

ReplayCache: Enabling Caches for Energy Harvesting Systems

Jianping Zeng¹, Jongouk Choi¹, Xinwei Fu²,
Ajay Paddayuru Shreepathi³, Dongyoon Lee³,
Changwoo Min², Changhee Jung¹

¹Purdue University

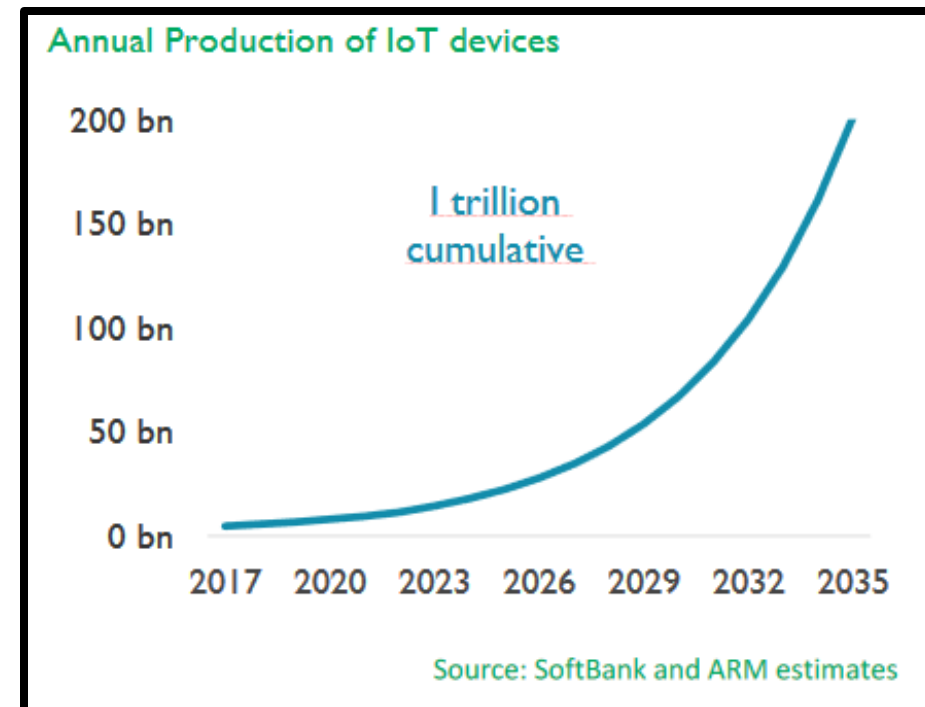
²Virginia Tech

³Stony Brook University


Battery is Not the Way to Go!



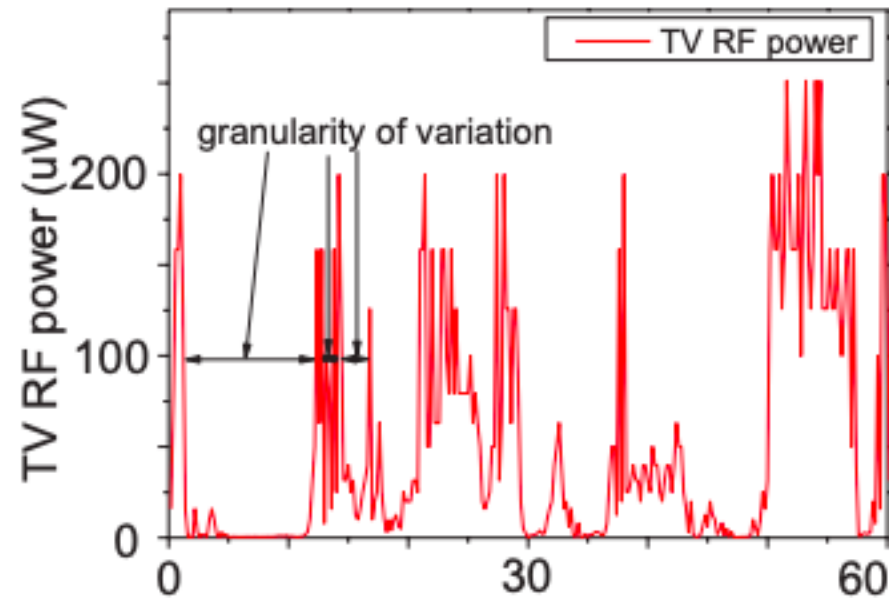
- Batteries are bulky
- They must be replaced



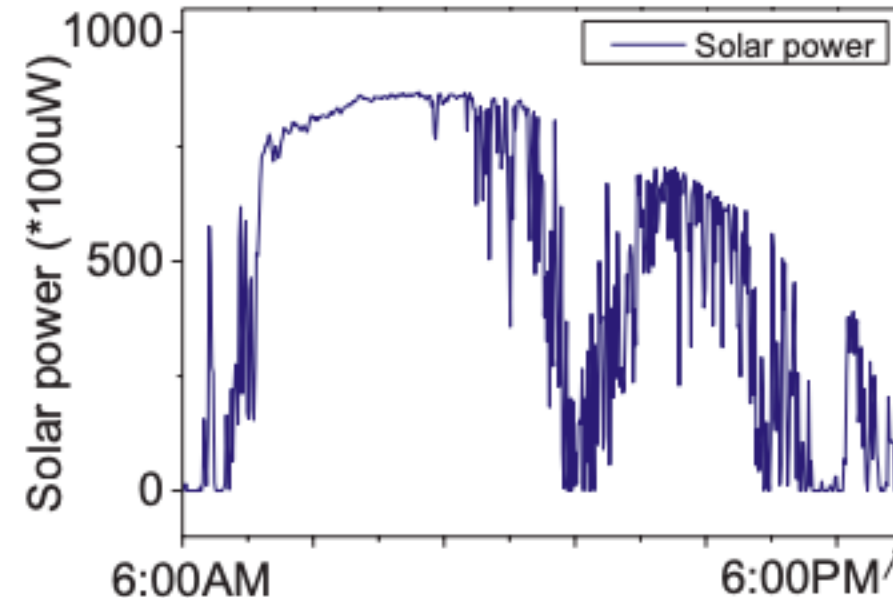
Energy Harvesting Systems (EHS)



Unreliability of Ambient Energy Sources

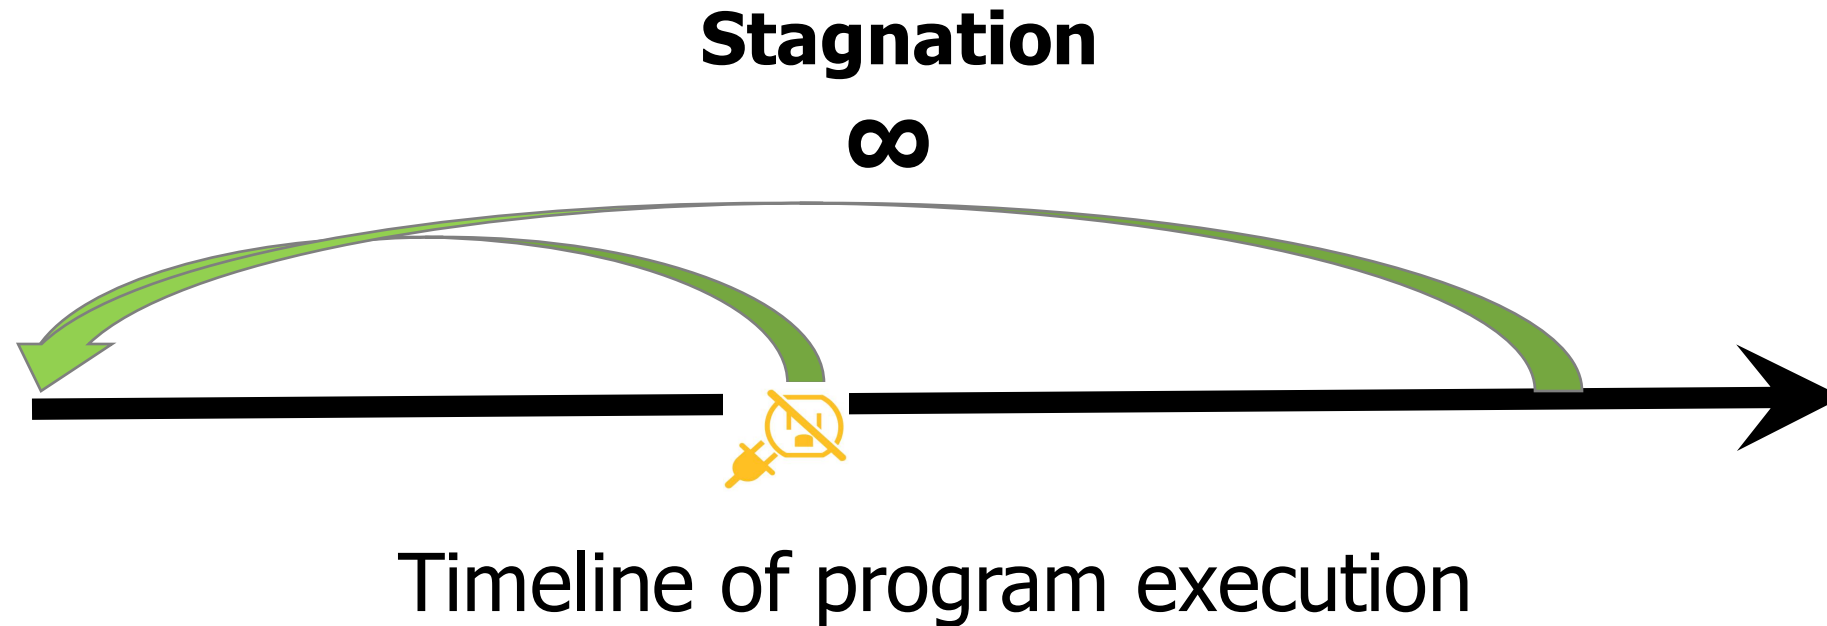


(a). Time (s) (Sample time: 0.33us)




(b). Time (s) (Sample time: 0.33us)

[Ma,HPCA'2015]




 : Power failure on which **all volatile registers lose their data**

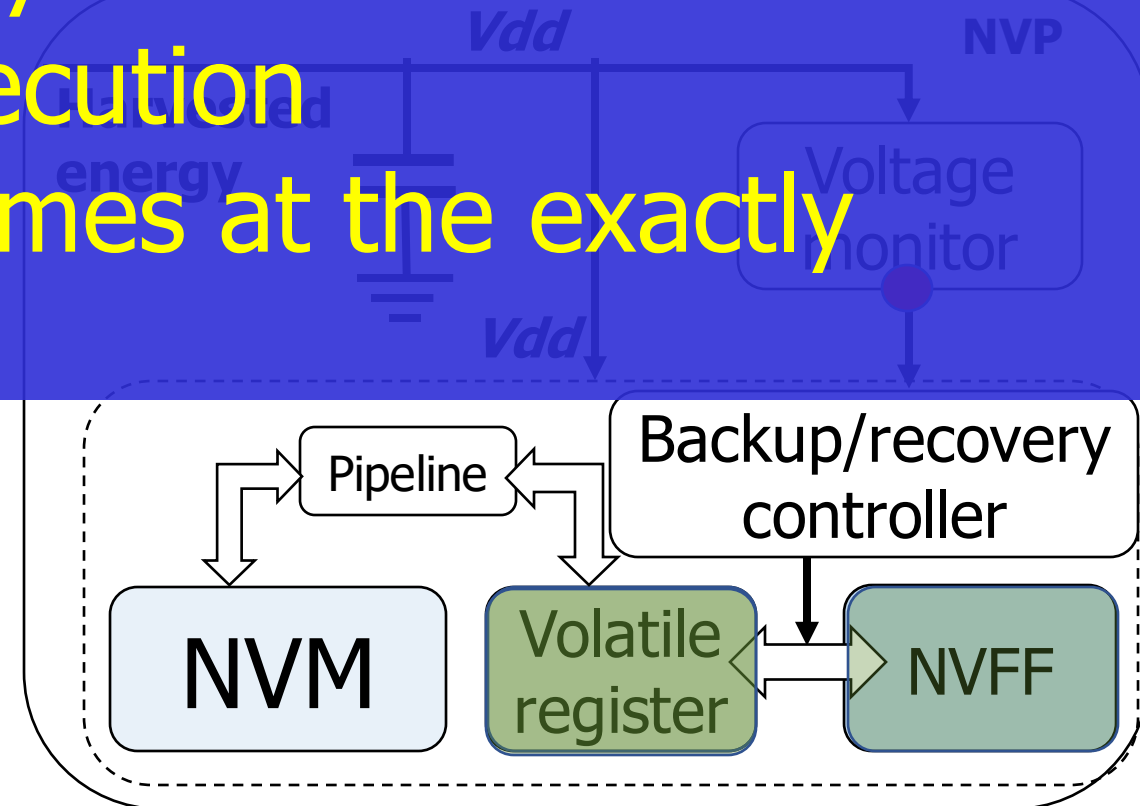
NVP: Just-in-time Register Checkpointing



- Register access is not delayed because of no NVFF access during program execution
- No-rollback: program resumes at the exactly failure point

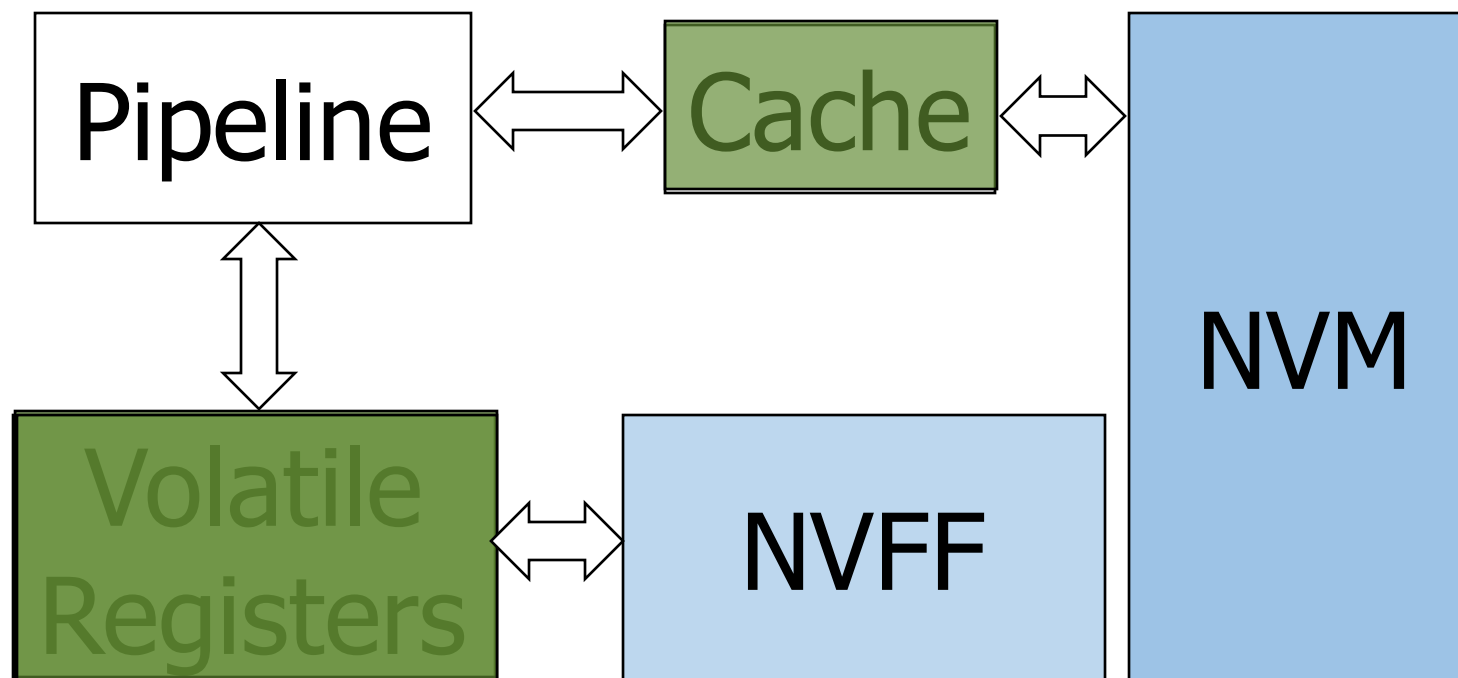


[Ma,HPCA'2015]




NVP Has No Cache Due to Its Crash Consistency Issue

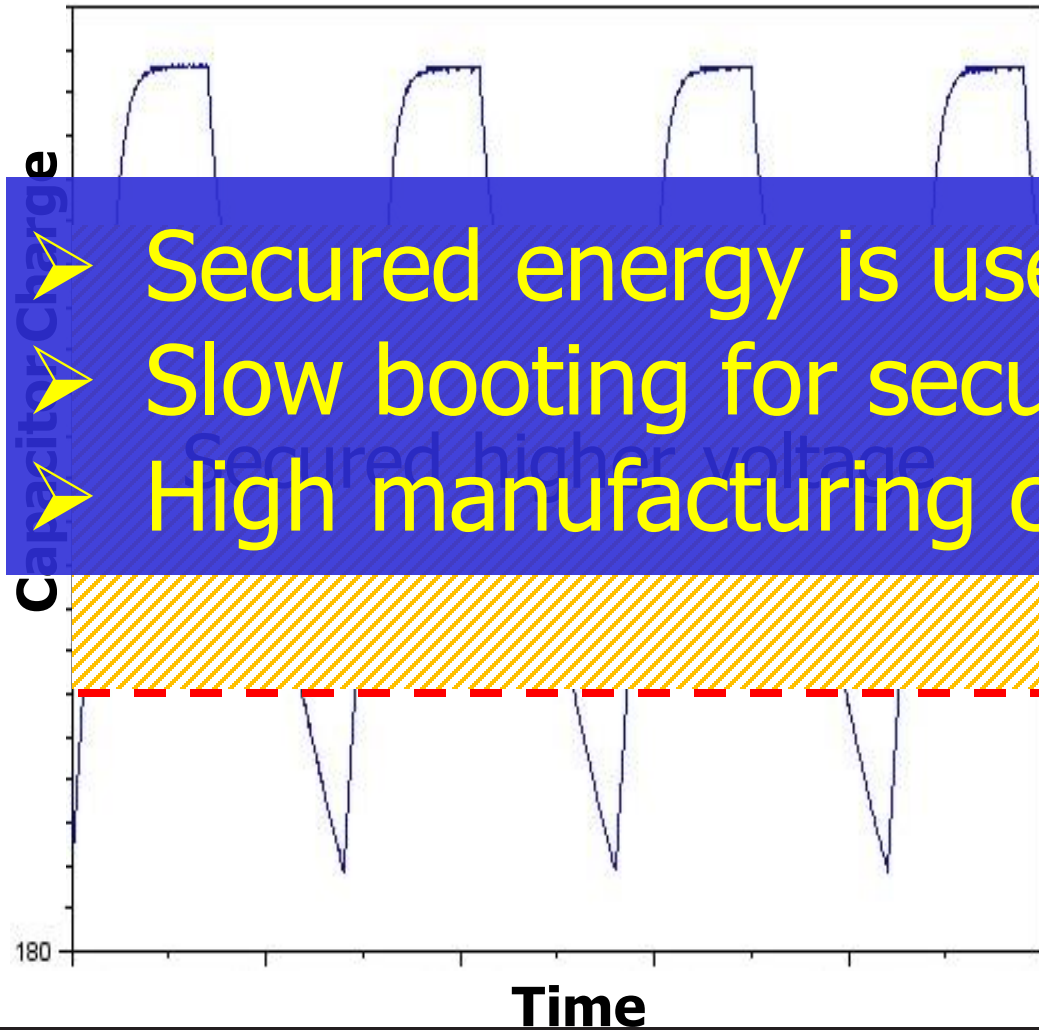
* We can't restart program from the exactly failure point as NVP does



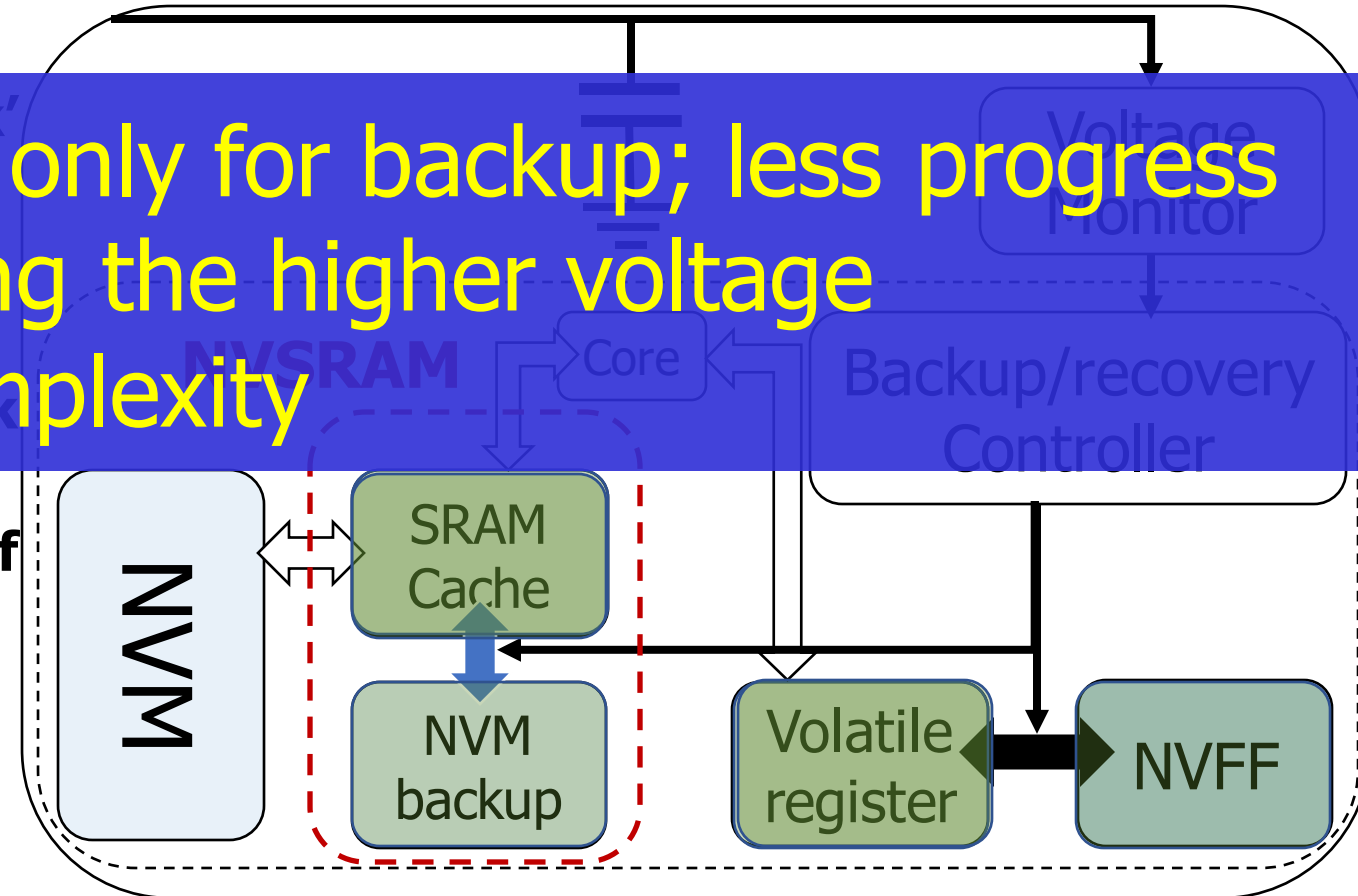
NVSRAM: The State-of-the-Art and its Limitations



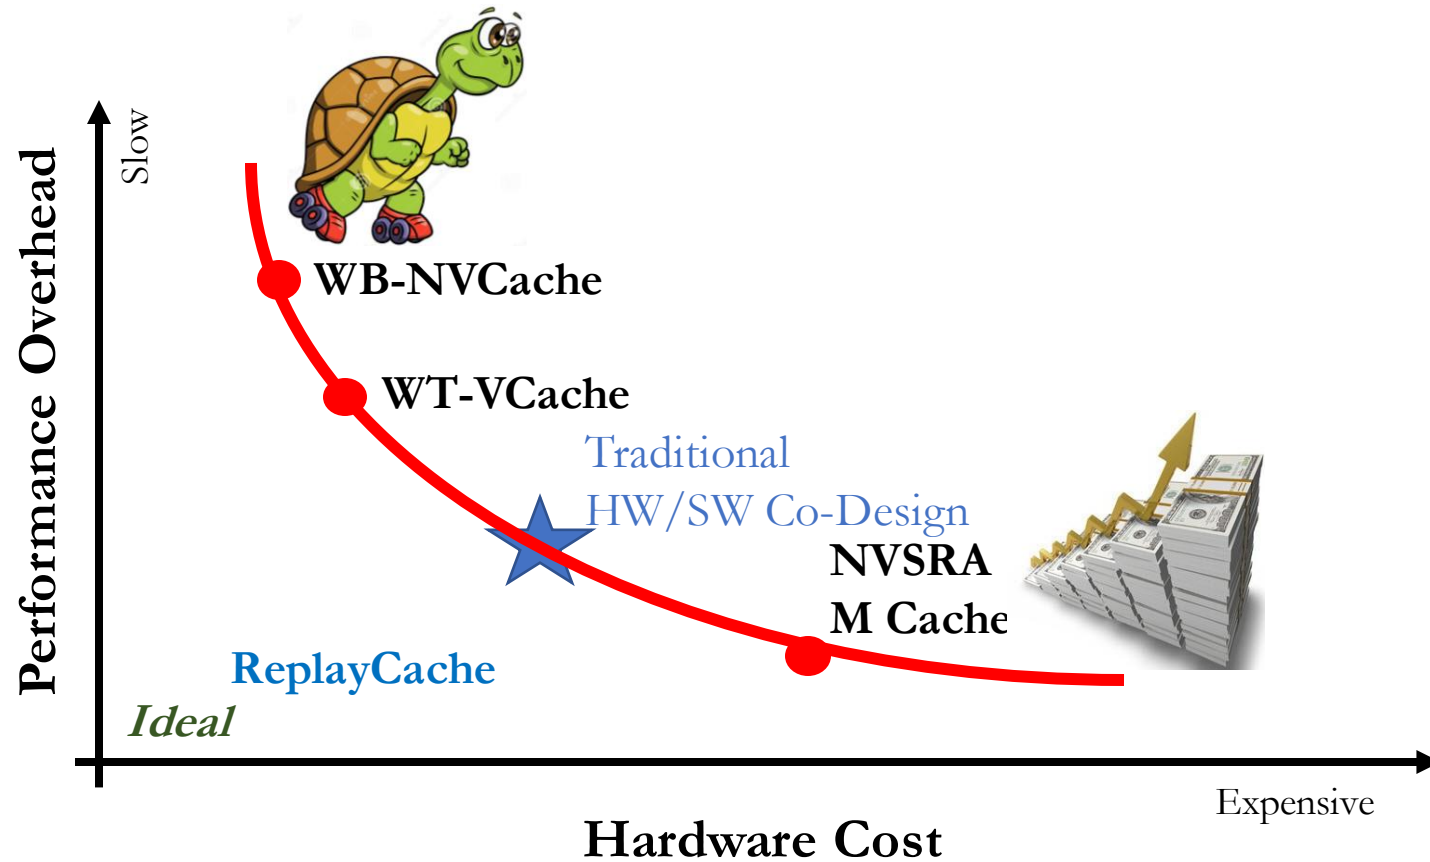
[Chui, VLSI'2010]



- Secured energy is used only for backup; less progress
- Slow booting for securing the higher voltage
- High manufacturing complexity



ReplayCache: A Pure Software Solution to Enabling Performant Volatile Cache for EHS




* Software undo/redo logging slow down 1.6-5x

Reason for Crash Inconsistency of Volatile Cache

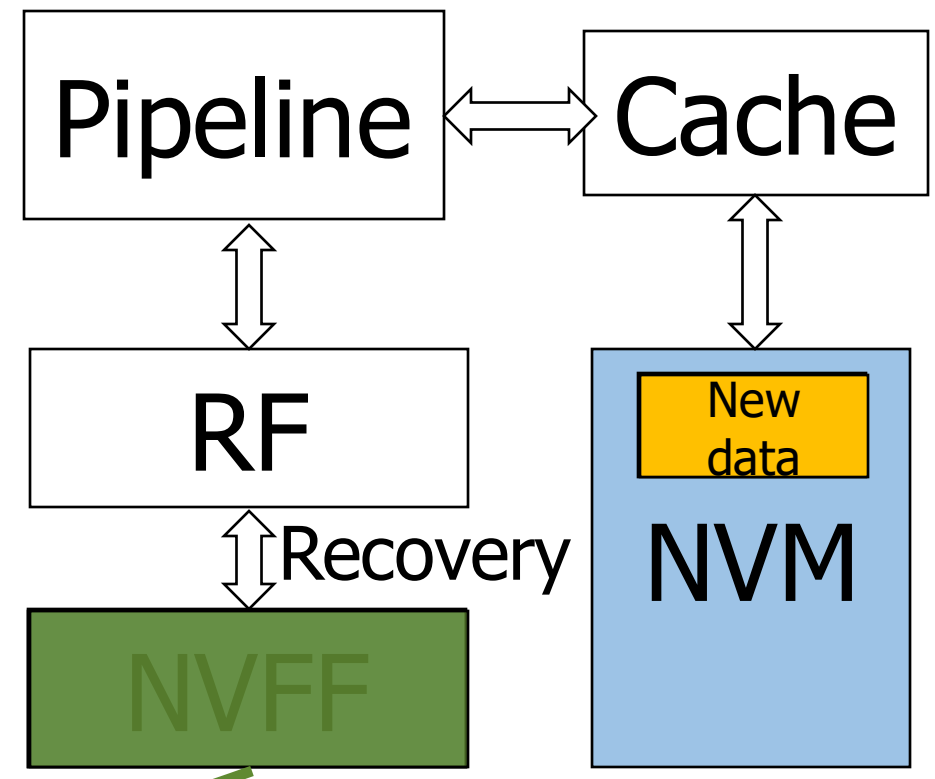
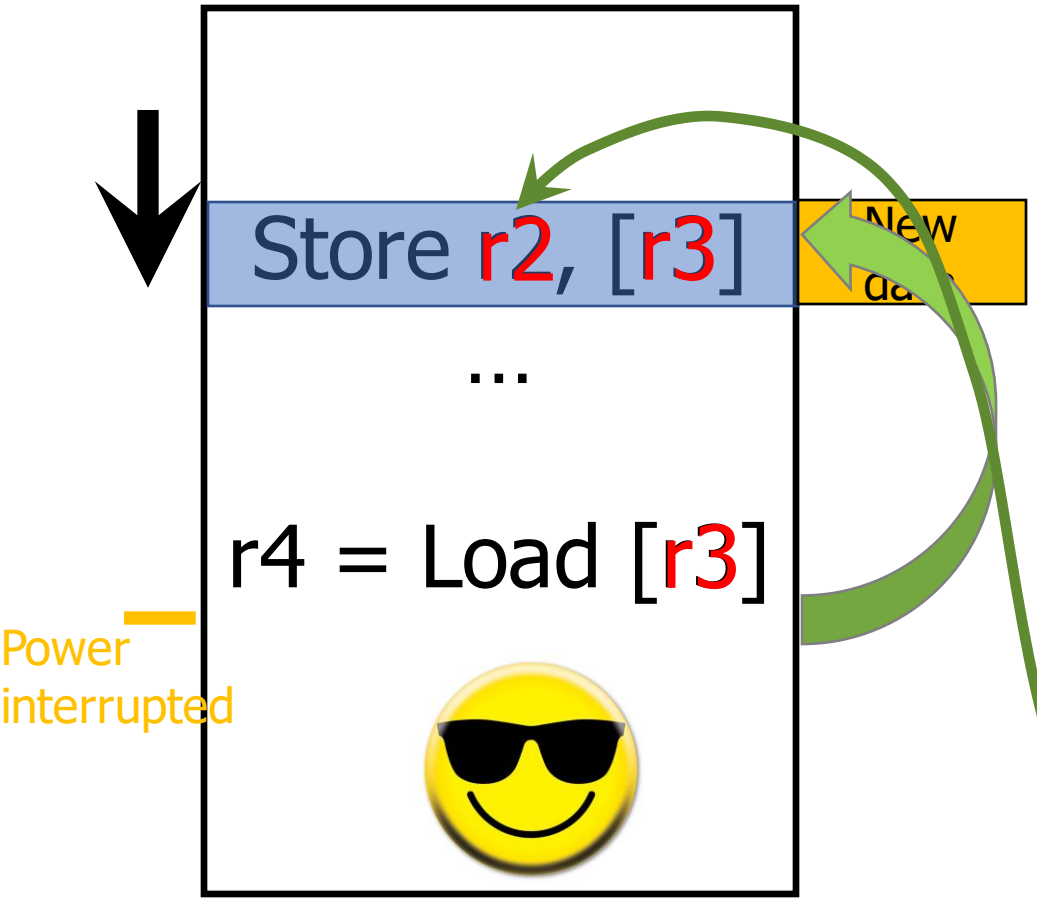
* r2, r3 are restored from NVFF

➤ Crash inconsistency caused by stores left behind power failure



ReplayCache Solution: Replaying Unpersisted Stores

* Recovery status after power failure happens



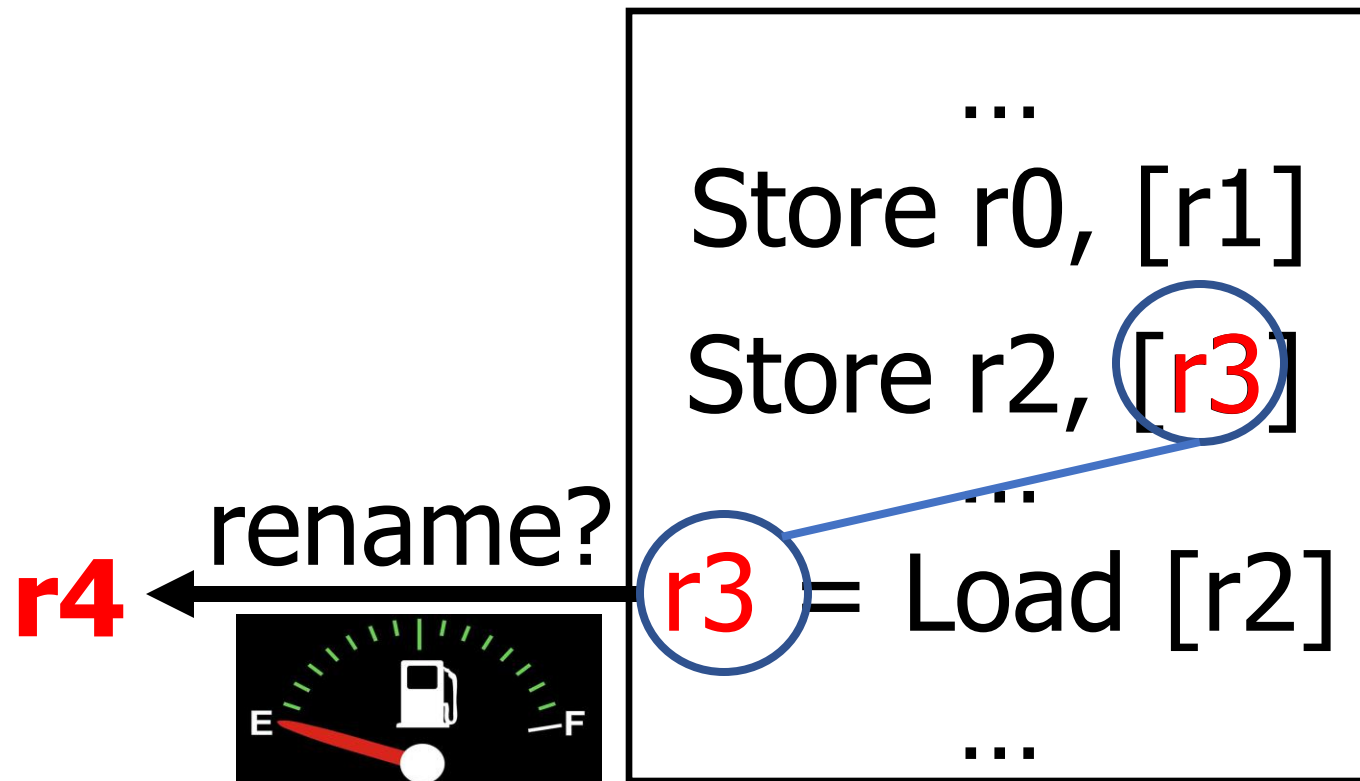
Store Integrity: A Property to Ensure Correct Replaying

- **Store integrity**: register operands of stores must be not overwritten by following definitions



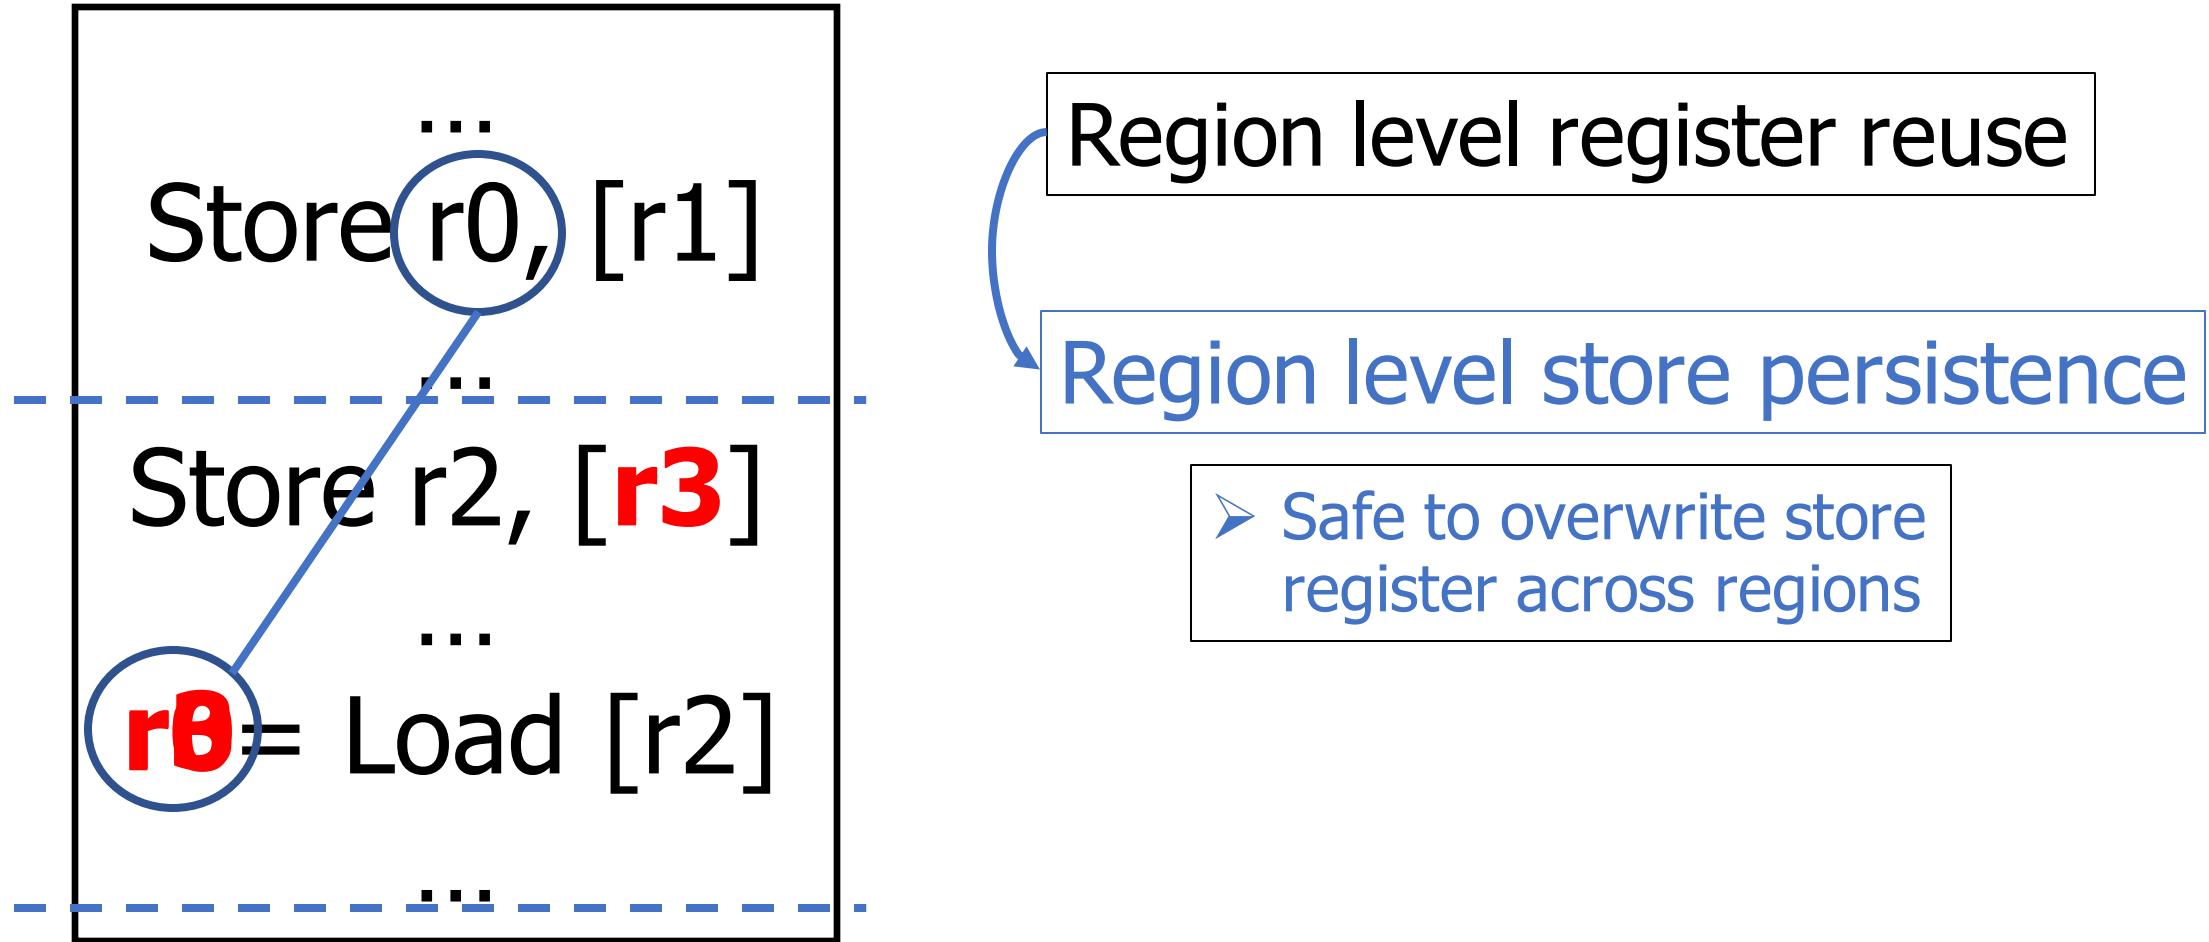
Challenge of Guaranteeing Store Integrity

- **Limited** number of registers prevents store integrity




Assume only 4 registers

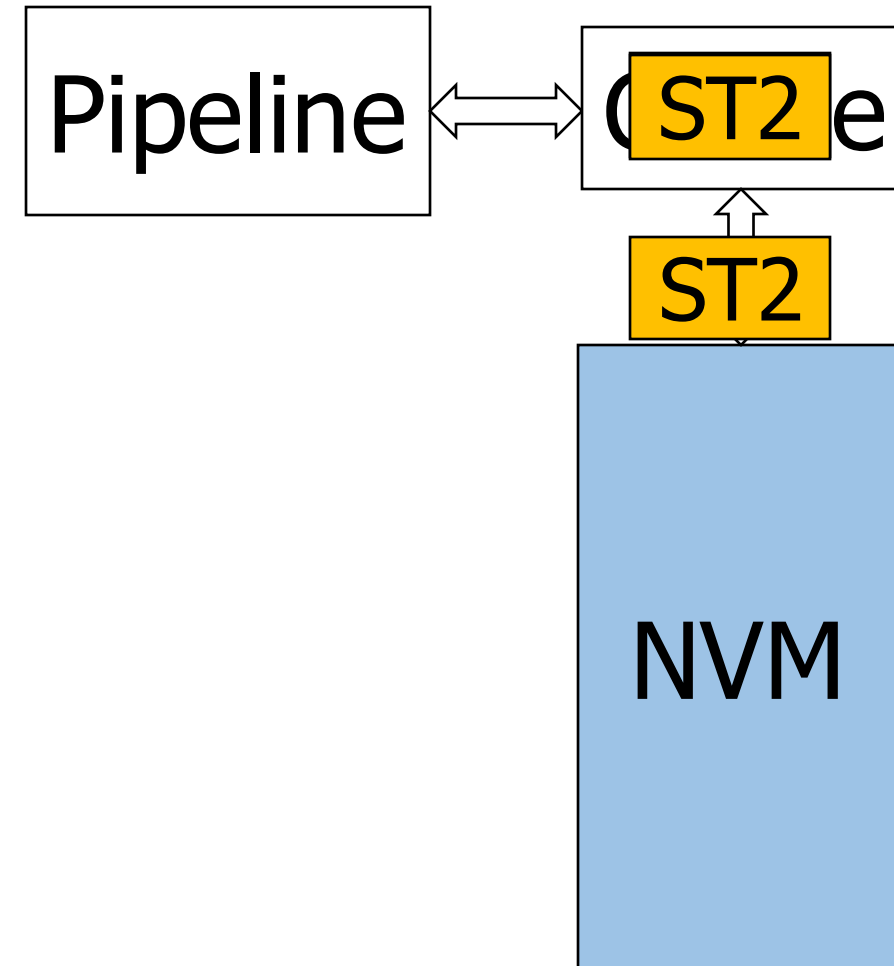
Write-after-read
(WAR) dep



Region-Level Store Persistence



* CLWB: asynchronously write back store to memory.

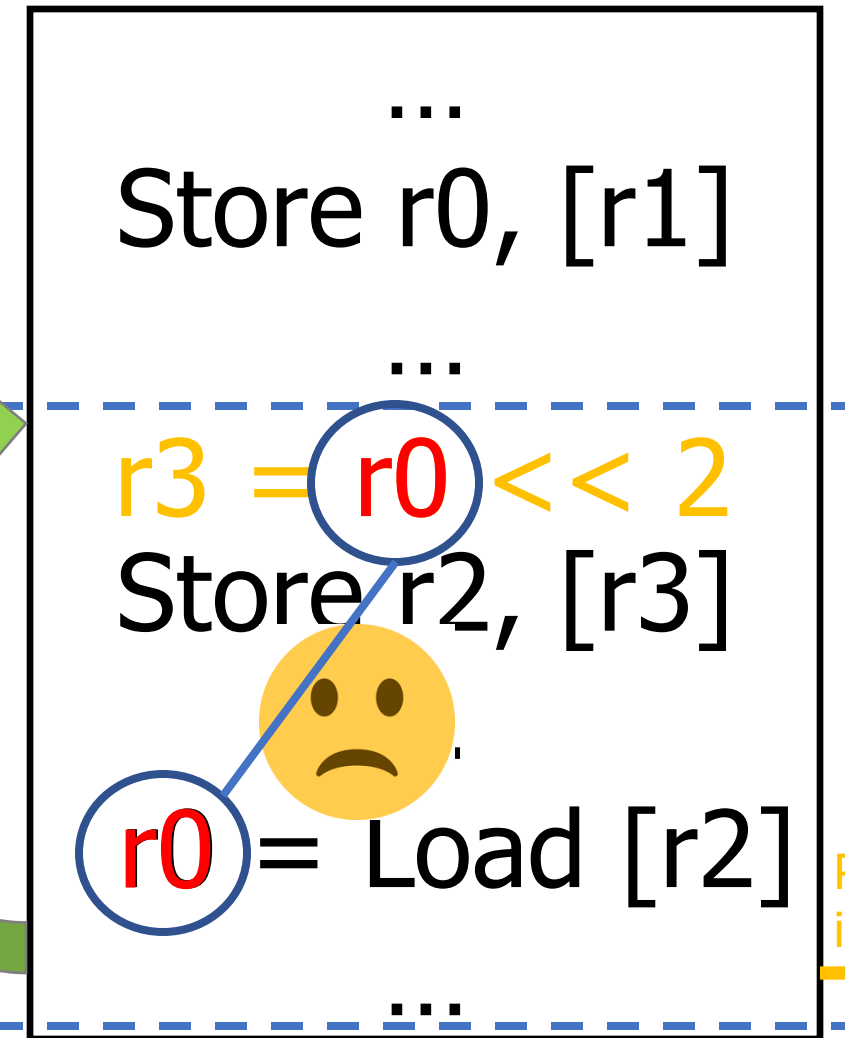



Recovery Protocol

- Stores left behind failure are the root of cause
- Replaying them in the wake of failure with store integrity
- Due to limited registers, region-level store integrity is introduced, which in turn requires region-level store persistence to allow each region to use all registers.
- **Recovery protocol**

Just Restart a Power-interrupted Region?

- Re-executing the region from its beginning **fails** to recover program status




Power interrupted

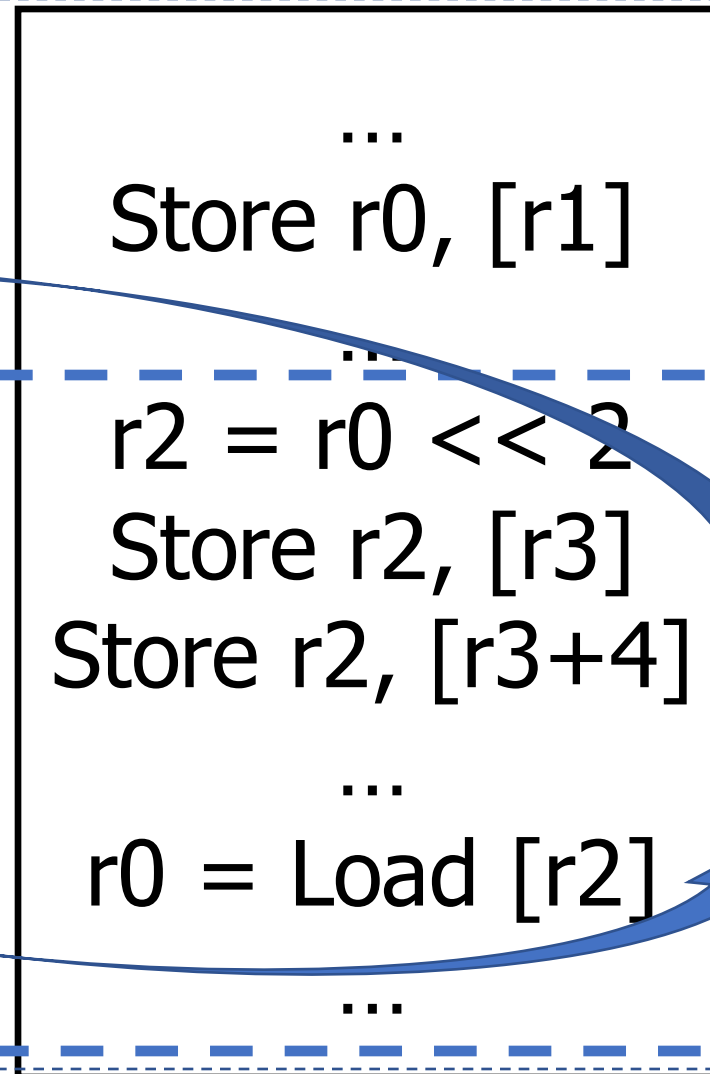
Region-Level Recovery Protocol

* Program Status and NVM status during recovery

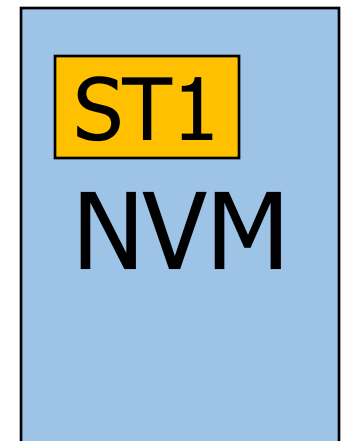
Recovery Code



access NVM registers
 r_2 and r_3 in recovery code

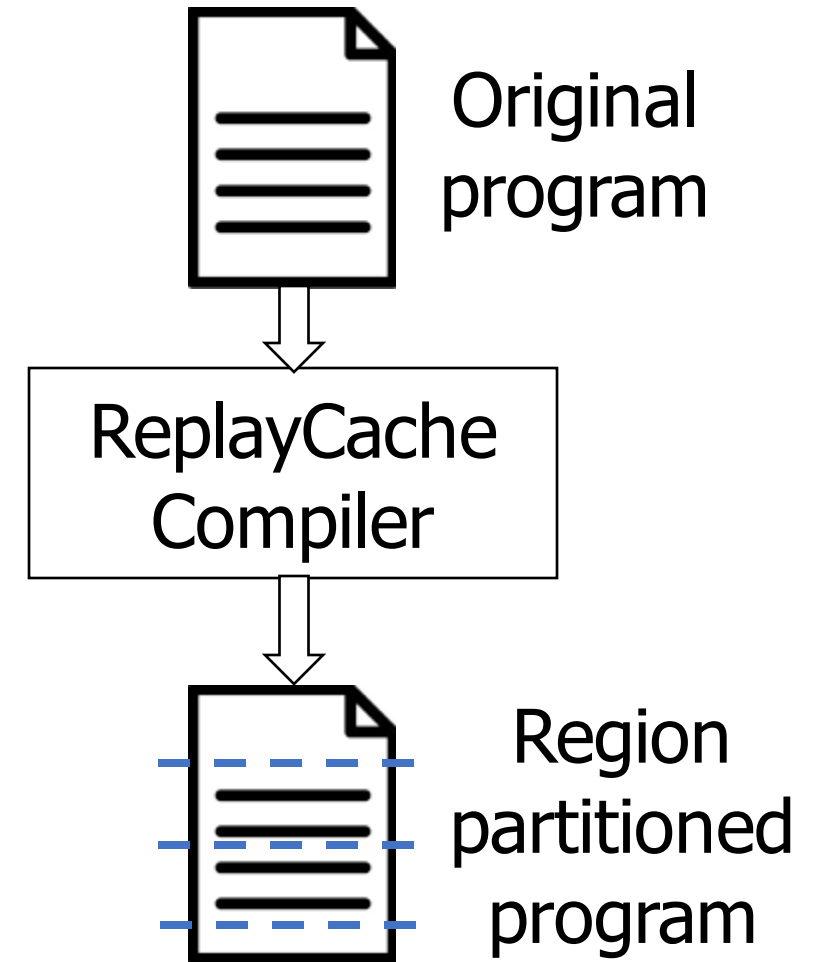


ST1

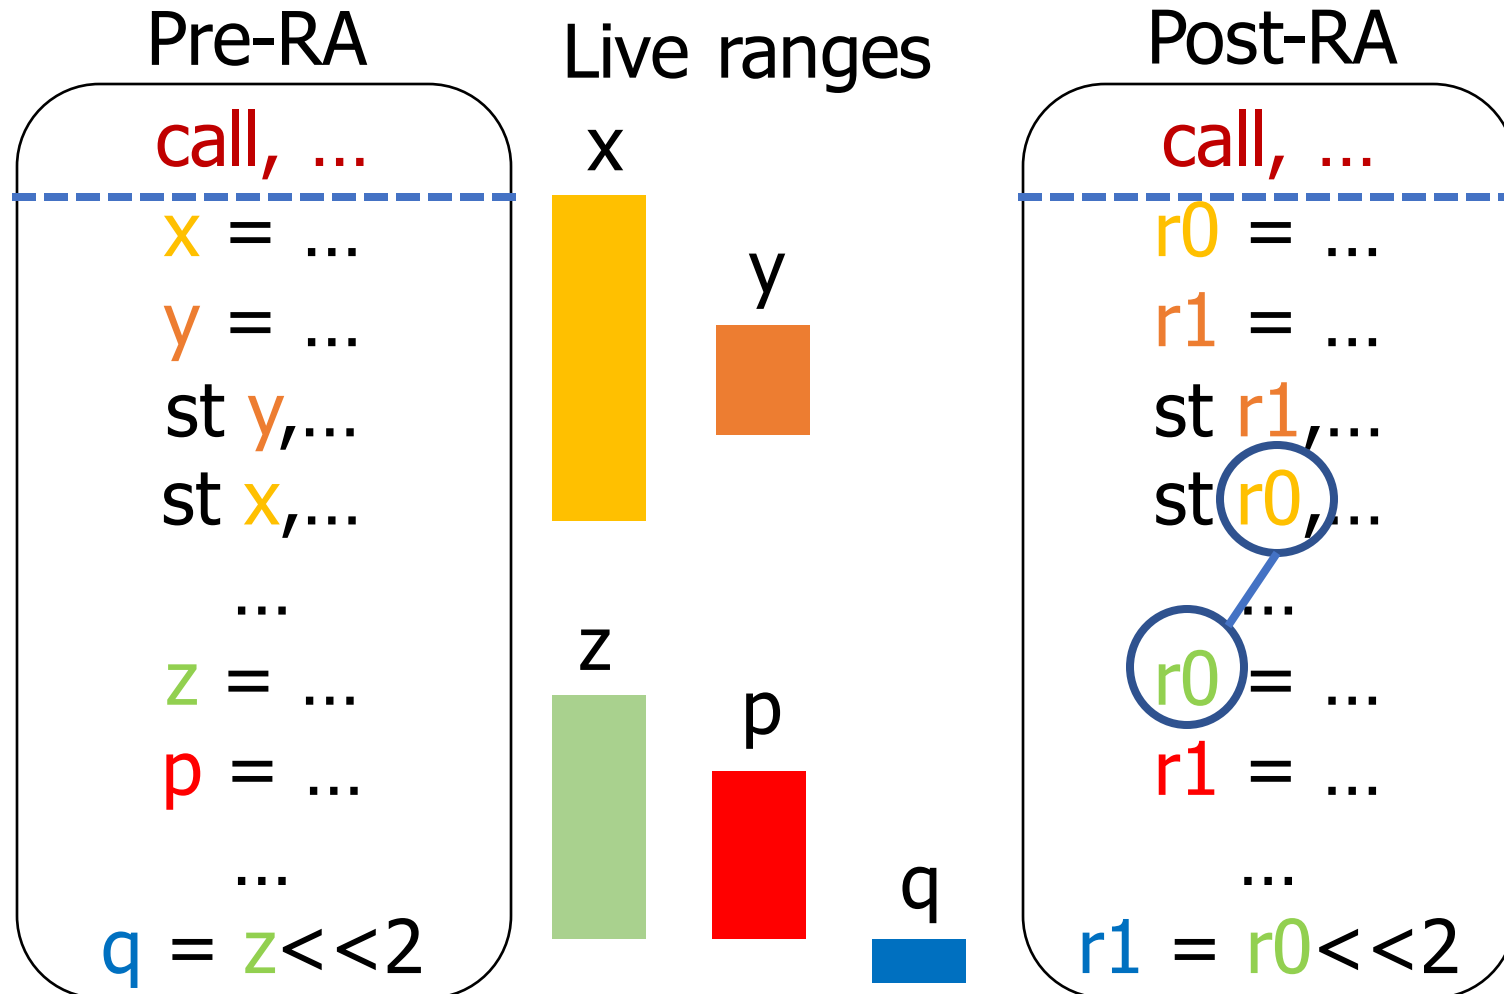


power interrupted

Challenge to Generate Store-Integrity Regions



No Store Integrity with Existing Register Allocation



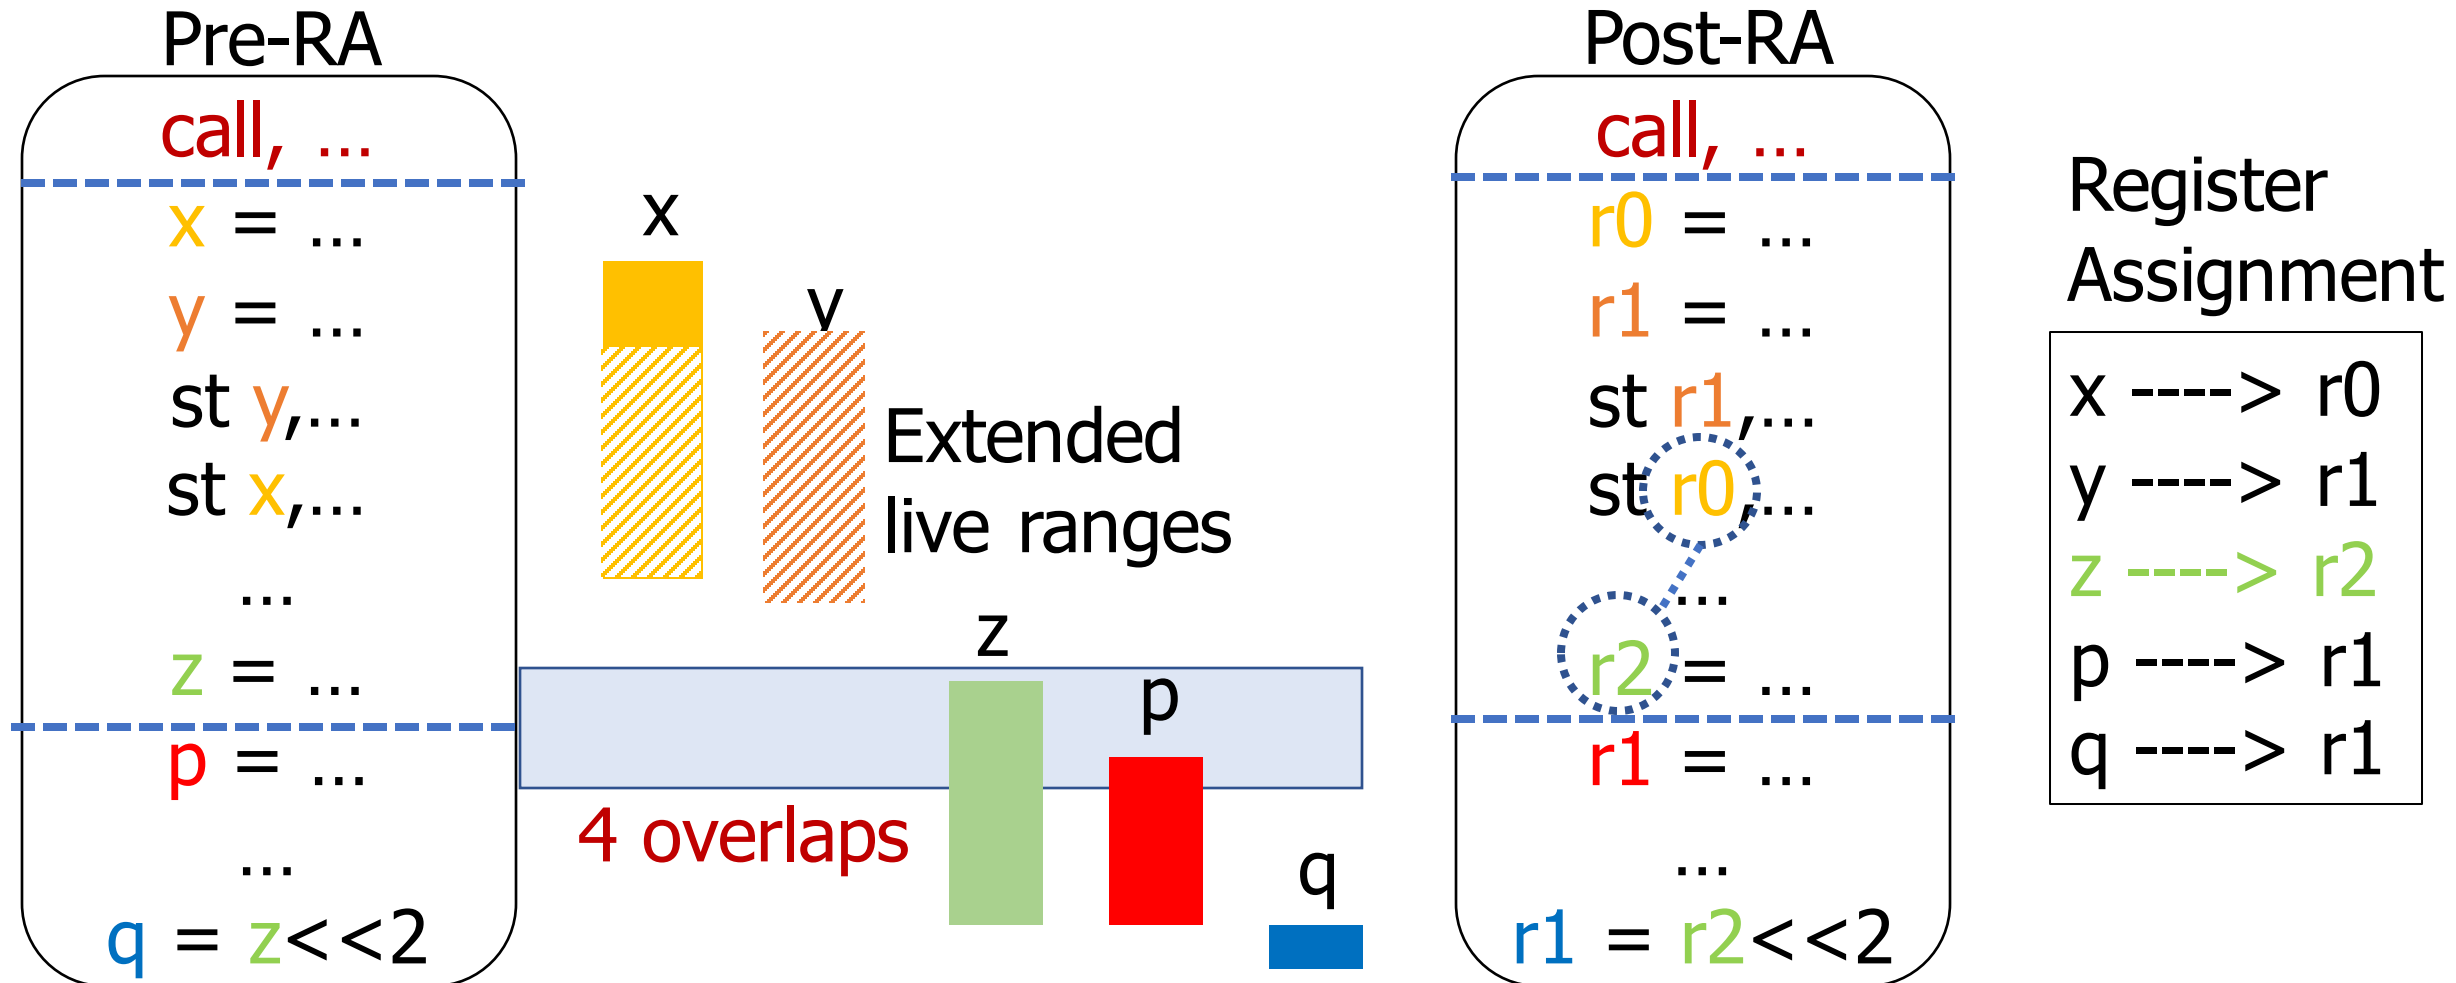
- Assume only **3** registers
- **Disjoint** live ranges can **share** the same physical register (x,z)

Register Assignment


x	---->	r0
y	---->	r1
z	---->	r0
p	---->	r1
q	---->	r1

Register-Renaming-Aware Region Partitioning

- Store operands must **not share same** registers with following definitions
- Assume 3 registers

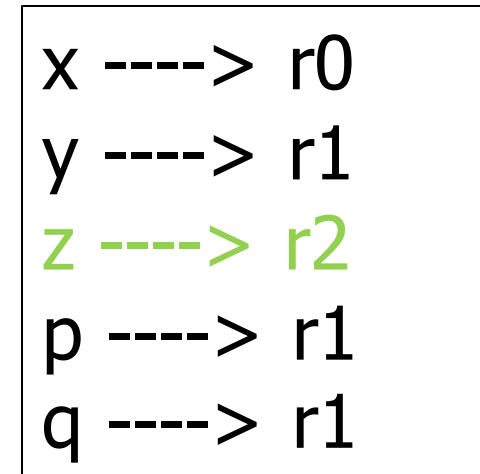


Spilling-Store-Registers Preservation




stack spill

Register Assignment

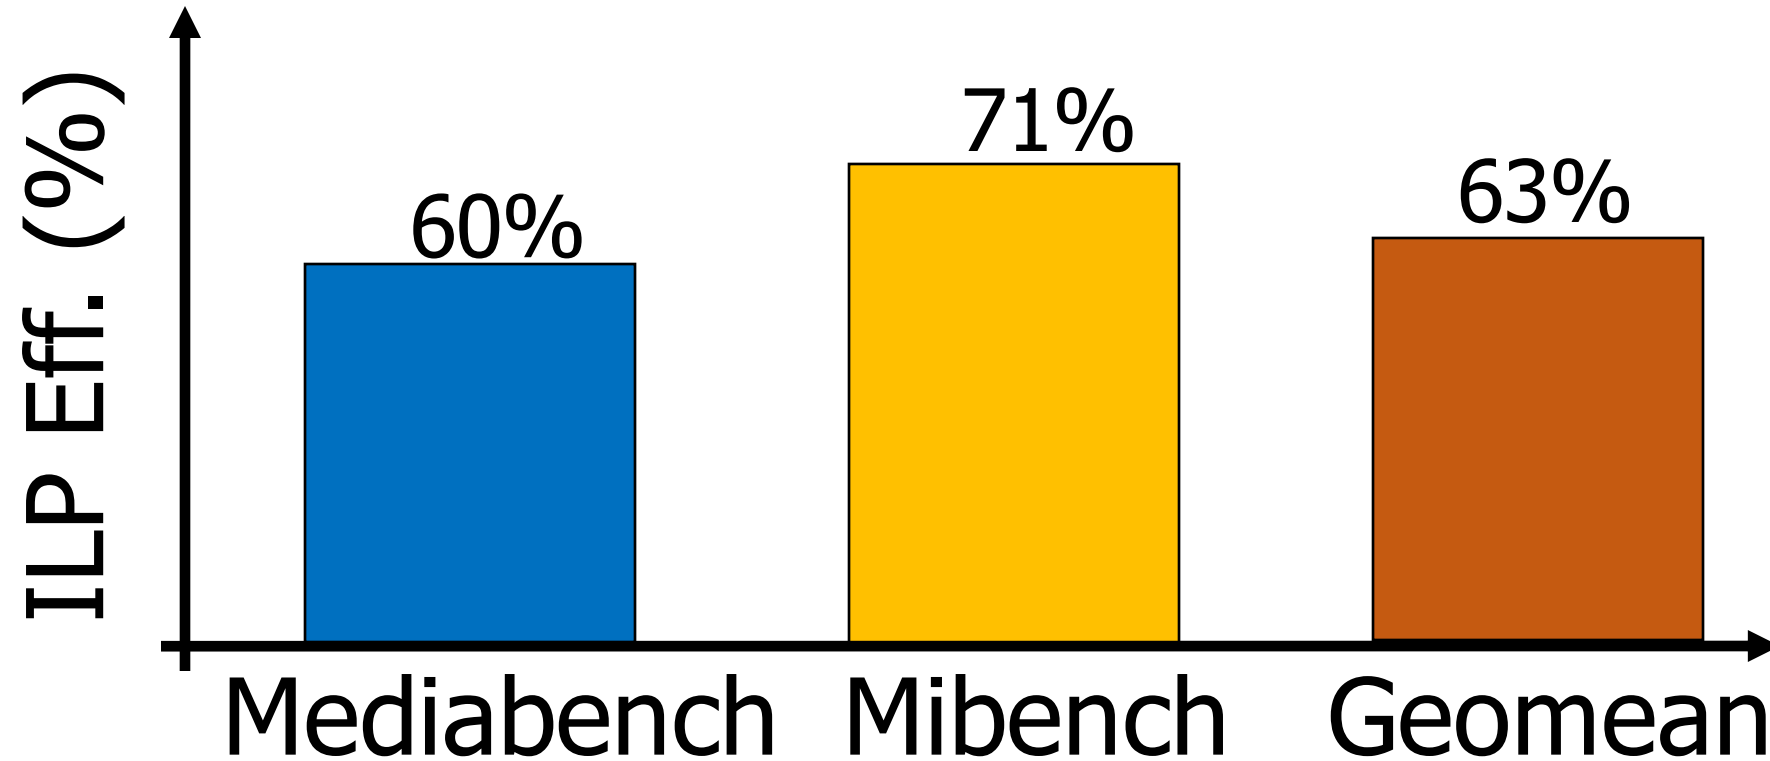


- Gem5-based **NVPSim** modeling a single core in-order ARMv7 processor with **8kB 2-way set-associative L1 I/D cache**, and 16MB ReRAM as main memory.
- **LLVM**-based region formation.
- Mediabench + Mibench.


Speedup over No-cache Baseline with Real Power Trace



ILP Efficiency (no power failure)



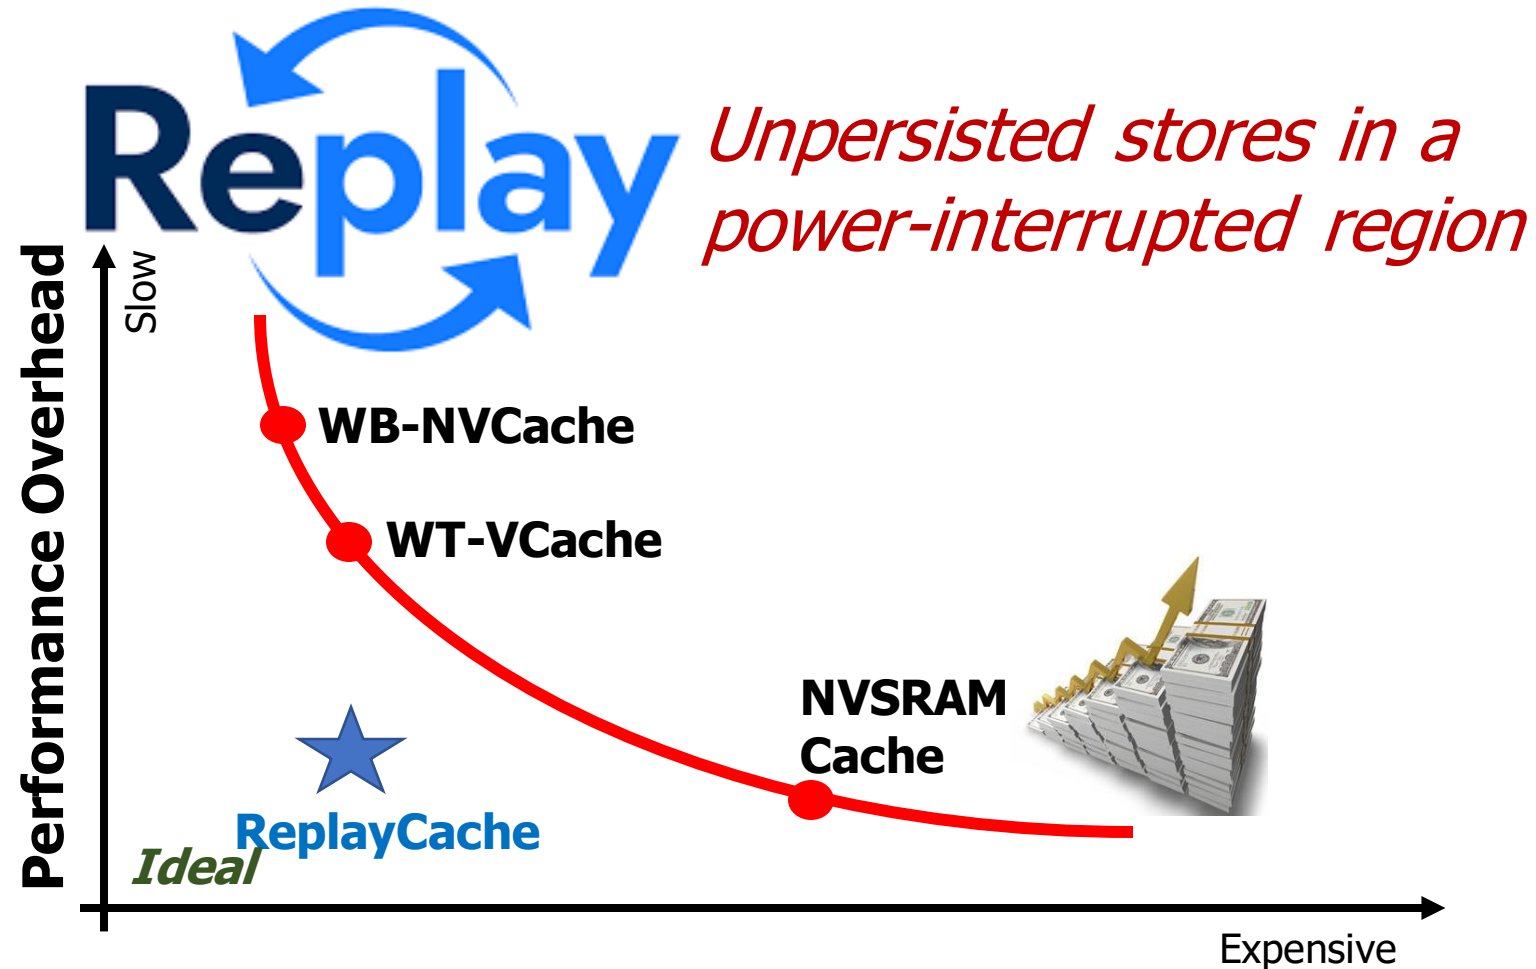
Sensitivity to Cache Size



* Office power trace

Conclusion

- A pure **software** design for enabling **WB volatile cache** with crash consistency guarantee.
- **Never amplify writes.**
- **Comparable performance** to an ideal NVSRAM cache for realistic power traces.



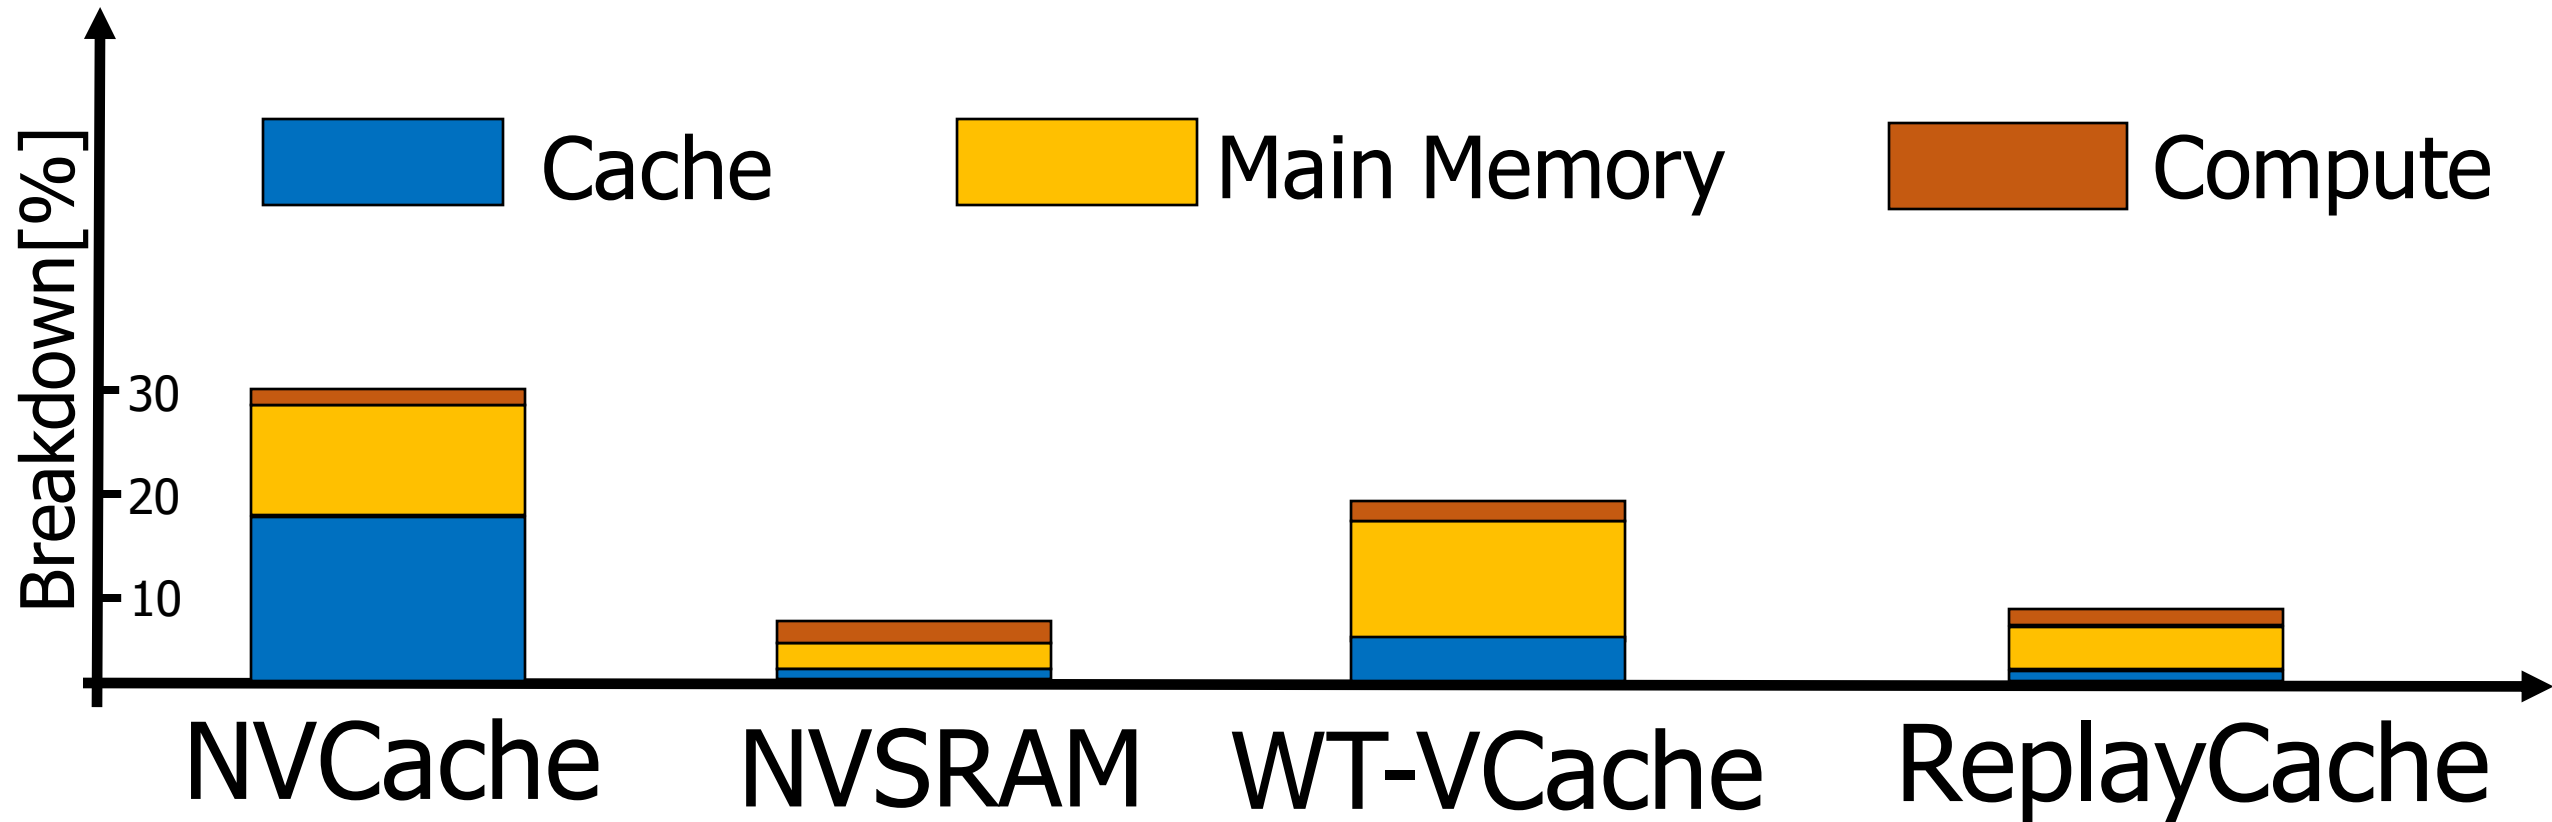
Thank You

Q&A

Jianping Zeng @Purdue University
(zeng207@purdue.edu)

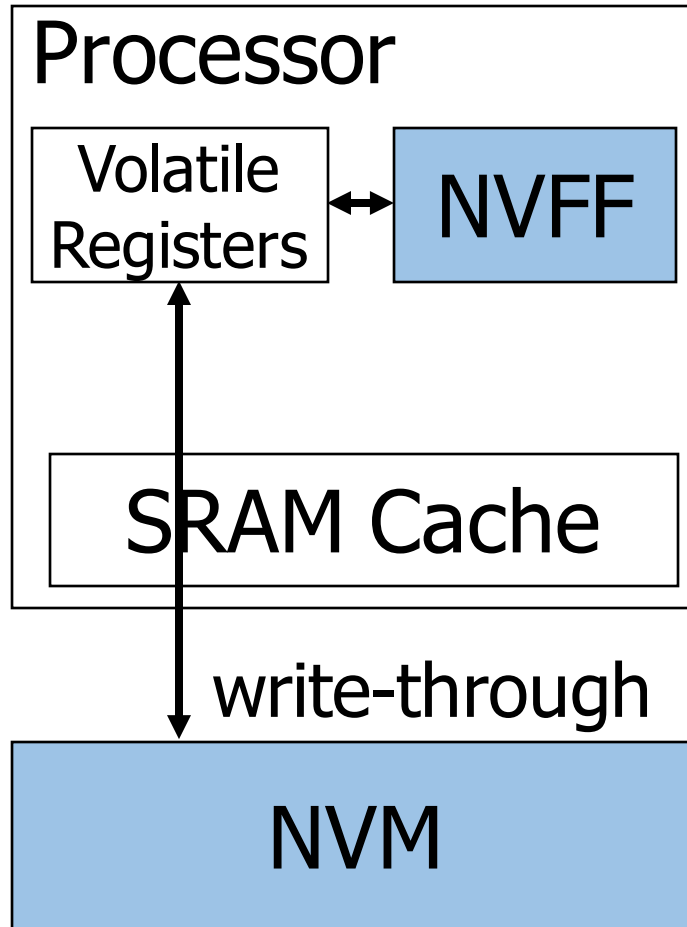
This presentation and recording belong to the authors. No distribution is allowed without the authors' permission.

Normalized Energy Consumption Breakdown

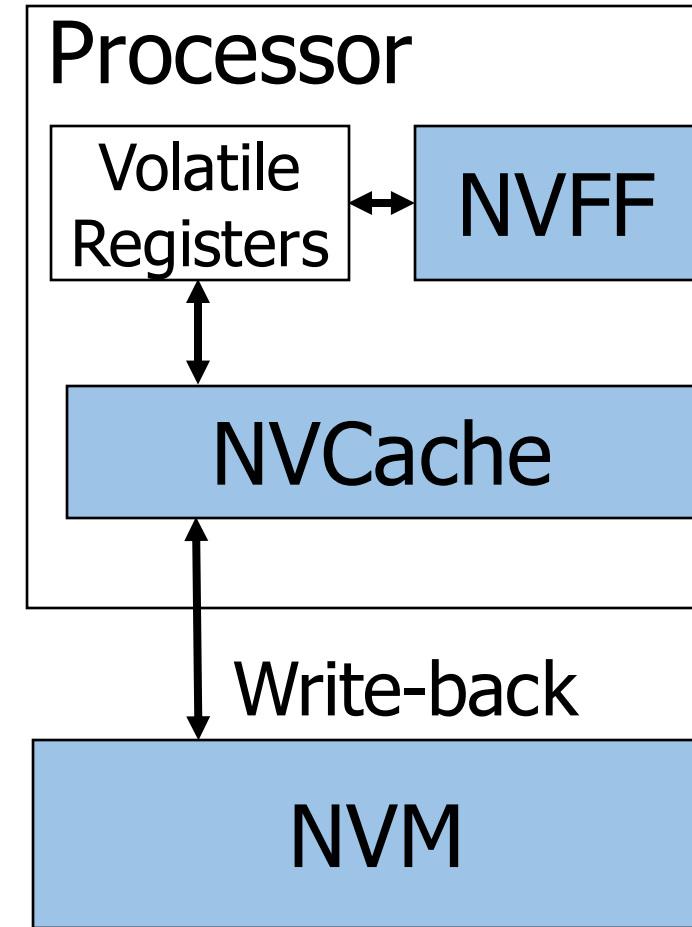


* Office power trace

Architecture of WT-VCache and WB NVCache



(a). Write-through (WT) volatile cache



(b). Write-back (WB) Nonvolatile cache

Speedup over Non-cache Baseline

