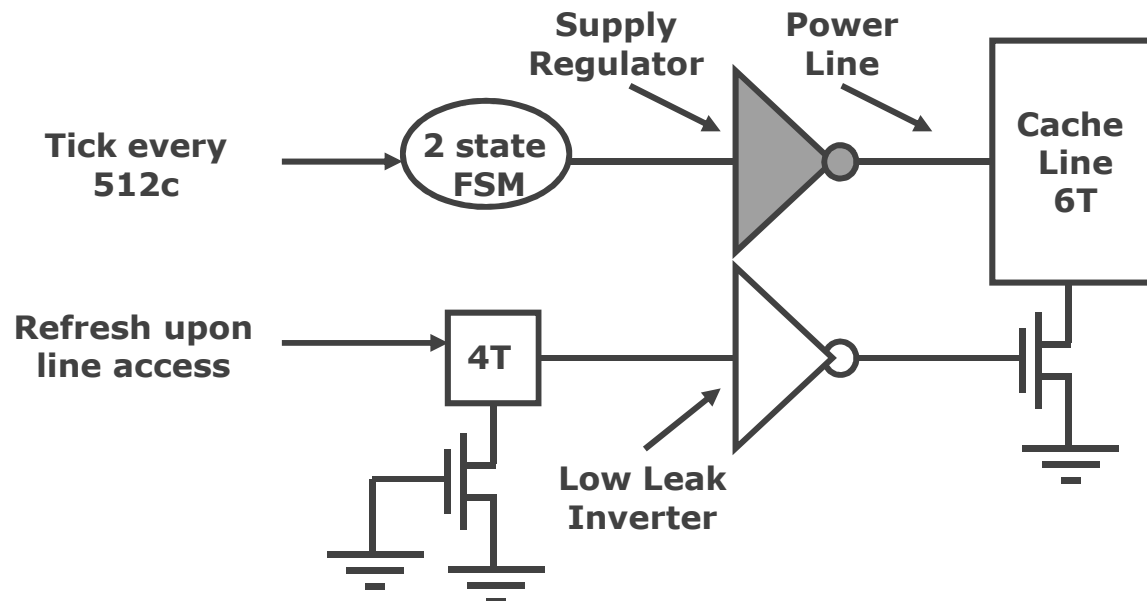
- Drowsy exploits DI
- But: Is this enough?

Hybrid Thermal Aware Mechanism

- We keep the drowsy interval const & vary the decay interval according to temperature




Timing Mechanism



- Drowsy timer: 2-state FSM, reset on access, ticked by global timer (512c)
 - Cache line idle 512c → drowsy mode
- Decay timer: **temperature-sensitive** 4T cell,
 - charged on access
 - “decays” after some temperature-depended decay-interval → gates cache line

4T Cell's Timing Nature



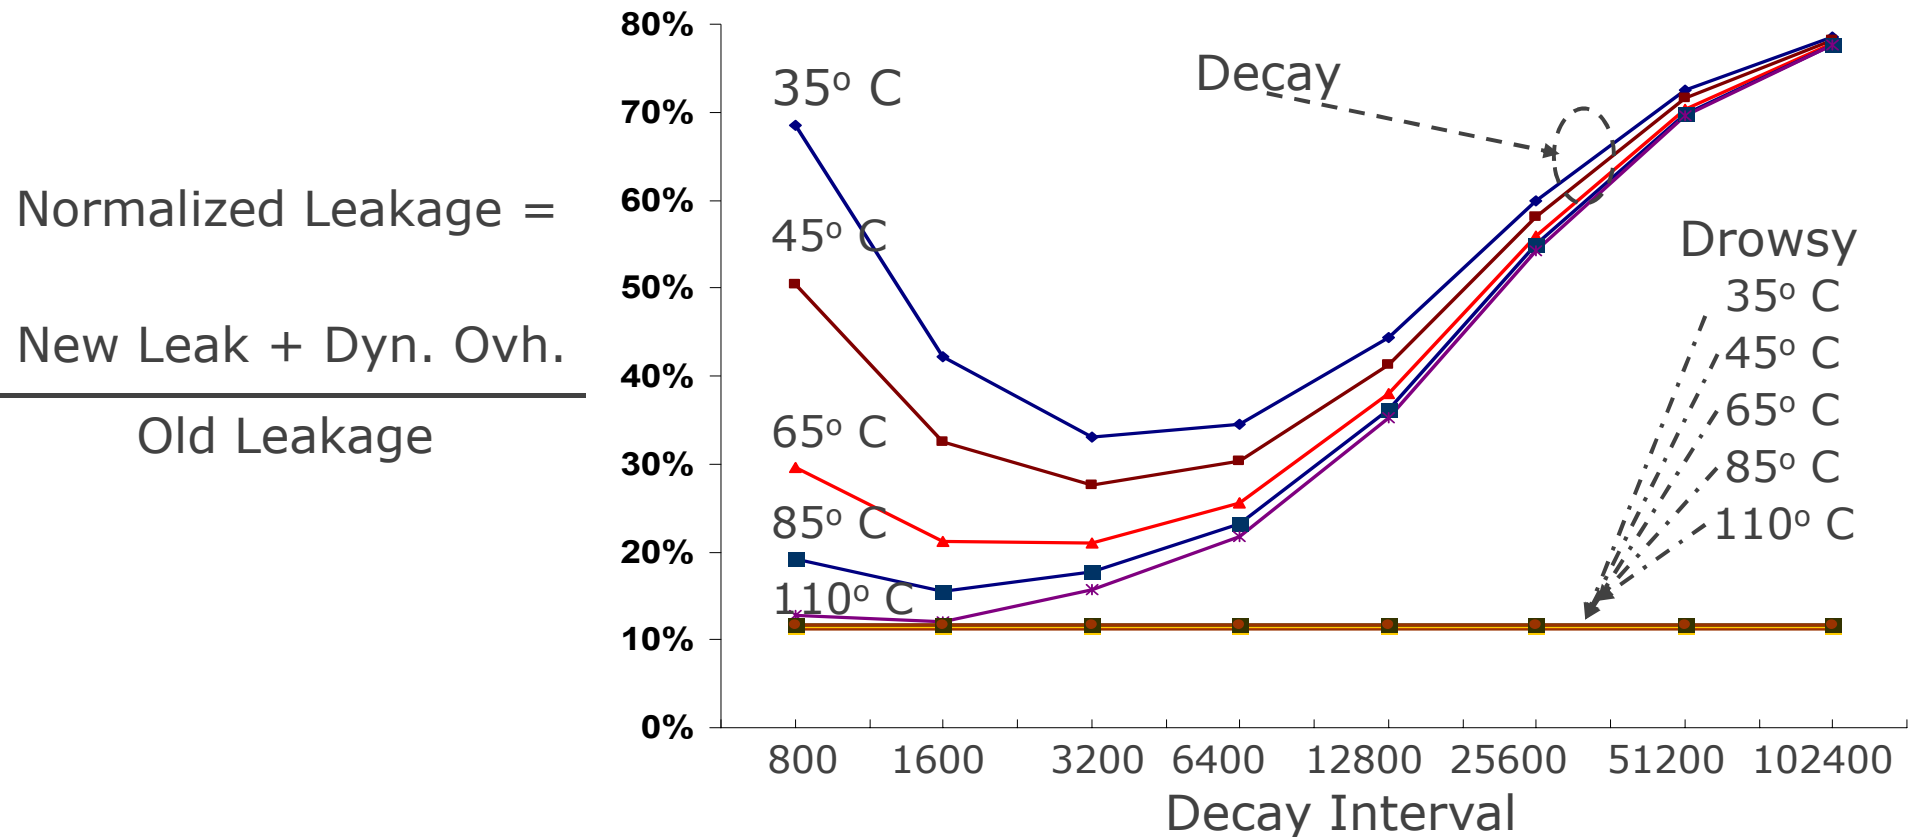
- 4T cell's *retention* time function of temperature
 - Previously proposed as **cheap temperature sensors**
- Can "design" retention times
 - gate transistors, geometry
- Automatic adaptation of retention times with T
- Retention times converge to the same value

This is what we want!

Hybrid-Decay-Drowsy, Which is Better?

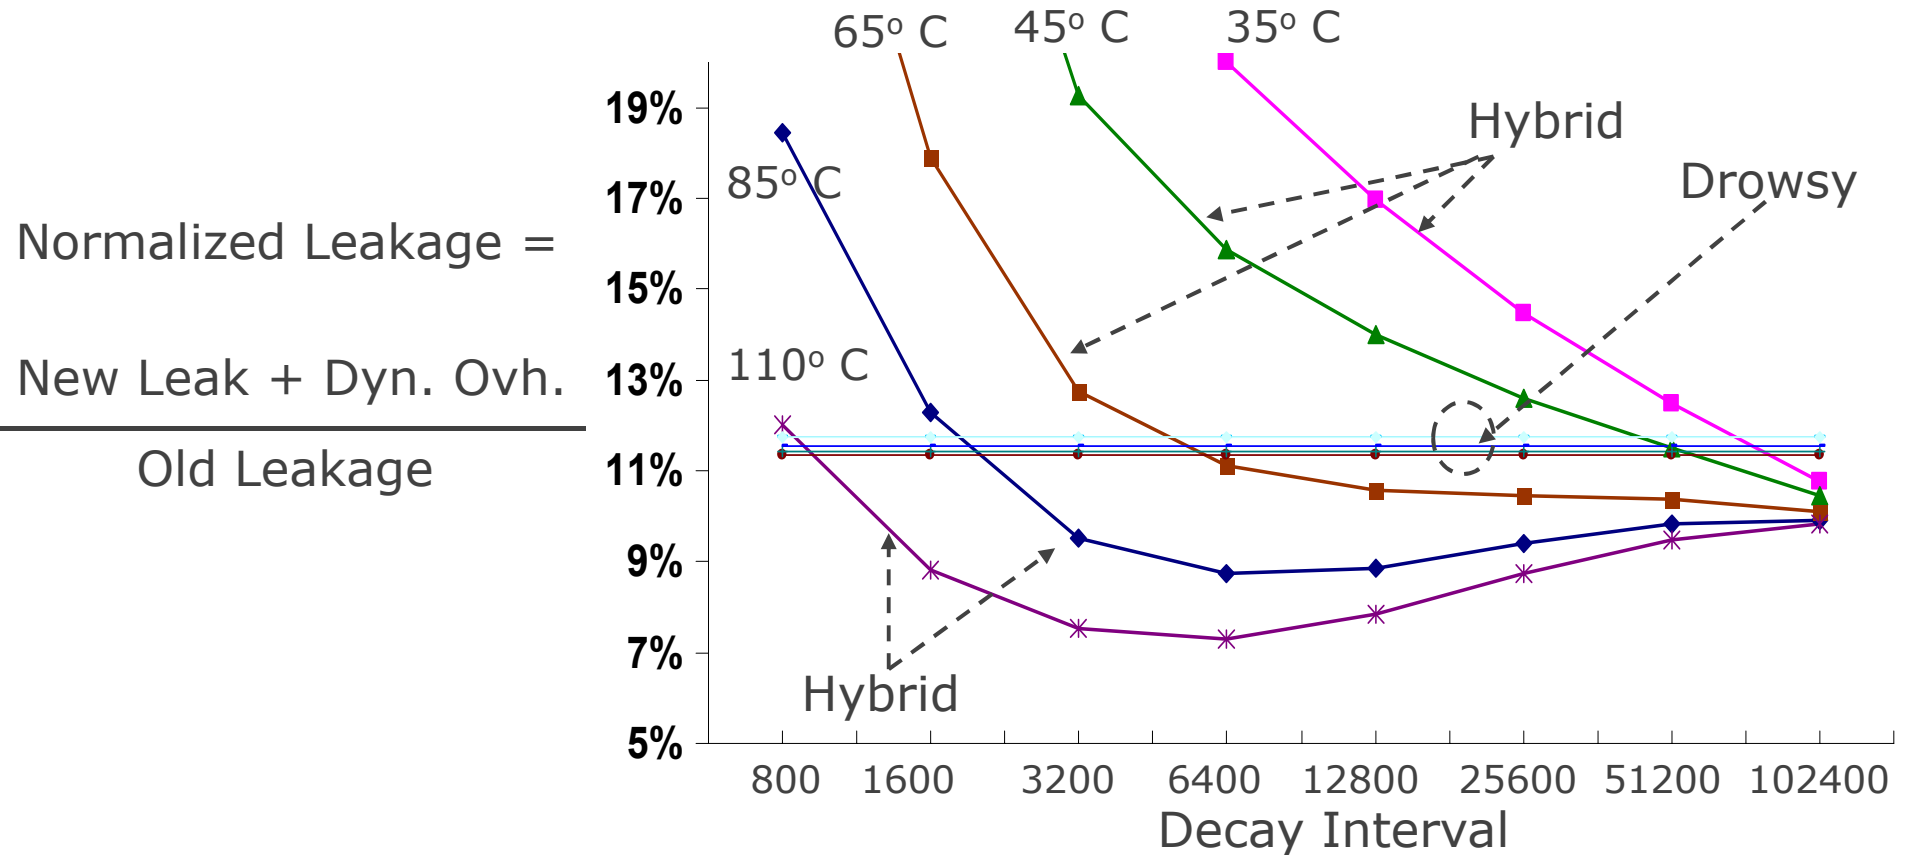
- HotLeakage, 6 Spec, PSpice
- Overall savings depend on:
 - L2 latency \rightarrow increases relative energy delay
 - \nearrow in latency favors drowsy
 - Tox \rightarrow reduces leakage
 - \searrow in Tox favors decay
- If decay is better than drowsy **hybrid is by default better than both**
- Parameters chosen to favor drowsy

Evaluation: Decay & Drowsy vs. Temp.



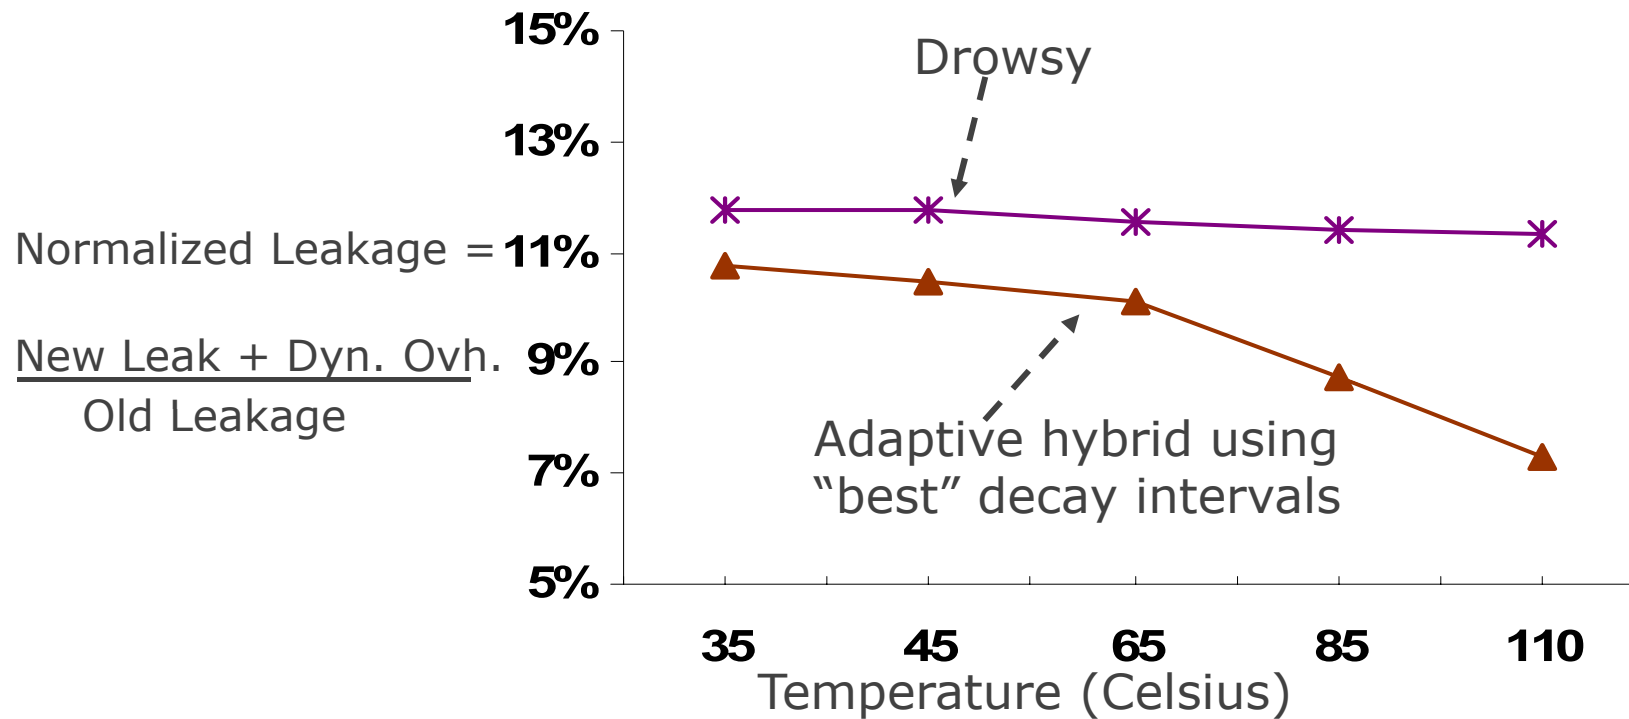
- Drowsy → flat, does not depend on T nor DI
- Decay → depends leak/dynamic ratio, affected by T (**leakage**) and DI (**dynamic**)

Evaluation: Hybrid vs. Drowsy

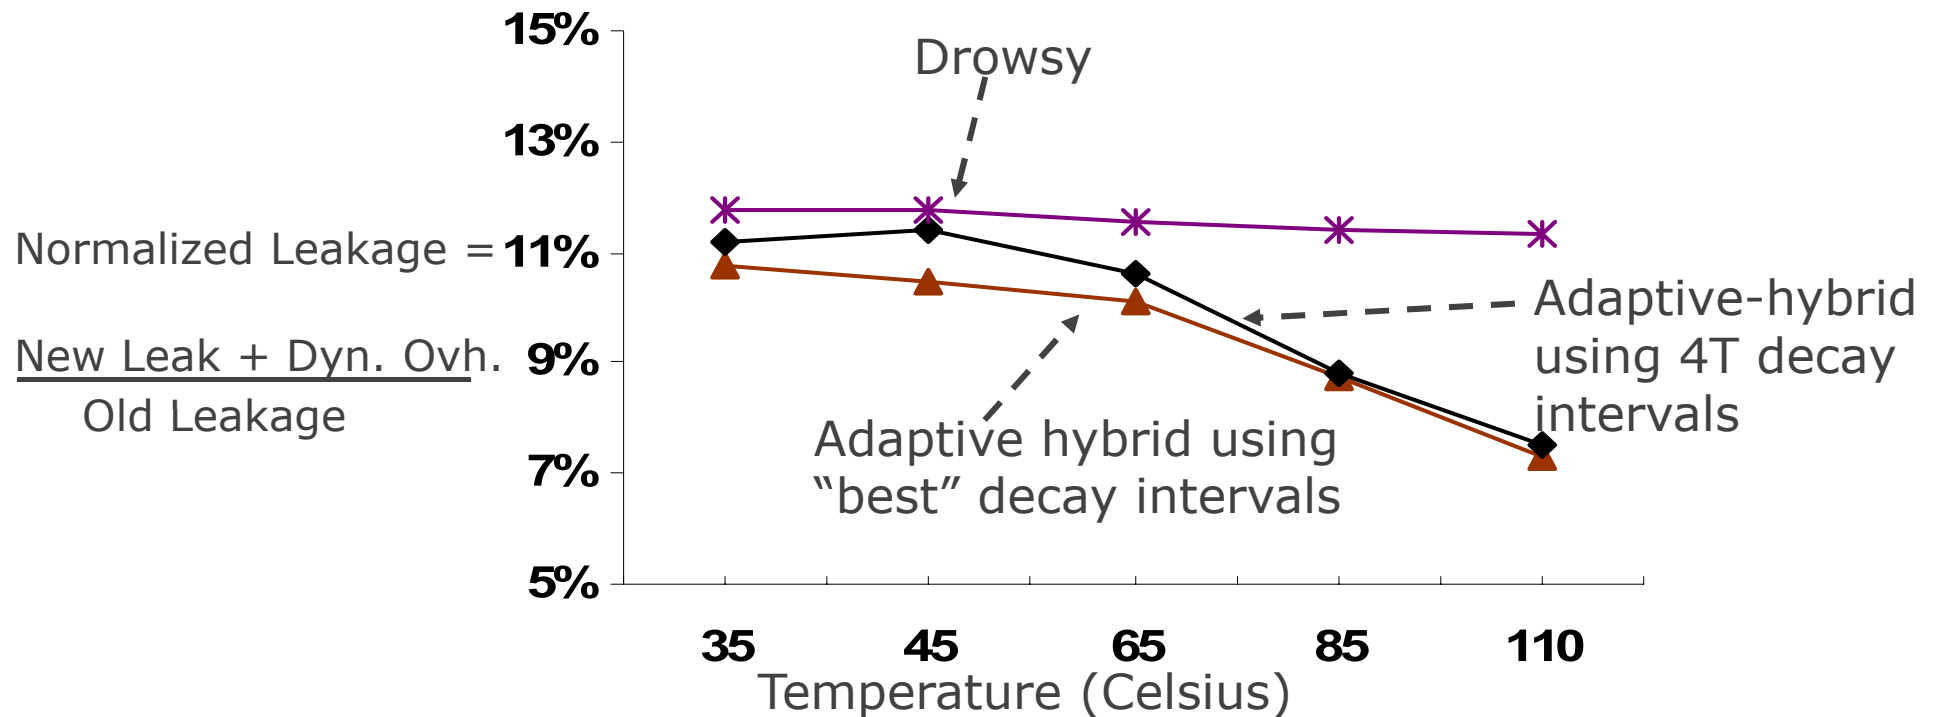


- Hybrid easily outperforms drowsy
- Hybrid adaptivity: simply chose the “best” decay intervals @ each temperature

Approximation of ideal hybrid with 4T



Approximation of ideal hybrid with 4T



- We approximate the ideal hybrid with the 4T decay mechanism
- Easily superior to drowsy, especially for high temps

Summary

- Hybrid Temperature Adaptive Leakage Control:
 - decay+drowsy
 - decay intervals vary upon temperature
 - High temperatures: aggressive decay
 - Low temperatures: employ drowsy much more, to keep the dynamic power low
- 4T decaying cells as timers to control decay intervals
- Outperform best of non-temperature adaptive schemes by as much as 33% @ high T

Future Work

- Thermal DVS
- Thermally Dynamic Threshold Transistors
- Varying Drowsy also, under some constraints may produce good results
- Reducing power means reducing Temperature
→ quantify the temperature reduction achieved by hybrid
- Model impact of sensor errors