



上海交通大学
SHANGHAI JIAO TONG UNIVERSITY



SHEEO: Continuous Energy Efficiency Optimization in Autonomous Embedded Systems



Presenter: Xinkai Wang (Shanghai Jiao Tong University)

Advisor: Chao Li, Minyi Guo

ICCD 2024, Milan, Italy







Content




1. Background and Motivation




2. Design of SHEEO



3. Evaluation Results



4. Conclusion and Vision







Content




1. Background and Motivation




2. Design of SHEEO



3. Evaluation Results



4. Conclusion and Vision





Autonomous Embedded System (AES)



AES is promising to minimize human intervention in various tasks.



Autonomous
Driving



Intelligent
Manufacturing



Unmanned
Drone



Space
Exploration

Support

*Changing
Environments*

Sensing



Perception



Energy

Actuation

*Adaptive
Actions*



Autonomous Embedded System (AES)



AES is promising to minimize human intervention in various tasks.



Autonomous
Driving



Intelligent
Manufacturing



Unmanned
Drone



Space
Exploration

Support

*Changing
Environments*

Sensing



Perception



Energy

Actuation

*Adaptive
Actions*




Background: Present AES Pipeline



AES has three-stage pipeline on limited heterogeneous hardware.


Three-Stage
Software
Pipeline



Middleware
Facilities



Limited
Hardware
Resources






Background: Present AES Pipeline



AES has three-stage pipeline on limited heterogeneous hardware.


Three-Stage
Software
Pipeline



Middleware
Facilities



Limited
Hardware
Resources






Background: Present AES Pipeline



AES has three-stage pipeline on limited heterogeneous hardware.


Three-Stage
Software
Pipeline



Middleware
Facilities



Limited
Hardware
Resources






Background: Present AES Pipeline



AES has three-stage pipeline on limited heterogeneous hardware.


Three-Stage
Software
Pipeline



Middleware
Facilities



Limited
Hardware
Resources







Background: Present AES Pipeline



AES has three-stage pipeline on limited heterogeneous hardware.





Background: AES Efficiency Optimization (EO)


AES faces varied environments and complex hardware to manage.

Temporal Variance: External Dynamics.

AES 	Urban Road	Traffic Jam	Overtake	Highway
Dynamics	Case 1	Case 2	Case 3	Case 4
Vehicle Speed	Low	Low	Fluctuated	High
Surrounding	Low	Fluctuated	Low	Low
Task Complexity	Low	High	High	Low
Current Method	V/F ↓		V/F ↑	
Expected Method	V/F ↓	Adapting V/F with time		V/F ↓

Require **Continuous** EO

Spatial Variance: Accelerator Behavior.




Require **Intelligent** EO

AES expects EO in **100ms** granularity within a large optimization space.


Researchers explore resource and energy optimization separately.

Increasing Resource Heterogeneity




SOV, MICRO'20 [1]

Accelerators offer unbalanced computing power




AutoPilot, MICRO'22 [2]

Reducing Energy Consumption



FIRM, OSDI'20 [3]



CALOREE, ASPLOS'18 [4]

On-device ML-based EO is costly for AES.

Lack of utilizing **ignored** hetero-computing resource for **costly** EO facility .

[1] Yu, Bo, et al. "Building the computing system for autonomous micromobility vehicles: Design constraints and architectural optimizations." *MICRO 2020*


[2] Krishnan, Srivatsan, et al. "Automatic domain-specific soc design for autonomous unmanned aerial vehicles." *MICRO 2022*

[3] Qiu, Haoran, et al. "FIRM: An intelligent fine-grained resource management framework for SLO-Oriented microservices." *OSDI 2020*

[4] Mishra, Nikita, et al. "Caloree: Learning control for predictable latency and low energy." *ASPLOS 2018*


Desirable efficiency optimization interfere with AES pipeline.

Complex ML models introduces great **training / inference overheads.**



FIRM, OSDI'20



Reinforcement Learning
Overhead=**420ms / 40ms**



CALOREE, ASPLOS'18

Transfer Learning
Overhead=**500ms / 2ms**

Continuous EO **disturbs execution** on resource-constrained AES.



Typical EO using normal computing power causes **~10%** slowdown of each iteration.




Complex ML models introduces great **training / inference overheads.**

[illegible]



Transfer Learning

Continuous EO **disturbs execution**
on resource-constrained AES.




AES presents underutilized shadow cycles for deploying EO.

Diversity of inference and accelerator present **misaligned** execution.



Typical inferences on GPU and DLA produces **~12%** underutilized resources.

Shadow cycles are **underutilized and ignored** computing resources in AES.



- **Vertical Shadow Cycles (VSC)**
Preemptive and Variable-sized
- **Horizontal Shadow Cycles (HSC)**
Periodic and Short-lived


The ignored resources in AES are promising for EO without slowdown.




Content




1. Background and Motivation




2. Design of SHEEO



3. Evaluation Results



4. Conclusion and Vision





Two key issues exist in SHEEO design...

Using Heterogeneous Resource

Reducing Energy Consumption

How to **Observe** AES?



Internal Runtime Status




External Environment





How to **Optimize** AES?

Costly Training

Agile Inference




Utilization of Shadow Cycles




Learning and Control

SHEEO exploits shadow cycles for continuous and intelligent EO.



- SHEEO has **two cooperative components**: real-time monitor and continuous learning manager.
- SHEEO works at OS, enhancing **power management facility** without intrusion into AES pipeline.

SHEEO observes internal and external status via real-time monitor.



(a) Multi-dimensional state-related features

State		Descriptions	Discrete Values
Workload Features	S _{CONV}	Number of CONV layers	L (<30), M (<50), H (≥50)
	S _{FC}	Number of FC layers	L (<10), H (≥10)
	S _{RC}	Number of RC layers	L (<10), H (≥10)
Runtime Dynamics	S _{Com}	Computing units' utilization	L (<25%), M (<75%), H (≥75%)
	S _{Mem}	Memory utilization	L (<25%), M (<75%), H (≥75%)
	S _{Speed}	Current vehicle speed	L (<10mph), M (<40mph), H (≥40mph)
	S _{Var}	Variation from last execution	L (<10%), M (<30%), H (≥30%)

➤ State Table for Internal Runtime Status


- Model software and hardware features.
- Adjustable for actual AES workloads.

➤ State Table for External Environment

- Model the external stochastic variance.
- Use discrete ranges to reduce complexity.

- Runtime Monitor observes the **start and end of AES tasks** for recording shadow cycles.
- It observes the environment changes and system status and records them into the **7-dimensional state table**.

SHEEO observes internal and external status via real-time monitor.



- Runtime Monitor observes the **start and end of AES tasks** for recording shadow cycles.
- It observes the environment changes and system status and records them into the **7-dimensional state table**.

(a) Multi-dimensional state-related features

State		Descriptions	Discrete Values
Workload Features	S _{CONV}	Number of CONV layers	L (<30), M (<50), H (≥50)
	S _{FC}	Number of FC layers	L (<10), H (≥10)
	S _{RC}	Number of RC layers	L (<10), H (≥10)
Runtime Dynamics	S _{Com}	Computing units' utilization	L (<25%), M (<75%), H (≥75%)
	S _{Mem}	Memory utilization	L (<25%), M (<75%), H (≥75%)
	S _{Speed}	Current vehicle speed	L (<10mph), M (<40mph), H (≥40mph)
	S _{Var}	Variation from last execution	L (<10%), M (<30%), H (≥30%)


➤ State Table for Internal Runtime Status

- Model software and hardware features.
- Adjustable for actual AES workloads.

➤ State Table for External Environment

- Model the external stochastic variance.
- Use discrete ranges to reduce complexity.

SHEEO optimizes efficiency via continuous learning manager.



Algorithm 1: Algorithm for continuous learning

Input: Pre-trained Q-table $Q(S, A)$, learning rate α , discount factor γ , exploration probability ϵ
Output: Fine-tuned Q-table and Action A for each iteration

```

1 while  $\exists$  VSC periods and  $\exists$  Store Table do
2   Calculate recorded reward  $R_{energy}$  and  $R_{latency}$ ;
3   For consecutive  $S$  and  $S'$  and chosen action  $A$ ;
4   Choose action  $A'$  with the largest  $Q(S', A')$ ;
5    $Q(S, A) \leftarrow Q(S, A) + \alpha[R + \gamma Q(S', A') - Q(S, A)]$ ;
6    $S \leftarrow S'$ 
7 end
8 foreach HSC periods do
9   Gather  $S$  and locate in  $Q(S, A)$ ;
10  if  $rand() < \epsilon$  then
11    Choose action  $A$  randomly;
12  else
13    Choose action  $A$  with the largest  $Q(S, A)$ ;
14  end
15  Store Table  $\leftarrow (A, S')$ ;
16 end
  
```

Utilizing preemptive and variable-sized VSC.

Two-fold Reward

- Minimize performance slack
- Maximize energy efficiency
- Used for updating Q-table

$$R = aR_{latency} + bR_{energy} \text{ s.t. } R_l \leq \text{Constraint}$$

$$R_{latency} = \max_{t \in \text{Perception}} T_{end}^t - T_{start}$$


$$R_{energy} = \sum_i E_{Unit}^i = \sum_i \int_{T_b}^{T_e} P_f \delta t + P_{idle} \times t_{idle}$$

- Continuous Learning Manager adopts **reinforcement learning (Q-learning)** to optimize energy efficiency.
- It embeds continuous learning into **VSC** and action control into **HSC** to reduce power management facility costs.

➤ SHEEO **learns** optimal action of given states.

- Consider status of consecutive iterations.
- Adjustable for Q-learning, DDPG, DRL, etc.

SHEEO optimizes efficiency via continuous learning manager.



Algorithm 1: Algorithm for continuous learning

Input: Pre-trained Q-table $Q(S, A)$, learning rate α , discount factor γ , exploration probability ϵ
Output: Fine-tuned Q-table and Action A for each iteration

```

1 while  $\exists$  VSC periods and  $\exists$  Store Table do
2   Calculate recorded reward  $R_{energy}$  and  $R_{latency}$ ;
3   For consecutive  $S$  and  $S'$  and chosen action  $A$ ;
4   Choose action  $A'$  with the largest  $Q(S', A')$ ;
5    $Q(S, A) \leftarrow Q(S, A) + \alpha[R + \gamma Q(S', A') - Q(S, A)]$ ;
6    $S \leftarrow S'$ 
7 end
8 foreach HSC periods do
9   Gather  $S$  and locate in  $Q(S, A)$ ;
10  if  $rand() < \epsilon$  then
11    Choose action  $A$  randomly;
12  else
13    Choose action  $A$  with the largest  $Q(S, A)$ ;
14  end
15  Store Table  $\leftarrow (A, S')$ ;
16 end
  
```

Utilizing short-lived and periodic HSC.

7-dimensional Action

- Divided into discrete levels.
- Consider heterogeneous HW.

Actions	Descriptions
S_X	Turn On/Off Component X
N_{CPU}	Number of Active Core
F_{CPU}	V/F Level of CPU Core
F_{GPU}	V/F Level of GPU
F_{DLA1}	V/F Level of DLA1
F_{DLA2}	V/F Level of DLA2
F_{MEM}	V/F Level of Memory

- Continuous Learning Manager adopts **reinforcement learning (Q-learning)** to optimize energy efficiency.
- It embeds continuous learning into **VSC** and action control into **HSC** to reduce power management facility costs.


- **SHEEO controls** the action in next iteration.
 - Use ϵ to control exploitation & exploration.
 - Agilely make decisions with largest $Q(S, A)$.




Content




1. Background and Motivation




2. Design of SHEEO



3. Evaluation Results



4. Conclusion and Vision





Evaluation Methodology and Settings



We evaluate SHEEO extensively on commercial edge platforms.

Evaluated Platform Specifications	
Device	<i>Nvidia Jetson AGX Orin Module</i>
CPU	<i>8-core ARM Cortex-A78 v8.2</i>
GPU	<i>1792-core Ampere GPU with 56 tensor cores</i>
Memory	<i>32GB 256-bit LPDDR5</i>
Accelerator	<i>2x NVDLA v2, 1x PVA v2</i>
System	<i>Linux 5.19.104-tegra with Jetpack 5.1.1</i>
Software	<i>CUDA 11.4 and TensorRT 8.5.2</i>

Evaluated Workloads	
SSD	<i>ssd-mobilenet-v1 for object detection.</i>
YOLO	<i>yolov3-tiny-416 for image detection.</i>
SRCNN	<i>super-resolution-bsd500 for image reconstruction.</i>
Evaluated Baselines	
Workload-Aware Control [1]	<i>Use static model to profile workloads.</i>
Learning-Based Control [2]	<i>Use ML model and historical statistics.</i>
Oracle	<i>Offline optimal energy efficiency scheme.</i>


- We evaluate SHEEO on **Nvidia Jetson AGX Orin** and consider shadow cycles on **GPU and NVDLA**.
- We set learning rate as 0.9, discount factor as 0.1, and exploration probability as 0.1.
- We mimic the perception stage of AES pipeline with three networks with different combinations.

[1] Bateni, Soroush, and Cong Liu. "NeuOS: A Latency-Predictable Multi-Dimensional Optimization Framework for DNN-driven Autonomous Systems." *USENIX ATC 20*

[2] Mishra, Nikita, et al. "Caloree: Learning control for predictable latency and low energy." *ASPLOS 2018*


SHEEO harvests shadow cycles in AES under various scenarios.

An example iteration in AES pipeline with SHEEO harvesting shadow cycles.



➤ *Q-learning is too lightweight to occupy all the shadow cycles, but more advanced facilities could utilize better.*

SHEEO harvests up to 88% vertical shadow cycles and 51% horizontal shadow cycles.



➤ *The harvesting ratio is higher under higher vehicle speed and higher runtime variance scenarios.*




Reducing Energy Consumptions




SHEEO improves energy efficiency of AES under various scenarios.

SHEEO provides a more agile and accurate power management solution.



- *SHEEO achieves better energy efficiency under higher runtime variance and medium vehicle speed.*

SHEEO achieves better energy-delay product (EDP) compared with baselines.




- *SHEEO improves energy efficiency with a little sacrifice of execution performance.*




Content




1. Background and Motivation




2. Design of SHEEO



3. Evaluation Results



4. Conclusion and Vision






Vision: Future of Autonomous Systems



Towards better accelerator usage for efficient system operations.

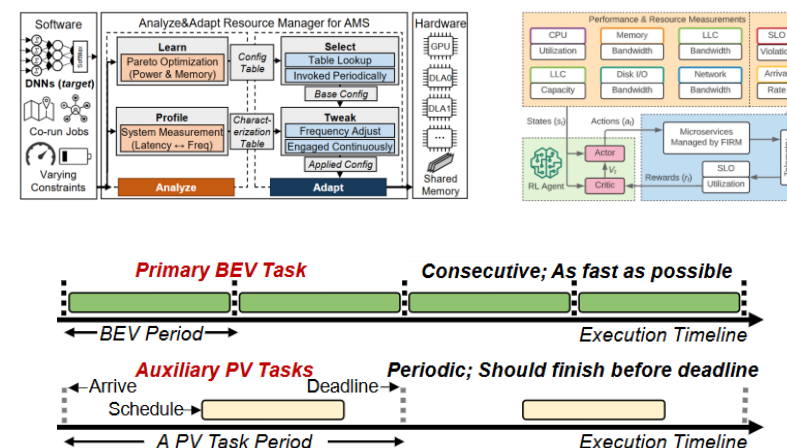
Future **accelerator-rich architecture** produce more shadow cycles to exploit!



Future **artificial brains** require more smart facilities and workloads to deploy!

ML-based Resource Mgmt. [2]

Perspective View Taks in BEV AES [3]



Moving Towards Efficient Autonomous Systems!


[1] Xinkai, Wang, et al. "Not All Resources are Visible: Exploiting Fragmented Shadow Resources in Shared-State Scheduler Architecture." SoCC 2023

[2] Lingyu, Sun, et al. "A2: Towards Accelerator Level Parallelism for Autonomous Micromobility Systems." TACO 2024


[3] Lingyu, Sun, et al. "Jigsaw: Taming BEV-centric Perception on Dual-SoC for Autonomous Driving." RTSS 2024



- We analyze **shadow cycles** in **autonomous embedded systems** and the need to **continuously optimizing** energy efficiency.



- SHEEO is an intelligent continuous energy efficiency optimizer that first **embeds the costly facility into shadow cycles**.
- SHEEO **harvests** the underutilized hetero-computing resources and **improves** AES energy efficiency under various scenarios.



- By exploiting **“free”** shadow cycles within more heterogeneous accelerators, AES can enable a wide range of intelligent **management facilities and autonomous workloads**.



Thank You! Q & A

SHEEO: Continuous Energy Efficiency Optimization in Autonomous Embedded Systems

Discussion: Xinkai Wang (Presenter), unbreakablewxk@sjtu.edu.cn



上海交通大学

SHANGHAI JIAO TONG UNIVERSITY

