Design Automation Conference 2025

DAC 2025

Author Response

Title: Intermittent Systems at Small Scale: Execution Model and Design Guidelines
Authors: Youngbin Kim and Yoojin Lim

Instructions

The author response period has begun. The reviews for your submission are displayed on this page. If you want to respond to the points raised in the reviews, you may do so in the boxes provided below.

Please note: *you are not obligated to respond to the reviews*.

Review #1

Summarize shortly the contributions of the paper in your own words.

The authors present the design guidelines to implement efficient intermittent systems with small energy storage. The proposed guidelines consider the buffering effects coming from the system's decoupling capacitors, which has been neglected in prior research. The design improves the end-to-end execution latency significantly under both static and dynamic checking pointing schemes, without incurring additional overhead.

Strengths

- Good illustrative examples to understand the issue related to the decoupling capacitors.
- The generality of the reference architecture is supported by experiment with multiple architectural setups.
- The weakness of the traditional execution model is well analyzed with supporting evidence.

Weaknesses

- The contributions involved in the proposed guidelines are unclear.
- The explanation of the execution model is too informal and high-level.

Main Discussion of Paper

[Summary] The intermittent batteryless systems have garnered

significant attention. However, the extremely small size of energy storage makes it difficult for the traditional execution model to provide a reasonable abstraction of actual execution behavior. The authors present the design guidelines to implement efficient intermittent systems with small energy storage. In particular, the proposed guidelines consider the buffering effects of the system's decoupling capacitors, which have been neglected in prior research. The design improves the end-to-end execution latency significantly under both static and dynamic checking pointing schemes, without incurring additional overhead.

[Critics] I have one major criticism. The proposed design guidelines are too high-level, making the overall contributions weak. I greatly appreciate the authors' providing ample supporting evidence to extract the general observations in designing the execution framework in intermittent computing. However, the guidelines given in Section 3 seem too obvious and remain at a too superficial level, which makes it difficult for readers to judge how to follow them. That is, most explanations in Section 3 list potential factors without detailed mechanisms to improve the efficiency of the intermittent systems. For example, the authors suggested that one idea is to delay checkpoint executions until the last possible moment (Section 3-A); it is unclear how one can compute the last possible moment and what factors interplay with each other. Regarding using V_dd as a checkpoint signal (Section 3-B), why V_ref/Vdd*2^n is the reasonable threshold to determine whether to execute a checkpoint? Regarding selecting hardware components (Section 3-C), what specific criteria should be used, and what tradeoff exists in selecting the hardware components? In short, the listed observations sound valid, but the proposed guidelines do not seem to deliver an easily applied solution since they abstract away many details. The authors should clarify this point.

Have a question for the Authors? (Optional)

Have the authors tried the design guidelines in hardware architectures that are not introduced in this paper?

Review #2

Summarize shortly the contributions of the paper in your own words.

The paper addresses the limitations of traditional execution models for intermittent systems operating on small energy storage, proposing a new model that accounts for the buffering effects of decoupling capacitors. It highlights inefficiencies in traditional models, which can lead to up to 5.62× higher power consumption and unsafe checkpoint operations. By incorporating this new model, the paper introduces design guidelines that significantly enhance the performance of static and dynamic checkpointing techniques.

Strengths

- The paper introduces a novel execution model that accounts for the buffering effects of decoupling capacitors, providing a more accurate representation of intermittent system behavior.

+The proposed guidelines enhance the performance of static and dynamic checkpointing techniques, achieving improvements of 3.04× and 2.85× on average, addressing key inefficiencies in traditional models. +The work targets a critical emerging area—batteryless, energy-harvesting IoT systems—making it highly relevant for sustainable and low-maintenance technologies. +The paper evaluates its model and guidelines across multiple benchmarks and system configurations, demonstrating robust performance gains and validating the approach. +It offers actionable design guidelines, such as delaying checkpoint execution and utilizing precise voltage monitoring, which can be readily applied to improve existing intermittent systems...

Weaknesses

- The proposed execution model and design guidelines are primarily evaluated through benchmarks and controlled experiments, with little emphasis on real-world deployment scenarios or diverse environmental conditions.
- The evaluation uses specific hardware configurations (e.g., STM32L5, MSP430FR5994), limiting the generalizability of the findings to other architectures or emerging hardware platforms.
- The paper discusses issues with sub-normal voltage operation but provides limited solutions for peripherals and components that may behave unpredictably under such conditions.
- The proposed guidelines, particularly those involving precise checkpoint timing and new voltage monitoring techniques, may increase design complexity and require additional hardware or software modifications, which might not be feasible for all systems.

Main Discussion of Paper

This is a well written paper and despite the weaknesses listed above it is an important paper which will be able to enhance the use of small systems in constrained scenarios. The examples used are a bit simple, they could have used more robust SPEC benchmarks. It is also not clear how a system such as this will work with differing power units and processors. The system is only tested with with limited hardware. Other things like temperature should be varied to check whetehr this would work under all conditions or whether there are limitations. The additional design complexity should be discussed.

Review #3

Summarize shortly the contributions of the paper in your own words.

This paper proposes a novel approach to enhance the efficiency of intermittent computing systems. It identifies limitations in traditional execution models that fail to account for energy buffering effects from decoupling capacitors, leading to inefficiencies and unsafe

checkpoints. The new execution model proposed by the authors addresses these issues, getting up to 5.62x improvement in power efficiency. This paper also provides design guidelines that optimize static and dynamic checkpoint techniques, achieving average performance gains of 3.04x and 2.85x.

Strengths

- Novel Findings and Proposed Model: Identifies buffering effects from decoupling capacitors and proposes delaying checkpoints and using Vdd with a reference voltage for signal detection.
- Clear Explanation of Findings: Clearly details hardware configurations and experimental setups to derive design considerations.
- Solid and comparative evaluation: The authors compare the proposed approach against two traditional approaches, static and dynamic schemes, with kernel and application benchmarks.

Weaknesses

- Paper Organization: The related work and literature review is scattered over multiple sections.
- Limited applicability: the proposed work has limited applicability to specific systems, and there is limited discussion on generalizing the approach to more diverse systems.
- Although minor, there are typos and inconsistencies throughout the paper.

Main Discussion of Paper

Although there are issues in paper organization, applicability, and readability, this paper proposes a novel execution model for batteryless intermittent systems, which was underexplored before, with a clear explanation of findings and a solid evaluation. Also, it should be straightforward to address the identified issues before the camera-ready submission. Therefore, this paper would be a valuable and interesting paper to be added to DAC this year.

Detailed comments on strengths:

- Novel Findings and Proposed Model: This paper has novel observations, such as the buffering effect caused by decoupling capacitors, which has been overlooked before. Based on this finding, this paper proposes a design practice to delay checkpoint executions as late as possible. Also, they propose to use the Vdd with a reference voltage for checkpoint signals.
- Clear Explanation of Findings: The paper clearly explains how they got the key observations. They provided all hardware configurations and experimental setups to get the results they observed. Also, they transparently explained the steps of how they derived the impacts of power efficiency, predicting power failures, and sub-normal voltage execution.
- Solid Evaluation: This paper evaluates the proposed solutions

based on rigorous tests. They test static and dynamic checkpointing schemes using multiple benchmarks, even considering other hardware configurations. This robust validation makes the results reliable. Also, using benchmarks shows its general applicability to various architectures and scenarios.

Detailed comments on weaknesses:

- Paper Organization -While the paper flows well overall, the paper organization can still be improved to avoid confusion, e.g., Related Work and literature review are scattered over multiple sections; however, they can be consolidated in one section to help readers understand the traditional approaches the author evaluated their approach against.
- Limited Impact: Although the proposed approach works well in the domain of this paper, it would have been great if the authors generalized the discussion of their work in similarly energy-constrained systems or at a larger scale.
- There are minor writing mistakes, typos, and inconsistencies.

Minor comments:

- A typo in the abstract: shows -> show
- The authors should add units in Fig. 9
- The authors should add a period after Table 1.
- Summarizing design guidelines using a table or list would be helpful if they are available.
- Fig 6's ratio scale is 0-1, and Fig 8. uses percentages. It would be better if it is consistent.
- Consistency in writing on units. Only "ms" has a space after the numbers.

Have a question for the Authors? (Optional)

What is the overhead of monitoring Vdd, and will this affect the performance?

Submit Response to Reviewers

Use the following boxes to enter your response to the reviews. Please limit the total amount of words in your comments to 600 words (longer responses will not be accepted by the system).

Response to Review #1:

Response to Review #2:

Response to Review #3:

General Response to Reviewers: