



UNIVERSITY OF LEEDS

This is a repository copy of *Dynamic GPU Energy Optimization for Machine Learning Training Workloads*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/182579/>

Version: Accepted Version

Article:

Wang, F, Zhang, W, Lai, S et al. (2 more authors) (2021) Dynamic GPU Energy Optimization for Machine Learning Training Workloads. *IEEE Transactions on Parallel and Distributed Systems*. p. 1. ISSN 1045-9219

<https://doi.org/10.1109/tpds.2021.3137867>

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Dynamic GPU Energy Optimization for Machine Learning Training Workloads

Farui Wang, Weizhe Zhang, *Senior Member, IEEE*, Shichao Lai, Meng Hao, and Zheng Wang

Abstract—GPUs are widely used to accelerate the training of machine learning workloads. As modern machine learning models become increasingly larger, they require a longer time to train, leading to higher GPU energy consumption. This paper presents GPOEO, an online GPU energy optimization framework for machine learning training workloads. GPOEO dynamically determines the optimal energy configuration by employing novel techniques for online measurement, multi-objective prediction modeling, and search optimization. To characterize the target workload behavior, GPOEO utilizes GPU performance counters. To reduce the performance counter profiling overhead, it uses an analytical model to detect the training iteration change and only collects performance counter data when an iteration shift is detected. GPOEO employs multi-objective models based on gradient boosting and a local search algorithm to find a trade-off between execution time and energy consumption. We evaluate the GPOEO by applying it to 71 machine learning workloads from two AI benchmark suites running on an NVIDIA RTX3080Ti GPU. Compared with the NVIDIA default scheduling strategy, GPOEO delivers a mean energy saving of 16.2% with a modest average execution time increase of 5.1%.

Index Terms—Dynamic energy optimization, online application iteration detection, multi-objective machine learning, GPU

1 INTRODUCTION

In recent years, deep neural networks (DNNs) have demonstrated breakthrough effectiveness in various tasks that were once deemed impossible [1], [2]. Training an effective DNN requires performing expensive training on a large number of samples. As the model training time increases, the energy consumption of machine learning training workloads also grows. This is a particular concern for high-performance GPUs because they are widely used to train DNN workloads but are less energy-efficient than their CPU counterparts. As a result, there is a critical need to reduce the energy consumption for GPUs when training machine learning workloads. Achieving this goal can reduce the maintenance cost and allow users to train larger models within the same budget.

Most of the existing studies in general GPU energy optimization are offline techniques [3]–[7]. These techniques require profiling applications ahead of time [3]–[5], [7] or analyzing their source code [6] in advance to collect application features. These approaches cannot adapt to the change of the program behavior and can incur expensive profiling overhead. Other approaches change the internal working mechanisms of machine learning (ML) models [8], [9], but they are specific to the ML algorithms used and do not generalize to other ML methods.

Recently, efforts have been made to target online GPU power optimization. The work presented in [10] proposes a dynamic power management method for integrated APUs. ODPP [11] is another online dynamic power-performance management framework. ODPP detects the periods of appli-

cations to build prediction models. However, their models can only detect coarse-grained phase change patterns.

This work aims to provide a better dynamic GPU energy optimization method for iterative ML training workloads. We present GPOEO, a micro-intrusive¹ GPU online energy optimization framework. GPOEO only requires the user to mark the beginning and the end of the target code region. GPOEO can then automatically collect the relevant information to monitor the program behavior to identify the repeat workload patterns and adjust the energy optimization scheme accordingly.

GPOEO trades execution time for energy efficiency. To this end, we employ multi-objective optimization models built on the XGBoost Classifier [12]. As a departure from existing online GPU power optimization schemes, GPOEO leverages the fine-grained runtime information provided by the hardware performance counters to model the system and program behavior to build more accurate decision models. We use the predictive model as a cost function to search for the optimal energy configuration for the current workload pattern. These supporting methods enable us to develop a period detection algorithm based on the fast Fourier transform (FFT) and feature sequence similarity evaluation to identify if a changed training iteration takes place and adjust our optimization accordingly.

We evaluate GPOEO by applying it to 71 machine learning workloads from three AI benchmark suites [13]. We test our approach on an NVIDIA RTX3080Ti GPU. Experimental results show that GPOEO can reduce GPU energy consumption by over 20% (16.2% on average), with most execution time increments less than 7% (5.1% on average) when compared to the NVIDIA default GPU scheduling strategy.

1. Micro-intrusive means we require minor changes to the program source code.

- F. Wang, W. Zhang, S. Lai, and M. Hao are with the School of Cyberspace Science at Harbin Institute of Technology, Harbin 150000, China.
E-mail: {wangfarui, wzzhang}@hit.edu.cn
- Z. Wang is with the School of Computing at University of Leeds, United Kingdom.
E-mail: z.wang5@leeds.ac.uk

This paper makes the following contributions:

- It is the first work to exploit performance counter metrics to perform online energy prediction and optimization for discrete GPUs. By minimizing the instrument and measurement overhead, our approach improves the practicability of online GPU profiling.
- It presents a robust detection algorithm to automatically detect the iterative period of ML applications from the GPU resource traces (Section 4.1).
- It showcases how an effective GPU online energy optimization framework can be constructed as a cost model to quickly identify a good energy configuration.

Online material: The GPOEO framework is publicly available at <https://github.com/ruixueqingyang/GPOEO>.

2 BACKGROUND AND MOTIVATION

2.1 GPU configuration and profiling tools

NVIDIA provides the NVML [14] and CUPTI [15] library for reconfiguring, measuring, and profiling its GPUs. The NVML [14] can set Streaming Multiprocessor (SM) clock frequency, global memory clock frequency (only certain types of NVIDIA GPUs), and powercap online. It also supports measuring power, SM utilization, and global memory utilization. The CUPTI [15] can profile performance counter. These tools make GPU online energy optimization possible. In this work, we adjust the clock frequency of the GPU SM and global memory to achieve energy saving.

2.2 Motivation

2.2.1 Energy Optimization Potential

The training phase of machine learning usually uses high-performance GPUs as accelerators. GPUs consume a lot of energy and time. To explore the energy-saving potential on GPUs, we run five applications in AIBench [13] and benchmarking-gnns [16] on all combinations of SM and memory clock frequencies and select the best configurations for minimal energy consumption within the slowdown constraint of 5%. Figure 1 shows the oracle results of energy saving, slowdown, and ED2P ($Energy \times Dealy^2$) saving. The AI_FE, AI_S2T, and SBM_GIN are relatively compute-intensive and save significant energy (14.9%-22.4%). The CLB_MLP and TSP_GatedGCN are relatively memory-intensive and save considerable energy (18.0%-26.4%). These results highlight that both compute-intensive and memory-intensive ML training workloads have the chance to save energy with acceptable slowdown.

2.2.2 Make Energy Optimization Practical

Most existing studies on GPU energy optimization are offline and need domain knowledge and complex offline profiling. Thus, users or researchers can hardly apply these studies in practice. An ideal online GPU energy optimization system should be non-intrusive to applications and transparent to users. Majumdar’s work [10] only applies to AMD APUs. Motivated by drawbacks, we propose an online energy optimization framework called GPOEO. Users only need to insert the Begin and End APIs, and the GPOEO automatically optimizes energy online.

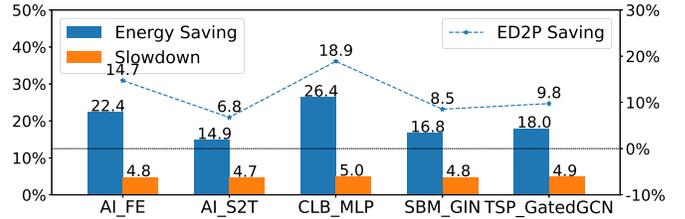


Fig. 1. Oracle ED2P saving of ML applications

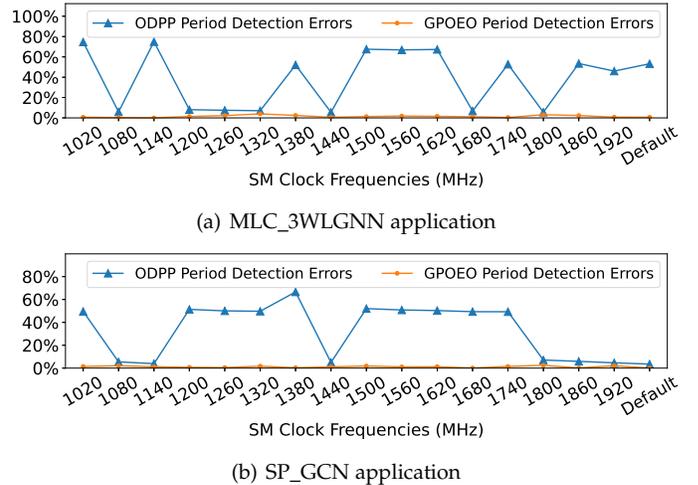


Fig. 2. Period detection errors of ODPP and GPOEO under different SM clock frequencies

2.2.3 Reduce Period Detection Error

Online sampling time series of power and utilization can reflect the periodic pattern. ODPP [11] uses the Fourier transform algorithm to detect periods, namely the execution time of one major iteration. However, according to our experimental results, the period error of ODPP [11] is quite large when the periodicity is not apparent. To solve this shortcoming, we propose a robust period detection algorithm. Figure 2 shows the absolute percentage errors of period detection with ODPP and our algorithm GPOEO under different SM clock frequencies on two ML applications. The “Default” means the NVIDIA default scheduling strategy. The period detection errors of ODPP are pretty significant and unstable, while the period detection errors of our GPOEO are less than 5% under all SM clock frequencies.

2.2.4 Using GPU Performance Counters

ODPP [11] uses power, SM utilization, and memory utilization as features to predict energy and execution time. However, these coarse-grained features are inadequate to give accurate predictions for some machine learning applications. We give an example in Figure 3. Applications in each pair have similar average power and utilization under the same reference clock frequency configuration, but their optimal SM clocks for ED2P are different (memory frequency fixed at 9251MHz). In light of this observation, we wish to use more fine-grained metrics to capture the subtle interactions between the application and the hardware. To this end, we use hardware performance counter metrics as features to build predictive models to estimate energy

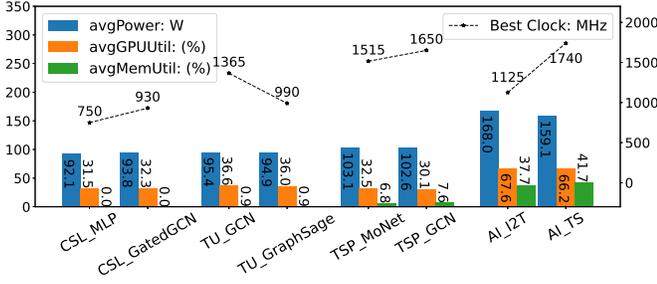


Fig. 3. ML applications with similar coarse-grained features and different optimal SM clocks for ED2P

and execution time. Later in the paper, we show that this strategy allows us to build more accurate prediction models, which in turn lead to better optimization decisions. The key challenges here are reducing the overhead of performance counter profiling and minimizing code modifications.

3 OVERVIEW OF OUR APPROACH

This section defines the scope of the work and gives an overview of our approach.

3.1 Problem Formulation

This work aims to find an optimal SM and memory clock frequency configuration to balance the GPU energy consumption and execution time for ML training workloads. Our goal can be formulated as follows:

$$\begin{aligned}
 \min \quad & F_{SM} = f_{obj}(\widehat{eng}_{SM}, \widehat{time}_{SM}) \\
 \text{s.t.} \quad & \widehat{eng}_{SM} = EngMdl_{SM}(\mathbf{w}_{SM}), \\
 & \widehat{time}_{SM} = TimeMdl_{SM}(\mathbf{w}_{SM}), \\
 & \mathbf{w}_{SM} = \{SMgear_i, \mathbf{Feature}\}, \\
 & SMgear_i \in \{SMgear_1, \dots, SMgear_m\}, \\
 & \mathbf{Feature} = \{feature_1, \dots, feature_q\}.
 \end{aligned} \tag{1}$$

$$\begin{aligned}
 \min \quad & F_{Mem} = f_{obj}(\widehat{eng}_{Mem}, \widehat{time}_{Mem}) \\
 \text{s.t.} \quad & \widehat{eng}_{Mem} = EngMdl_{Mem}(\mathbf{w}_{Mem}), \\
 & \widehat{time}_{Mem} = TimeMdl_{Mem}(\mathbf{w}_{Mem}), \\
 & \mathbf{w}_{Mem} = \{Memgear_i, \mathbf{Feature}\}, \\
 & Memgear_i \in \{Memgear_1, \dots, Memgear_p\}.
 \end{aligned} \tag{2}$$

Table 1 lists the variables used in our problem formulation. We note that our approach can be applied to an arbitrary objective function. We select the energy consumption with an execution time increase constraint as our objective function to explicitly control the performance loss. Like [17], we assume that the search space of SM and memory clocks is convex and optimize the SM frequency and memory frequency in order. We use four prediction models to estimate energy and time with different SM and memory clock frequencies. Then we select the SM and memory clock configuration that can best satisfy the optimization goal.

TABLE 1
Description of Parameters

Notation	Description
f_{obj}	objective function
\widehat{eng}_{SM}	energy consumption predicted with the SM clock
\widehat{time}_{SM}	execution time predicted with the SM clock
$EngMdl_{SM}$	energy prediction model with the SM clock
$TimeMdl_{SM}$	time prediction model with the SM clock
$SMgear_i$	one SM gear represents an SM clock frequency
\mathbf{w}_{SM}	the input vector containing SM clock
\widehat{eng}_{Mem}	energy consumption predicted with the memory clock
\widehat{time}_{Mem}	execution time predicted with the memory clock
$EngMdl_{Mem}$	energy prediction model with the memory clock
$TimeMdl_{Mem}$	time prediction model with the memory clock
$Memgear_i$	one memory gear represents a memory clock frequency
\mathbf{w}_{Mem}	the input vector containing memory clock
Feature	feature vector measured under the reference clock configuration from the ML applications

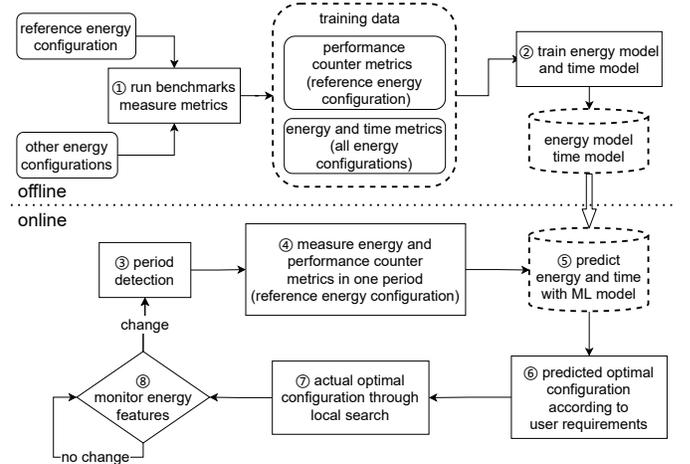


Fig. 4. Overview of the GPOEO workflow.

3.2 Framework Overview

GPOEO includes two stages, shown in Figure 4. In the offline training stage, we run representative benchmarks on each SM and memory clock frequencies (1) to collect performance counter metrics on the reference SM and memory frequencies and energy-time data on all frequencies. Then we train multi-objective models (2). In the online optimization stage, we measure energy and performance counter metrics in one detected period (3, 4), then predict the optimal SM and memory frequencies (5, 6). After that, we try frequencies around the predicted optimal configuration and compare the measured energy-time data to search for the actual optimal configuration (7). Finally, we set the actual optimal configuration and monitor energy characteristics (8). If the fluctuation of energy characteristics exceeds the threshold, another optimization process will start.

4 DESIGN AND IMPLEMENT

4.1 Robust Period Detection

Iterative ML workloads exhibit periodic behavior across iterations. We treat one period, namely one iteration, as basic unit for online metric measurement and evaluation. Our approach does not rely on offline code instrumentation to reduce developer involvement.

4.1.1 Fourier Transform: Get Candidate Periods

To detect the period change, we leverage the Fourier transform. This established methodology is widely used in signal processing to detect the dominant frequency and period. In our scenario, we assume the periodic feature of ML applications, such as GPU power and utilization, can be expressed as a time-domain function $\mathcal{F}(t)$. Generally speaking, $\mathcal{F}(t)$ is periodic and can be expressed as a linear combination of trigonometric series (know as Fourier Series) in exponential form as:

$$\mathcal{F}(t) = \sum_{n=-\infty}^{\infty} C_n e^{2\pi j n t / T_{iter}} \quad (3)$$

where C_n are Fourier coefficients.

We can get a series of frequency components via Fourier transform:

$$\mathcal{F}(\omega) = 2\pi \sum_{n=-\infty}^{\infty} C_n \delta(\omega - 2\pi n / T_{iter}) \quad (4)$$

the frequency $\mathcal{F}(\omega)$ is a series of impulse functions (i.e., δ) located at $2\pi n / T_{iter}$, with amplitude proportional to C_n . These impulses represent different frequency components in $\mathcal{F}(t)$. The one with the largest amplitude is the major frequency component. The iterative period T_{iter} of the ML application can be calculated as $T_{iter} = 1 / f_{major}$, where f_{major} is the frequency of the major frequency component.

4.1.2 Feature Sequence Similarity: Find Accurate Periods

As mentioned in Section 2.2.3, the period detection results of the Fourier transform may have vast errors. Our period calculation algorithm (Algorithm 1) combines the fast Fourier transform (FFT) and feature sequence similarity (Algorithm 2) to reduce these errors. We find amplitude peaks and corresponding periods T_{Peak} in the result of FFT (line 1-3). Then we select periods T_{peak_m} with relatively high peaks as candidate periods T_{cand} (line 4-5). We set the peak coefficient c_{peak} to 0.6-0.7 empirically. Then we use feature sequence similarity (Algorithm 2) to evaluate and select the best candidate period with minimal error. Finally, we perform a local search around the best candidate period to further improve accuracy.

We propose the feature sequence similarity algorithm (Algorithm 2) to evaluate candidate periods. We divide the feature sampling curve into several sub-curves, and the time duration of each sub-curve is equal to the candidate period (line 1-4). The more similar the curves are, the closer the candidate period is to the actual period. So we evaluate the similarity of each pair of adjacent sub-curves (line 5-18).

Each sub-curve is a time series, and calculating the Euclidean distance between corresponding points is a common method to measure the similarity of time series. However, in

Algorithm 1 Period calculation algorithm based on FFT and feature sequence similarity

Input: feature sampling sequence $\mathbf{Smp} = \{s_1, \dots, s_N\}$, sampling interval T_s

Output: iterative period T_{iter} , error of period err

- 1: $Freq, Ampl = \{freq_i\}, \{ampl_i\} = \text{FFT}(\mathbf{Smp}, T_s)$
- 2: $T = \{T_i\} = \{1/freq_i\}$
- 3: Find peaks in $Ampl$, get the set of peaks $AmplPeak = \{amplpeak_m\}$ and the set of corresponding periods $T_{Peak} = \{T_{peak_m}\}$
- 4: $IndexCand = \{idxcand_k\}$
 $= \{m | amplpeak_m > c_{peak} \times \max(AmplPeak)\}$
- 5: $TCand = \{T_{cand_p}\} = \{T_{peak_{idxcand_k}}\}$
- 6: **for** $T_{cand_p} \in TCand$ **do**
- 7: $err_p = \text{Algorithm2}(T_{cand_p}, \mathbf{Smp}, T_s)$
- 8: **end for**
- 9: Find $err_{idxmin} = \min(\{err_p\})$
- 10: $T_{cand_{opt}} = T_{cand_{idxmin}}$
- 11: $N_T = T_s \times (N - 1) / T_{cand_{opt}}$,
 $T_{low} = T_{cand_{opt}} \times (1 - 1 / (N_T + 1))$
- 12: $T_{up} = T_{cand_{opt}} \times (1 + 1 / (N_T - 1))$
- 13: Construct an arithmetic sequence $T_{Local} = \{T_{local_q}\}$,
where $T_{local_q} = T_{low} + (q - 1) \times T_{accu}$,
 $q = 1, \dots, (T_{up} - T_{low}) / T_{accu}$
- 14: **for** $T_{local_q} \in T_{Local}$ **do**
- 15: $err_{local_q} = \text{Algorithm2}(T_{local_q}, \mathbf{Smp}, T_s)$
- 16: **end for**
- 17: Find $err = err_{idxmin} = \min(\{err_{local_q}\})$
- 18: $T_{iter} = T_{local_{idxmin}}$
- 19: **return** T_{iter}, err

our scenario, the Euclidean distance is susceptible to high-frequency interference and reports the wrong similarity. Inspired by this, instead of calculating the distance of each pair of corresponding points, we group the sampling data within a series and evaluate the distance between the corresponding groups. Specifically, we use the Gaussian mixture model algorithm to cluster each sub-curves into N groups (line 8-11). Then, for each group, we calculate the relative average amplitudes (line 12-13). The average operation can eliminate the influence of high-frequency interference. Later, we calculate the symmetric mean absolute percentage error (SMAPE) of the current group pair (line 14). Finally, we calculate the weighted average of all SMAPEs as similarity error (line 16-17). The weights are the numbers of sampling points in groups. When all adjacent sub-curves are evaluated, we treat the mean similarity error as the error of T_{iter} (line 19).

4.1.3 Online Robust Period Detection Algorithm Framework

To get accurate T_{iter} dynamically, we design an online robust period detection algorithm framework (Algorithm 3). The framework calculates T_{iter} in a rolling form while sampling the power and utilization of the GPU until the T_{iter} is stable. First, we get the initial period T_{init} , using Algorithm 2. If sampling duration is short than $c_{measure}$ times T_{init} ($c_{measure}$ is set to 2 empirically), it is too short

Algorithm 2 Feature sequence similarity algorithm

Input: Period to be evaluated T_{iter} , feature sampling sequence $\mathbf{Smp} = \{s_1, \dots, s_N\}$, sampling interval T_s

Output: error of period Err_T

```

1:  $Num_T = \lfloor s_N / T_{iter} \rfloor$ ;  $Num_s = \lfloor T / T_s \rfloor$ 
2: for  $i \in \{1, \dots, Num_T\}$  do
3:    $\mathbf{Smp}_i = \{s_{(i-1) \times Num_s + 1}, \dots, s_{i \times Num_s}\} = \{ss_{i,j}\}$ 
      ( $j = 1, \dots, Num_s$ )
4: end for
5: for  $i \in \{1, \dots, Num_T - 1\}$  do
6:    $Mean_{prev} = \text{mean}(\mathbf{Smp}_i)$ 
7:    $Mean_{back} = \text{mean}(\mathbf{Smp}_{i+1})$ 
8:    $\{\mathbf{GaGrp}_1, \dots, \mathbf{GaGrp}_{Num_G}\} = \text{Gauss}(\mathbf{Smp}_i, Num_G)$ 
      where  $\mathbf{GaGrp}_j = \{idx_{j,1}, \dots, idx_{j,Num_j}\}$ 
9:   for  $j \in \{1, \dots, Num_G\}$  do
10:     $\mathbf{Grp}_{prev} = \{ss_{i,idx_{j,1}}, \dots, ss_{i,idx_{j,Num_j}}\}$ 
11:     $\mathbf{Grp}_{back} = \{ss_{i+1,idx_{j,1}}, \dots, ss_{i+1,idx_{j,Num_j}}\}$ 
12:     $RelValPrev_j = \text{mean}(\mathbf{Grp}_{prev}) - Mean_{prev}$ 
13:     $RelValBack_j = \text{mean}(\mathbf{Grp}_{back}) - Mean_{back}$ 
14:     $grperr_j = \text{SMAPE}(RelValPrev_j, RelValBack_j)$ 
15:   end for
16:    $Weight = \{|\mathbf{GaGrp}_1|, \dots, |\mathbf{GaGrp}_{Num_G}|\}$ 
17:    $err_i = \text{avg}(\{grperr_1, \dots, grperr_{Num_G}\}, Weight)$ 
18: end for
19: return  $Err_T = \text{mean}(\{err_1, \dots, err_{Num_T-1}\})$ 

```

Note: Gauss = the Gaussian mixture model clustering algorithm

to do a rolling calculation, so we calculate the required sampling duration, skip rolling calculation, and return (line 2-6). If the sampling duration is long enough, we get some updated samples by setting t_{start} (line 7). Samples before t_{start} are ignored because they may be outdated. We set $step = 0.5$ and $c_{eval} = 6.5$ according to our experiments. Then, the rolling period calculation begins. In each iteration, we calculate the period (line 9-11) and increase t_{start} to exclude more outdated samples for the next iteration (line 12). With the rolling period calculation, we get a sequence of periods with errors (line 14). If an ML application runs in the regular iteration phase, these periods should be similar. If the difference of these periods is less than the threshold, set $SmpDur_{next} = -1$ to indicate stop sampling. Otherwise, calculate $SmpDur_{next}$ (line 16-21). The best period, namely T_{iter} , is the period with the minimal error (line 15). If $SmpDur_{next} > 0$, we ignore T_{iter} , keep sampling, and call Algorithm 3 again after $SmpDur_{next}$. Otherwise, we use T_{iter} for feature measurement (Section 4.2).

4.2 Micro-intrusive Online Feature Measurement

Based on the robust period detection, we propose the adaptive feature measurement algorithm (Algorithm 4). The feature measurement causes overhead and extends T_{iter} compared with the normal run of the application. So, we adaptively detect T_{iter} of applications while the measurement is running (line 1-7). In line 5, we collect $\mathbf{Feature}_{dect}$ after a delay of $SmpDur_{Next}$. We use the sampling sequence of $\mathbf{Feature}_{dect}$ to form the curve for the period detection, so we must sample $\mathbf{Feature}_{dect}$ continuously and uniformly. On the NVIDIA GPU platform, the instantaneous power, SM utilization, and memory utilization can meet

Algorithm 3 Online robust period detection algorithm framework

Input: feature sampling sequence $\mathbf{Smp} = \{s_1, \dots, s_N\}$, sampling interval T_s

Output: iterative period T_{iter} , next sampling duration $SmpDur_{next}$

```

1:  $T_{init}, err_{init} = \text{Algorithm1}(\mathbf{Smp}, T_s)$ 
2:  $SmpDur = (N - 1) \times T_s$ 
3: if  $SmpDur < c_{measure} \times T_{init}$  then
4:    $T_{iter} = T_{init}$ 
5:    $SmpDur_{next} = c_{measure} \times T_{init} - SmpDur$ 
6: end if
7:  $t_{start} = \max(0, (SmpDur - (2 + c_{eval} \times step) \times T_{init}))$ 
8: while  $(SmpDur - t_{start}) / T_{init} \geq c_{measure}$  do
9:    $istart = 1 + \lfloor t_{start} / T_s \rfloor$ 
10:   $\mathbf{SubSmp} = \{s_{istart}, \dots, s_N\}$ 
11:   $T_j, err_j = \text{Algorithm1}(\mathbf{SubSmp}, T_s)$ 
12:   $t_{start} = t_{start} + step \times T_{init}$ 
13: end while
14:  $T = \{T_j\}, Err = \{err_j\}$ 
15: Find  $\min(Err) = err_k$  then  $T_{iter} = T_k$ 
16:  $Diff = \text{abs}((\max(T) - \min(T)) / \text{mean}(T))$ 
17: if  $Diff < Diff_{threshold}$  then
18:    $SmpDur_{next} = -1$ 
19: else
20:    $SmpDur_{next} =$ 
       $\lfloor SmpDur / \max(T) \rfloor \times \max(T) - SmpDur$ 
21: end if
22: return  $T_{iter}, SmpDur_{next}$ 

```

Algorithm 4 Adaptive feature measurement

Input: features to be measured $FeatureName = \{name_1, \dots, name_n\}$, the feature for iterative period detection $\mathbf{Feature}_{dect}$

Output: feature data $\mathbf{Feature} = \{feature_1, \dots, feature_n\}$

```

1: Start measurement for  $FeatureName$  and  $\mathbf{Feature}_{dect}$ 
2:  $SmpDur_{next} = SmpDur_{init}$ 
3: while  $SmpDur_{next} > 0$  do
4:   Delay  $SmpDur_{next}$ 
5:   Collect sampling sequence of  $\mathbf{Feature}_{dect}$ 
6:    $T_{iter}, SmpDur_{next} = \text{Algorithm3}(\mathbf{Feature}_{dect}, T_s)$ 
7: end while
8: Restart measurement for  $FeatureName$ 
9: Delay  $T_{iter}$  then stop measurement
10: Collect feature data  $\mathbf{Feature}$ 
11: return  $\mathbf{Feature}$ 

```

this requirement. However, the performance counters cannot meet this requirement because its minimum sampling granularity is a CUDA kernel. According to our experimental results, we use the composite feature of power, SM utilization, and memory utilization as $\mathbf{Feature}_{dect}$, whose traces show more obvious periodicity. In line 6, we call Algorithm 3 to detect T_{iter} and determine whether T_{iter} is stable. If $SmpDur_{Next} > 0$, we redo line 4-6 with the new $SmpDur_{Next}$. Otherwise, we think T_{iter} is stable and restart the feature measurement (line 8). After a delay of T_{iter} , we stop measurement and collect the feature data (line 8-10).

We implement micro-intrusive online feature measure-

ment based on Algorithm 4. Users only need to insert the Begin and End APIs of GPOEO at the beginning and end of their source code, rather than instrument loops.

4.3 Training Data Collection and Model Generation

We establish energy consumption and execution time prediction models with the boosting machine learning method. Our work aims to find the optimal SM and memory clock gears that optimize the objective function. Therefore, these four models only need to predict the relative energy consumption and execution time relative to the NVIDIA default scheduling strategy rather than the absolute energy and time [4].

To train the models defined in Equation 1 and 2, we define and collect four training data sets:

$$EngTr_{SM} = \{EngSM_j^i\} \quad (5)$$

$$TimeTr_{SM} = \{TimeSM_j^i\} \quad (6)$$

$$EngTr_{Mem} = \{EngMem_j^i\} \quad (7)$$

$$TimeTr_{Mem} = \{TimeMem_j^i\} \quad (8)$$

where $EngSM_j^i$ or $TimeSM_j^i$ represents a piece of training data collected from application i under the SM clock gear $SMgear_j$ and the memory clock controlled by the NVIDIA default scheduling strategy. $EngMem_j^i$ or $TimeMem_j^i$ is collected under the memory clock gear $Memgear_j$ and the optimal SM clock gear. $eng_{default}^i$ and $time_{default}^i$ are the default energy consumption and execution time under the NVIDIA default scheduling strategy. Each input vector $\mathbf{w}_{SM_j^i}$ or $\mathbf{w}_{Mem_j^i}$ contains $SMgear_j$ or $Memgear_j$ and the feature vector $\mathbf{Feature}^i$. For each application, the feature vector $\mathbf{Feature}^i$ is measured under the reference clock configuration.

$$EngSM_j^i = \{\mathbf{w}_{SM_j^i}, eng_{SM_j^i} / eng_{default}^i\} \quad (9)$$

$$TimeSM_j^i = \{\mathbf{w}_{SM_j^i}, time_{SM_j^i} / time_{default}^i\} \quad (10)$$

$$\mathbf{w}_{SM_j^i} = \{SMgear_j, \mathbf{Feature}^i\} \quad (11)$$

$$EngMem_j^i = \{\mathbf{w}_{Mem_j^i}, eng_{Mem_j^i} / eng_{default}^i\} \quad (12)$$

$$TimeMem_j^i = \{\mathbf{w}_{Mem_j^i}, time_{Mem_j^i} / time_{default}^i\} \quad (13)$$

$$\mathbf{w}_{Mem_j^i} = \{Memgear_j, \mathbf{Feature}^i\} \quad (14)$$

4.3.1 Feature Selection

As mentioned in Section 2.2.4, the high-level features, such as power, SM utilization, and memory utilization, cannot provide enough information for modeling. So we introduce performance counter metrics as input features.

The CUPTI [15] supports profiling over 1100 performance counters on NVIDIA high-end GPUs, such as RTX2080Ti and RTX3080Ti. Inspired by Arafa's work [7], we treat each Parallel Thread Execution (PTX) instruction as a basic unit of energy consumption and focus on performance counter metrics that reflect the density of different types of PTX instructions. We list the selected metrics in Table 2. A metric with prefix `sm__inst_executed_pipe_` represents a type of PTX instruction. The suffix `pct_of_peak_sustained_active` means the percentage of actual mean instruction throughput to the theoretical maximum sustained instruction throughput in an activity clock cycle [15].

TABLE 2
Selected Features

Notation	Full name or expression
IPCpct	<code>sm__inst_executed.PctSus</code>
L1MissPerInst	<code>l1tex__t_sectors_lookup_miss.sum / sm__inst_executed.sum</code>
L1MissPct	<code>(l1tex__t_sectors_lookup_miss.sum + l1tex__t_sectors_lookup_hit.sum)</code>
L2MissPerInst	<code>lts__t_sectors_lookup_miss.sum / sm__inst_executed.sum</code>
L2MissPct	<code>(lts__t_sectors_lookup_miss.sum + lts__t_sectors_lookup_hit.sum)</code>
ALUPct	<code>sm__inst_executed_pipe_alu.PctSus</code>
ADUPct	<code>sm__inst_executed_pipe_adu.PctSus</code>
FP16Pct	<code>sm__inst_executed_pipe_fp16.PctSus</code>
FMAPct	<code>sm__inst_executed_pipe_fma.PctSus</code>
FP64Pct	<code>sm__inst_executed_pipe_fp64.PctSus</code>
XUPct	<code>sm__inst_executed_pipe_xu.PctSus</code>
TNSPct	<code>sm__inst_executed_pipe_tensor.PctSus</code>
CBUPct	<code>sm__inst_executed_pipe_cbu.PctSus</code>
LSUPct	<code>sm__inst_executed_pipe_lsu.PctSus</code>
TEXPct	<code>sm__inst_executed_pipe_tex.PctSus</code>
UNIPct	<code>sm__inst_executed_pipe_uniform.PctSus</code>

Note: $PctSus = sum.pct_of_peak_sustained_active$

According to related work [4], DRAM and L2 cache miss information may help predict energy and time. However, profiling performance counters of DRAM and FBPA need several kernel replays which cannot be implemented online. We design several composite metrics as alternatives, also shown in Table 2. The `L1MissPerInst` or `L2MissPerInst` means the number of L1 or L2 cache miss per instruction. The `L1MissPct` or `L2MissPct` means the percentage of L1 or L2 cache miss. In conclusion, we use metrics listed in Table 2 as features.

4.3.2 Benchmark Selection and Training Data Collection

We use the PyTorch Benchmarks [18], which contain over 40 mini ML applications, for training. We select the AIBench Training Component Benchmark [13], the benchmarking-gnns [16], a ThunderSVM [19] workload, and a ThunderGBM [20] workload as the testing set. The benchmarking-gnns is a graph neural networks (GNN) benchmarking framework including seven datasets (CLB, CSL, SBM, TSP, TU, MLC, and SP) and nine models.

For each application in the training set, we profile features under the reference clock configuration. We measure the actual energy ($eng_{SM_j^i}$) and time ($time_{SM_j^i}$) under all available SM clock gears and the memory clock controlled by the NVIDIA default scheduling strategy. We also measure the actual energy ($eng_{Mem_j^i}$) and time ($time_{Mem_j^i}$) under all available memory clock gears and the optimal SM clock gear. Feature, energy, and time data construct four training data sets ($EngTr_{SM}$, $TimeTr_{SM}$, $EngTr_{Mem}$, and $TimeTr_{Mem}$). We conduct each measurement ten times and take the average to reduce the training data noise.

4.3.3 Prediction model

We construct energy and time prediction models with the XGBoost [12], a widely used supervised machine learning

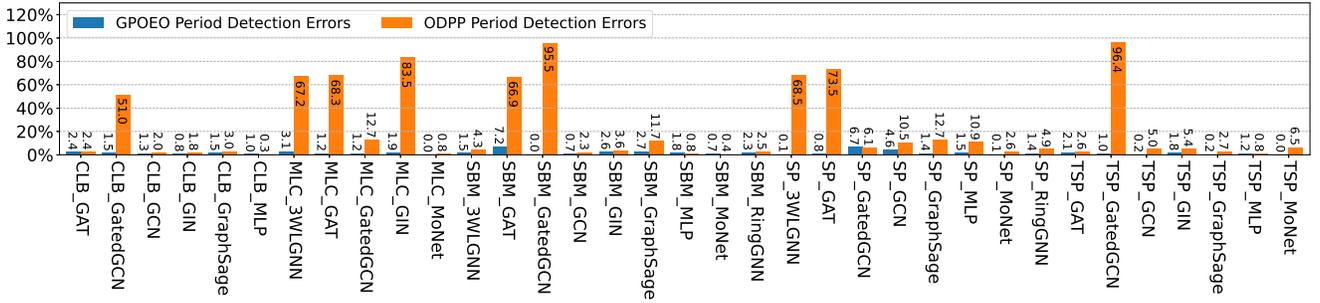


Fig. 5. The period detection errors of GPOEO and ODPP. GPOEO gives lower errors compared to ODPP.

model improved from the gradient tree boosting algorithm. Combined with our scenario, we apply the XGBoost to build four prediction models shown in Equation 1 and 2. The model construction processes are similar, so we take the energy consumption model $EngMdl_{SM}$ as an example to introduce these processes. For $EngMdl_{SM}$, the XGBoost weights and sums several regression trees and can be defined as follows:

$$\widehat{eng}_{SM_j^i} = \sum_{k=1}^K f_k(\mathbf{w}_{SM_j^i}), f_k \in \mathcal{F} \quad (15)$$

where $\widehat{eng}_{SM_j^i}$ is the predicted energy consumption of the input vector $\mathbf{w}_{SM_j^i}$, which contains the features of application i and the SM gear j . f_k represents the k th regression tree, and K is the number of regression trees. We use the training datasets $EngTr_{SM}$ to train the model $EngMdl_{SM}$. During training, we use the grid search method to tune the hyper-parameters such as the minimum loss reduction, the maximum depth of a tree, the minimum sum of instance weight, and the maximum number of nodes to be added.

4.3.4 Online Local Search

To fix minor errors caused by multi-objective predicted models, we perform a local search around the predicted optimal clock configuration. We use online detected period and energy data measured during one period to evaluate different clock frequencies and find the actual optimal clock configuration. We first conduct a local search for the memory clock because wrong memory clock frequencies (too low) can lead to severe slowdowns. Based on the optimal memory clock frequency, we conduct another local search for the SM clock. According to our experimental results and Schwarzrock’s work [17], the relationship between energy optimization objectives and clock gears is generally a convex function. Based on this observation, we use the golden-section search method [21] to accelerate our local search. We first find a gear with a worse objective value on each side of the predicted optimal gear to determine the search interval. Then we follow the classic golden section-search process [21] to attempt different clock gears. Due to possible errors of measured energy and periods, we fit the attempted points to obtain a convex function. Then we use the convex function to determine the optimal clock gear.

4.3.5 Optimize aperiodic applications

The power and utilization traces of some applications are aperiodic, such as applications in CSL and TUs datasets,

the ThunderSVM [19] workload, and the ThunderGBM [20] workload. We can not evaluate their execution time or energy consumption with data measured in one period. In this case, we fix the measurement time interval and the use measured mean number of instructions executed per second (IPS) and power to evaluate the execution time and energy consumption. If the number of instructions in the program is $Inst_{sum}$, the execution time can be calculated as $time = Inst_{sum}/IPS$, and the energy consumption can be calculated as $energy = power * time = Inst_{sum} * power/IPS$. Then we can evaluate different clock frequency configurations.

Based on these analyses, we apply GPOEO to aperiodic applications. We first measure performance counter metrics in the fixed time interval and predict the optimal clock frequency configuration. In the online local search, we measure and calculate $time_{default}$ and $energy_{default}$ under the NVIDIA default scheduling strategy as the baseline. Then, we follow the golden-section search strategy and collect $time_i$ and $energy_i$ under different clock frequency configurations near the predicted optimal clock frequency configuration. Finally, we compare these $time_i$ and $energy_i$ with the baseline to find the optimal clock frequency configuration.

5 EVALUATION

We evaluate our GPOEO system in this section. We first introduce the experimental step. Then, we evaluate the accuracy of the period detection algorithm and energy-performance prediction models. Finally, we analyze the online energy-saving results on various ML benchmarks.

5.1 Experimental Setup

5.1.1 Hardware and Software Platforms

Our experimental platform is a GPU server equipped with one NVIDIA RTX3080Ti GPU, one AMD 5950X CPU, and 64 GB memory. The software environment is NVIDIA Driver 470.57 and CUDA 11.3. RTX3080Ti supports continuously adjustable SM clock frequency, from 210 MHz to 2,025 MHz, and the step is 15 MHz. We find that some higher frequencies are not practical or stable. Under lower frequencies, applications cannot save energy while suffering severe slowdowns. Therefore, we only consider the frequencies in the middle part, which can run stably and may improve energy efficiency. We treat each SM clock frequency as an SM clock gear, from $SMgear_{16} = 450$ MHz to $SMgear_{114}$

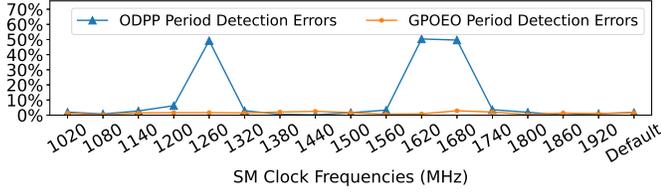


Fig. 6. Period detection errors of ODPP and GPOEO on the CLB_GAT application under different SM clock frequencies

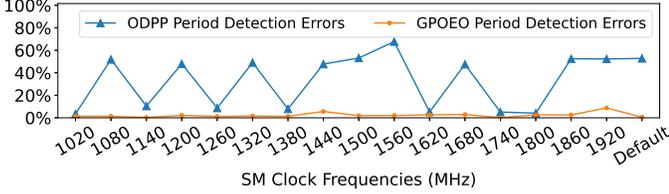


Fig. 7. Period detection errors of ODPP and GPOEO on the SBM_3WLGNN application under different SM clock frequencies

= 1,920 MHz. RTX3080Ti supports 5 global memory clock frequencies: $Memgear_0 = 450$ MHz, $Memgear_1 = 810$ MHz, $Memgear_2 = 5,001$ MHz, $Memgear_3 = 9,251$ MHz, and $Memgear_4 = 9,501$ MHz. We select $SMgear_{106} = 1,800$ MHz and $Memgear_3 = 9,251$ MHz as reference frequencies, which are used in performance counter profiling.

5.1.2 Benchmarks

We train our machine learning models using data collected from PyTorch Benchmarks [18]. We then test the trained models on AIBench [13] and benchmarking-gnns [16]. In addition to deep neural networks, we also use two classical ML workloads, the ThunderSVM [19] and ThunderGBM [20], in our evaluation.

5.2 Accuracy and Sensitivity of Period Detection

We analyze our period detection algorithm on 34 different ML applications. Figure 5 shows the period detection errors of GPOEO and ODPP [11] using the NVIDIA default scheduling strategy. GPOEO is far more accurate than ODPP. The mean period error of GPOEO is 1.72%, while the mean error of ODPP is 23.16%. The maximum error of GPOEO is 7.2%, and 32 errors are within 5%. For ODPP, nine errors are above 50%, nine errors are among 5-13%, and sixteen errors are within 5%. GPOEO is more accurate than ODPP on 29 applications. For the other three applications, the accuracy differences are within 1%.

Figure 6, 7, and 8 show the periods detection errors of GPOEO and ODPP under varying SM clock frequencies and fixed memory clock frequency. GPOEO exhibits stable low errors (within 5% for most cases). This phenomenon indicates that GPOEO is sensitive to clock frequency changes. However, ODPP shows poor stability. On CLB_GAT and SBM_3WLGNN, ODPP gains huge errors under three and ten SM clock frequencies, respectively. This instability causes inaccurate prediction models, which ODPP builds online. For TSP_GatedGCN, ODPP gains huge errors (nearly 100%) under all frequencies because it is

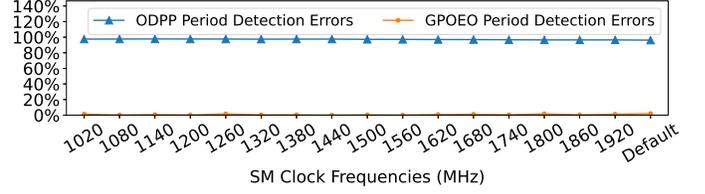
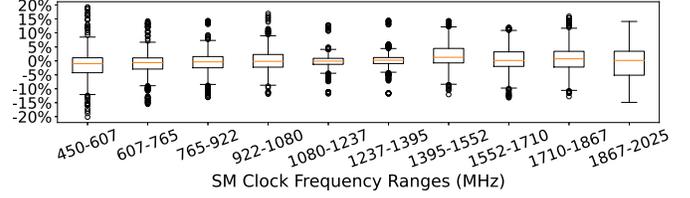
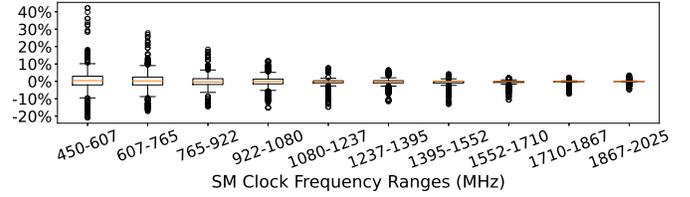


Fig. 8. Period detection errors of ODPP and GPOEO on the TSP_GatedGCN application under different SM clock frequencies

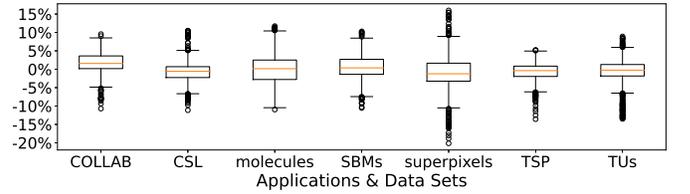


(a) Energy consumption prediction errors

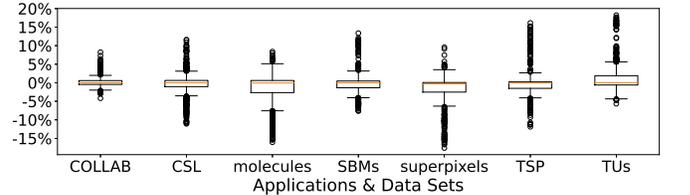


(b) Execution time prediction errors

Fig. 9. Prediction errors by varying the SM clock (grouped by different SM clock ranges)



(a) Energy consumption prediction errors



(b) Execution time prediction errors

Fig. 10. Prediction errors by varying the SM clock (grouped by different datasets)

disturbed by high-frequency noise, only detects the local short period, and ignores the real global period.

5.3 Accuracy of Energy and Execution Time Prediction

In this section, we evaluate four prediction models. We measure selected features (Table 2) during one iteration of each ML application under the reference SM and memory clock frequency (1800 MHz and 9251MHz). Then we use these features and different SM and memory clock frequencies as input of our multi-objective models to predict the relative energy consumption and execution time relative

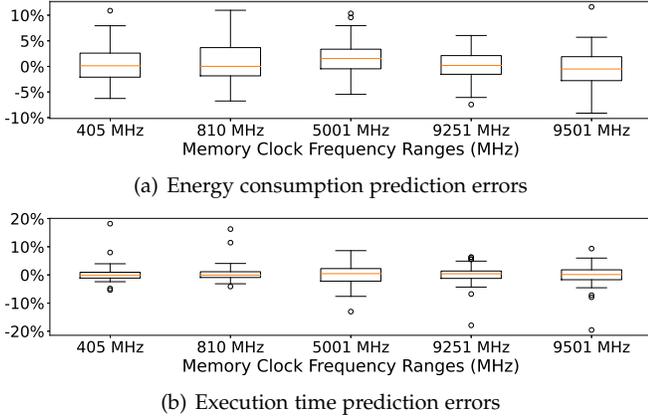


Fig. 11. Prediction errors by varying the memory clock (grouped by different memory clock frequencies)

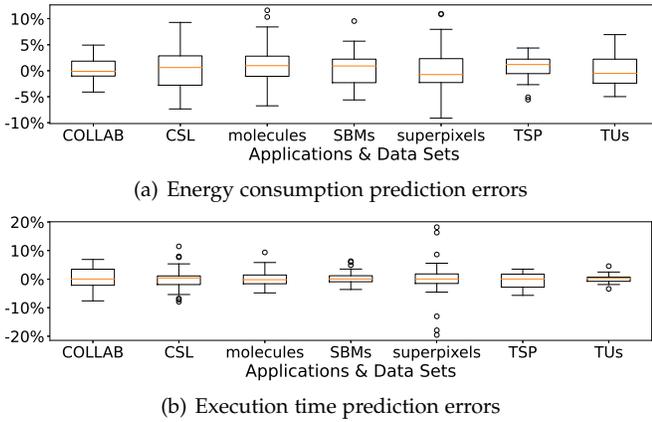


Fig. 12. Prediction errors by varying the memory clock (grouped by different datasets)

to the NVIDIA default scheduling strategy. Note that for the SM clock models $EngMdl_{SM}$ and $TimeMdl_{SM}$, we let the NVIDIA default scheduling strategy control the memory clock and then predict and select the optimal SM clock frequency. For the memory models $EngMdl_{Mem}$ and $TimeMdl_{Mem}$, we assume that the SM clock is already set to the optimal frequency and then predict and select the optimal memory clock frequency.

For the SM clock models, we collect 10,890 prediction values in total (2 objectives, 55 applications, and 99 SM clock frequencies). The mean prediction errors of energy consumption and execution time are 3.05% and 2.09%, respectively. Next, we analyze the distributions of errors sorted by SM clock frequency ranges and application datasets. Figure 9 shows the prediction errors of energy consumption and execution time separately grouped by different SM clock frequency ranges. Each clock frequency range contains 495 or 550 prediction values depending on the number of energy gears included in the clock range (nine or ten gears). For all clock ranges, most energy and time prediction errors are less than 5%. Most scattered errors are within 20%. This phenomenon proves that our energy prediction model and execution time prediction model are accurate for different SM clock frequencies. Figure 10 shows the prediction errors grouped by different application datasets. Each group

contains 594 to 891 prediction values depending on the number of applications included in the datasets (six to nine applications). For all application datasets, most energy and execution time prediction errors are less than 5%. This phenomenon demonstrates that our SM clock models are applicable to all these applications.

For the memory clock models, we collect 550 prediction values in total (2 objectives, 55 applications, and 5 memory clock frequencies). The mean prediction errors of energy consumption and execution time are 2.72% and 2.31%, respectively. Figure 11 shows the prediction errors of energy consumption and execution time separately grouped by different memory clock frequencies. For all clock frequencies, most energy and time prediction errors are less than 5%. Figure 12 shows the prediction errors grouped by different application datasets. For all application datasets, most prediction errors are also within 5%. These phenomena demonstrate that our memory clock models are applicable to all memory frequencies and all these applications.

5.4 Results of Online Optimization

Our multi-objective models predict the energy consumption and execution time of one iteration under different SM and memory clock gears. Users can use diverse energy efficiency objectives to select the predicted optimal energy gears and guide local search to find the actual optimal energy gears. In this experiment, we set the objective function to minimize the energy consumption within the slowdown constraint of 5%. For each ML training application, we use the energy consumption and execution time under the NVIDIA default scheduling strategy as the baseline. Then we run these applications with our GPOEO system to optimize the energy consumption online. We also implement the ODPP [11] as a comparison.

5.4.1 Medium benchmark suite

Figure 13 shows the online optimization results of the AIBench, the ThunderSVM workload, and the ThunderGBM workload. We use energy saving, slowdown (execution time increase), and ED^2P ($ED^2P = Energy \times Time^2$) saving to evaluate GPOEO and ODPP [11]. GPOEO achieves an average energy saving of 14.7% and an average ED^2P saving of 6.8%, with an average execution time increase of 4.6%. Especially, GPOEO achieves significant energy saving ($\geq 20\%$) on four out of sixteen applications. On AI_I2T, GPOEO achieves 29.5% energy saving and 0.1% slowdown. Table 3 shows the online optimization process for SM and memory clocks. On AI_I2T, the error of our multi-objective prediction models is only -2 SM gears. This error is eliminated by the online local search within five steps so that the vast energy saving can be achieved. For AI_3DFR, AI_3DOR, AI_I2IP, and AI_TS, GPOEO obtains considerable energy savings (18-24%) with reasonable slowdowns (less than 8%). The online optimization processes of these applications are similar to AI_I2T. The prediction errors of SM gear are relatively small (within 11 SM gears), and the online local search reduces these errors in three to five steps.

AI_FE, AI_ICMP, AI_IGEN, AI_LRK, AI_OBJ, AI_S2T, and AI_ST have medium energy-saving opportunities (10-21%) with the slowdown constraint of 5%. GPOEO gains

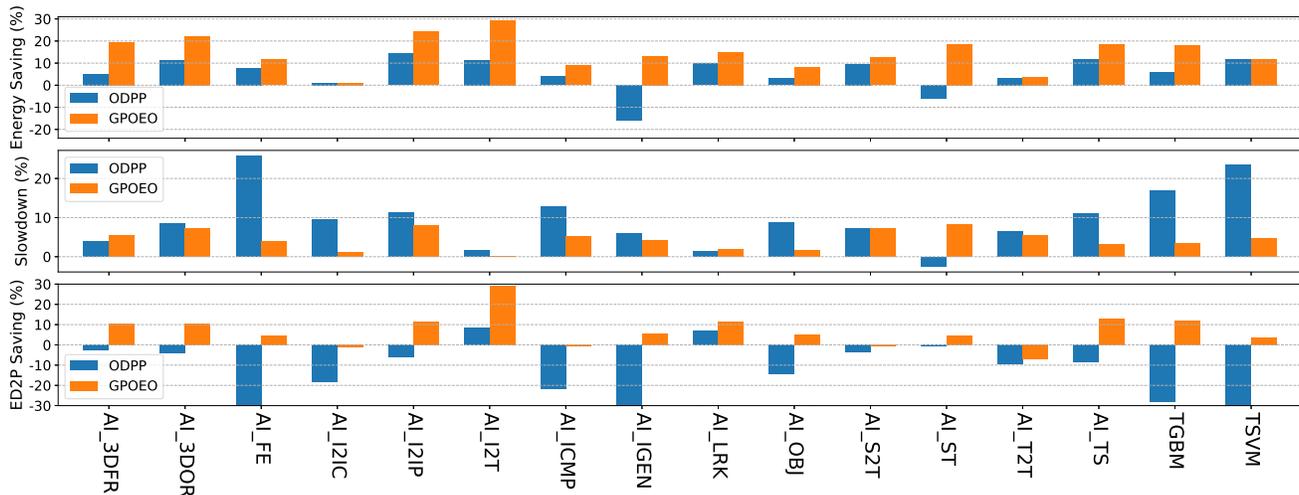


Fig. 13. Energy savings, slowdown, and ED2P saving compared to the NVIDIA default scheduling strategy on AIBench and traditional ML applications

TABLE 3
Online optimization process for SM and memory clock on AIBench

	3DFR	3DOR	FE	I2IC	I2IP	I2T	ICMP	IGEN	LRK	OBJ	S2T	ST	T2T	TS
Oracle SM Gear	104	102	109	104	91	94	107	71	88	100	108	39	100	107
Prediction Error (SM gear)	-11	5	-13	2	3	-2	-13	22	24	-4	-15	11	5	-5
Search Error (SM gear)	-2	-3	-8	1	-2	0	-2	-3	-2	0	-7	4	-2	-1
# of Search Steps (SM)	5	4	4	3	4	3	6	8	9	4	5	5	4	4
Oracle Mem clock (MHz)	9501	9251	9501	9501	9251	9251	9251	405	5001	9251	9251	810	9251	9501
Predicted Mem clock	9501	9501	9251	9501	9501	9501	9501	405	5001	9251	9501	405	9251	9501
Searched Mem clock	9501	9501	9501	9501	9501	9251	9251	405	5001	9251	9251	810	9251	9501
# of Search Steps (Mem)	2	2	3	2	2	2	2	3	3	3	2	3	3	2

medium energy savings (8-14%) with reasonable slowdowns (2-7%) on these seven applications. Two reasons cause GPOEO does not make full use of energy-saving potentials. For AI_FE and AI_S2T, the period detection and feature measurement get inaccurate data due to some abnormal iterations. Therefore prediction errors are significant and local search can not eliminate these errors. For the convenience of experiments, we just train each machine learning model for a few iterations, and only iterations after optimization processes can enjoy the optimal SM gear. However, in real-life scenarios, the model will be trained on many more iterations, for which our approach can give higher energy saving. GPOEO does not perform well on AI_I2IC and AI_T2T. These two applications are computation-intensive with little energy-saving opportunities (1% and 4%). For non-periodical applications TGBM and TSVM, GPOEO also achieves good results, saving 17.9% and 12.0% energy with slowdowns of 3.5% and 4.7%, respectively. As for the memory clock, the oracle frequencies of thirteen applications are 9,251 MHz or 9,501 MHz. Our method finds oracle frequencies on eleven applications and finds nearby 9,501 MHz on the other two applications within two or three search steps. In fact, the energy consumption and execution time are similar under these two memory clock frequencies. GPOEO also finds low oracle memory frequencies on AI_GEN, AI_LRK, and AI_ST within three steps. Our method meets the performance loss constraint of 5% on

nine applications, while ODPP only achieves the same goal on four applications. For GPOEO, only iterations after optimization processes are guaranteed to meet the performance loss constraint, and the previous iterations may result in a violation of the performance loss constraint.

Compared to ODPP, our approach can improve the energy saving by 9.2%, improve the ED^2P saving by 20.4%, and reduce the slowdown by 5.0% on average. GPOEO performs better than ODPP on all sixteen applications. Especially, our approach gains both larger energy savings and lighter slowdowns on eight applications. For AI_I2IC and TSVM, ODPP builds two piecewise linear models online to predict energy and time. The accuracy of these two models is highly dependent on the period detection accuracy. Poor period detection accuracy of ODPP causes significant prediction errors, so ODPP shows heavier slowdowns and less energy saving. On non-periodical applications TGBM and TSVM, ODPP shows much worse results than GPOEO. This phenomenon demonstrates that ODPP cannot handle non-periodical applications, while GPOEO can.

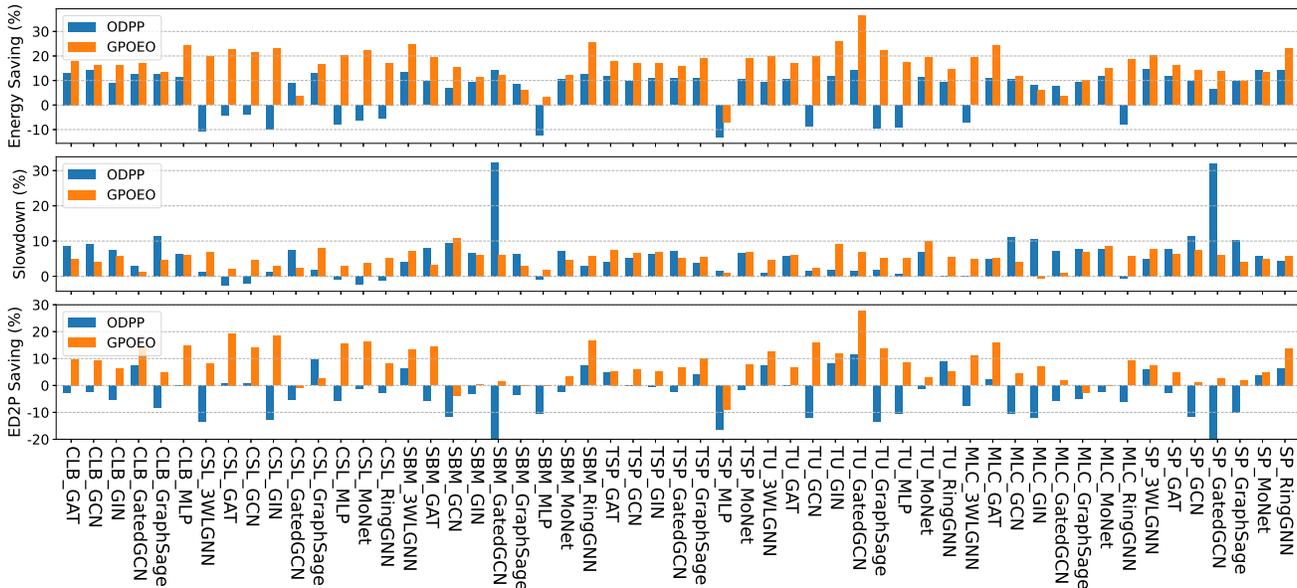


Fig. 14. Energy saving, slowdown, and ED2P saving compared to the NVIDIA default scheduling strategy on benchmarking-gnns applications

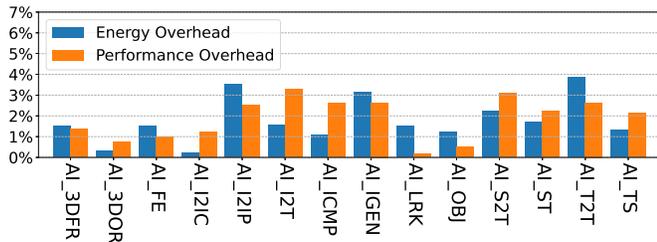


Fig. 15. Energy and time overhead of GPOEO on AIBench

5.4.2 Large benchmark suite

We also use the benchmarking-gnns suite to evaluate our method and ODPP [11] extensively. Figure 14 shows the online optimization results of 55 applications in the benchmark-gnns. GPOEO achieves an energy saving of 16.6% and an ED^2P saving of 7.8%, with an execution time increase of 5.2% on average. In comparison, ODPP gains an energy saving of 6.1% and an ED^2P saving of -4.5%, with an execution time increase of 5.6% on average. GPOEO saves ED^2P on 50 applications, while ODPP only gains ED^2P saving on 19 applications. Our method meets the performance loss constraint of 5% on 22 applications. ODPP achieves the same goal on 27 applications because it does not profile performance counters and has a lower measurement overhead. With the help of performance counter features, GPOEO saves more energy than ODPP on 48 applications. Applications in CSL and TU datasets are non-periodical. ODPP can not handle these applications and only saves 0.73% of energy with a slowdown of 1.32% on average. On the contrary, GPOEO gains an energy saving of 20.0% with a slowdown of 5.2% on average. These phenomena show that GPOEO has better performance and broader applicability than ODPP.

5.5 Overhead Evaluation

To evaluate the overhead of the GPOEO system, we run AIBench applications with the entire GPOEO system except adjusting SM and memory clocks. We also manually set the number of local search steps shown in Table 3. As shown in Figure 15, all energy and time overheads are within 4%.

6 RELATED WORK

GPU performance counter profiling is expensive. Most GPU energy optimization work [4], [6], [7], [22] is offline. Most offline work [4], [6], [22] concentrates on modeling and predicting energy and performance under different configurations. Other work studies GPU energy and performance from the view of assembly instructions [6], [7].

A few studies [10], [11] realize GPU online energy consumption optimization. Majumdar’s *et al.* [10] exploit low-overhead fine-grained profiling functions supported by the APU platform. It obtains the stream of energy, execution time, throughput, and performance counter metrics at kernel granularity without worrying about time and energy overhead. Therefore it can catch energy-saving chances at kernel granularity and save considerable energy with low overhead. Majumdar’s work cannot be transplanted to the NVIDIA GPU platform, considering the significant amount of code instrumentation (around each kernel) and the high overhead of continuously profiling. Our GPOEO targets the NVIDIA GPU platform with high-overhead online profiling (time overhead > 8% and energy overhead > 10%). We profile performance counter metrics in only one period to minimize overhead. We set one optimal energy configuration for one iteration phase, so we may miss energy saving chances within the iteration period. ODPP [11] proposes an online dynamic power-performance management framework for GPUs using coarse-grained features.

ML dedicated energy optimization work [8], [9], [23]–[25] also has emerged. These studies explore hyperparame-

ter and model structure configurations in the training stage to improve the energy efficiency in the inference stage.

7 CONCLUSION

We have presented GPOEO, a new online GPU energy optimization framework for iterative machine learning (ML) workloads. GPOEO detects iterative periods, measures performance counter features, and predicts the optimal energy configurations automatically. We evaluate GPOEO on 71 ML applications running on an NVIDIA RTX3080Ti GPU. GPOEO achieves a mean energy saving of 16.2% at a modest cost of 5.1% performance loss compared to the NVIDIA default scheduling strategy. The multi-objective prediction models need offline training, so future work includes exploring model free methods.

ACKNOWLEDGMENTS

This work was supported in part by the National Key Research and Development Program of China (2020YFB1406902), the Key-Area Research and Development Program of Guangdong Province (2020B0101360001), the Shenzhen Science and Technology Research and Development Foundation (JCYJ20190806143418198), the National Natural Science Foundation of China (NSFC) (61872110 and 61872294), the Fundamental Research Funds for the Central Universities (Grant No. HIT.OCEF.2021007), the Peng Cheng Laboratory Project (PCL2021A02), and the CCF-Huawei fund (CCF-HuaweiHP2021002). Professor Weizhe Zhang is the corresponding author.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] M. Z. Alom, T. M. Taha, C. Yakopcic, S. Westberg, P. Sidike, M. S. Nasrin, M. Hasan, B. C. Van Essen, A. A. Awwal, and V. K. Asari, "A state-of-the-art survey on deep learning theory and architectures," *Electronics*, vol. 8, no. 3, p. 292, 2019.
- [3] J. Guerreiro, A. Ilic, N. Roma, and P. Tomas, "Gpgpu power modeling for multi-domain voltage-frequency scaling," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2018, pp. 789–800.
- [4] —, "Modeling and decoupling the gpu power consumption for cross-domain dvfs," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 11, pp. 2494–2506, 2019.
- [5] —, "Dvfs-aware application classification to improve gpgpus energy efficiency," *Parallel Computing*, vol. 83, pp. 93–117, 2019.
- [6] —, "Gpu static modeling using ptx and deep structured learning," *IEEE Access*, vol. 7, pp. 159 150–159 161, 2019.
- [7] Y. Arafat, A. ElWazir, A. ElKanishy, Y. Aly, A. Elsayed, A.-H. Badawy, G. Chennupati, S. Eidenbenz, and N. Santhi, "Verified instruction-level energy consumption measurement for nvidia gpus," in *Proceedings of the 17th ACM International Conference on Computing Frontiers*, 2020, pp. 60–70.
- [8] Y. Wang, Q. Wang, and X. Chu, "Energy-efficient inference service of transformer-based deep learning models on gpus," in *2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics (Cybermatics)*. IEEE, 2020, pp. 323–331.
- [9] S. M. Nabavinejad, S. Reda, and M. Ebrahimi, "Batchsizer: Power-performance trade-off for dnn inference," in *Proceedings of the 26th Asia and South Pacific Design Automation Conference*, 2021, pp. 819–824.
- [10] A. Majumdar, L. Piga, I. Paul, J. L. Greathouse, W. Huang, and D. H. Albonese, "Dynamic gpgpu power management using adaptive model predictive control," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2017, pp. 613–624.
- [11] P. Zou, L. Ang, K. Barker, and R. Ge, "Indicator-directed dynamic power management for iterative workloads on gpu-accelerated systems," in *2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*. IEEE, 2020, pp. 559–568.
- [12] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [13] F. Tang, W. Gao, J. Zhan, C. Lan, X. Wen, L. Wang, C. Luo, Z. Cao, X. Xiong, Z. Jiang *et al.*, "Aibench training: balanced industry-standard ai training benchmarking," in *2021 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2021, pp. 24–35.
- [14] NVIDIA Corporation, "Nvml api reference manual," <https://docs.nvidia.com/deploy/nvml-api/index.html>, August 2021, accessed August 19, 2021.
- [15] —, "Cupti documentation," <https://docs.nvidia.com/cupti/Cupti/index.html>, August 2021, accessed August 19, 2021.
- [16] V. P. Dwivedi, C. K. Joshi, T. Laurent, Y. Bengio, and X. Bresson, "Benchmarking graph neural networks," *arXiv preprint arXiv:2003.00982*, 2020.
- [17] J. Schwarzrock, C. C. de Oliveira, M. Ritt, A. F. Lorenzon, and A. C. S. Beck, "A runtime and non-intrusive approach to optimize edp by tuning threads and cpu frequency for openmp applications," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 7, pp. 1713–1724, 2020.
- [18] pytorch.org, "Pytorch benchmarks," <https://github.com/pytorch/benchmark>, August 2021, accessed August 19, 2021.
- [19] Z. Wen, J. Shi, Q. Li, B. He, and J. Chen, "ThunderSVM: A fast SVM library on GPUs and CPUs," *Journal of Machine Learning Research*, vol. 19, pp. 797–801, 2018.
- [20] Z. Wen, J. Shi, B. He, Q. Li, and J. Chen, "ThunderGBM: Fast GBDTs and random forests on GPUs," *Journal of Machine Learning Research*, vol. 21, 2020.
- [21] H. Walser, *The golden section*. MAA, 2001.
- [22] K. Fan, B. Cosenza, and B. Juurlink, "Predictable gpus frequency scaling for energy and performance," in *Proceedings of the 48th International Conference on Parallel Processing*, 2019, pp. 1–10.
- [23] Z. Wang and M. O'Boyle, "Machine learning in compiler optimization," *Proceedings of the IEEE*, vol. 106, no. 11, pp. 1879–1901, 2018.
- [24] D. Stamoullis, E. Cai, D.-C. Juan, and D. Marculescu, "Hyperpower: Power-and memory-constrained hyper-parameter optimization for neural networks," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2018, pp. 19–24.
- [25] A. E. Brownlee, J. Adair, S. O. Haraldsson, and J. Jabbo, "Exploring the accuracy-energy trade-off in machine learning," in *2021 IEEE/ACM International Workshop on Genetic Improvement (GI)*. IEEE, 2021, pp. 11–18.