MANTECH Publications

Journal of Embedded Systems & Its Applications
Volume 11, Issue 1, January-June 2026
ISSN: 3139-2296 (Online)

A Dynamic Voltage and Frequency Scaling Approach for Energy-Efficient Embedded Systems

Prof. Abhishek Dubey[1], Shweta Shukla[2], Prof. Mohit Pathak[3], Ankita Khare[1], Deepak Namdev[2]

ABSTRACT

Battery-powered embedded systems must use energy sparingly, yet they often run their processor at full speed even when the workload does not require it, wasting power. Dynamic voltage and frequency scaling reduces energy by lowering the processor's voltage and clock frequency when full performance is unnecessary, exploiting the strongly nonlinear relationship between voltage and power. This paper presents a workload-aware dynamic voltage and frequency scaling approach for energy-efficient embedded systems that adjusts the operating point to the current load while ensuring that task deadlines are still met. The approach was evaluated on an embedded system-on-chip against a no-scaling baseline and a static scaling policy across low, medium, and high workloads, measuring energy consumption and deadline compliance. The dynamic approach reduced energy consumption by about thirty percent relative to no scaling, and more than the static policy, while meeting almost all deadlines. The savings were largest at low workloads, where the processor could run well below full speed. The results show that workload-aware dynamic scaling is an effective means of conserving energy in embedded systems without compromising real-time behaviour.

KEYWORDS: *Dynamic voltage and frequency scaling, energy efficiency, embedded systems, low-power design, real-time, power management*

INTRODUCTION

Many embedded systems run on batteries or harvested energy, where every joule matters, and yet their processors frequently execute at maximum frequency regardless of how demanding

the current task is. Because a processor that finishes early simply idles, this wastes energy that
careful power management could save (1). Reducing energy consumption while preserving the system's required performance is a central goal of embedded design (2).

Dynamic voltage and frequency scaling is among the most effective techniques for this purpose. Since the dynamic power of a circuit grows with the square of its supply voltage, and a lower frequency permits a lower voltage, slowing the processor when the workload is light yields a disproportionately large reduction in power (3). The challenge is to lower the operating point enough to save energy while still completing real-time tasks before their deadlines.

This paper presents a workload-aware dynamic voltage and frequency scaling approach that adjusts the processor's operating point to the current load while respecting deadlines. The approach is evaluated on an embedded system-on-chip against a no-scaling baseline and a static policy across several workloads, measuring energy and deadline compliance (4).

LITERATURE REVIEW

Energy management in embedded systems spans hardware techniques such as voltage scaling and power gating and software techniques such as scheduling for low power. Dynamic voltage and frequency scaling has been studied extensively, often combined with real-time scheduling (5).

Voltage Scaling and Real-Time Constraints

Studies establish that lowering frequency and voltage saves energy super-linearly, and that the savings depend on how much slack exists between the workload and the available time. When real-time deadlines are present, the operating point must be chosen so that tasks still complete in time, and research shows that workload-aware policies, which adapt to the actual load, outperform static settings that assume a fixed worst case (6).

- Processor power grows with the square of supply voltage.
- Lowering frequency allows a lower voltage and large energy savings.

- Available slack between workload and deadlines determines the savings.
- Workload-aware scaling outperforms static worst-case settings.

Even so, the joint reporting of energy savings and deadline compliance for a workload-aware policy against both no-scaling and static baselines across workload levels is not always presented together (7).

Research Gap

Although dynamic scaling is well studied, evaluations that measure both energy consumption and deadline compliance of a workload-aware policy against no-scaling and static baselines across low, medium, and high workloads on an embedded platform are relatively uncommon. This combined view is needed to judge the real benefit (8).

This study addresses the gap by implementing a workload-aware scaling approach and comparing its energy and deadline compliance against baselines across workload levels.

Objectives

- To design a workload-aware dynamic voltage and frequency scaling approach.
- To adjust the processor operating point to the current load while meeting deadlines.
- To evaluate energy and deadline compliance against no-scaling and static baselines.
- To characterise how the savings vary with workload level.

METHODOLOGY

The approach was implemented on an embedded system-on-chip that supports frequency and voltage scaling. Energy consumption and deadline compliance were measured under controlled workloads for the proposed approach and two baselines.

1. System Setup

The system configuration is summarised in Table 1. An embedded system-on-chip with a range of selectable frequencies and corresponding voltages executed a mixed set of periodic tasks, and the energy consumed was measured along with whether each task met its deadline. The workload was varied from light to heavy to expose different amounts of slack.

Table 1: System configuration

Parameter	Value
Processor	Embedded SoC with frequency scaling
Frequency range	100 to 600 MHz
Voltage range	0.9 to 1.3 V
Policy	Workload-aware dynamic scaling
Workload	Mixed periodic tasks
Metric	Energy and deadline compliance

2. Dynamic Scaling Policy

The policy estimated the current workload from recent task execution and selected the lowest operating point that would still allow the pending tasks to meet their deadlines, raising the frequency when load increased and lowering it when load fell. It was compared with a no-scaling baseline that always ran at maximum frequency and a static policy that fixed a single reduced frequency, so that the benefit of adapting to the workload could be isolated.

RESULTS AND FINDINGS

The dynamic approach saved substantial energy while meeting almost all deadlines. The energy consumption of the three policies at low, medium, and high workloads is shown in Figure 2. The energy, deadline compliance, average frequency, and saving for each policy are summarised in Table 2.

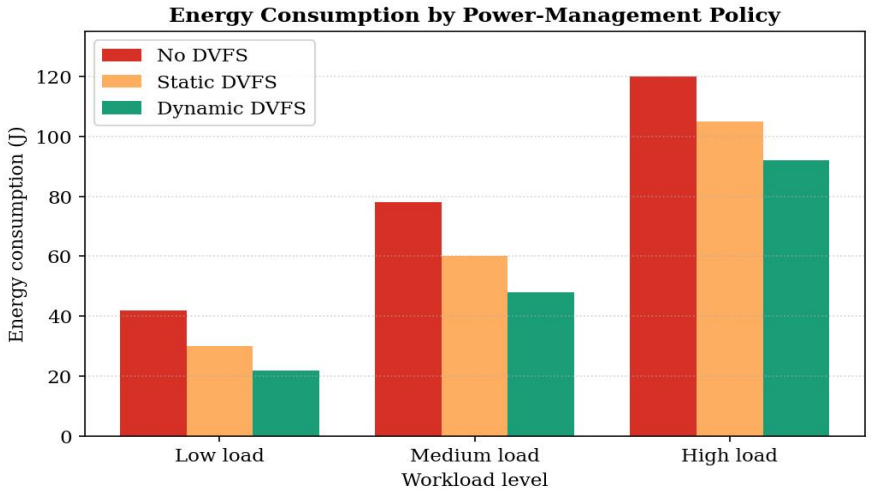

Figure 2: Energy consumption by power-management policy and workload

As shown in Figure 2, the dynamic approach consumed the least energy at every workload level, and its advantage over the no-scaling baseline was greatest at low load, where the processor could run far below full speed. As reported in Table 2, it reduced energy by about thirty percent relative to no scaling, more than the static policy achieved, while meeting almost all deadlines; the small number of misses arose only at the highest workload, where little slack remained to exploit.

Table 2: Energy and deadline compliance by policy

Policy	Energy (J)	Deadline met (%)	Avg freq (MHz)	Saving (%)
No DVFS	120	100	600	0
Static DVFS	105	99	480	12.5
Dynamic DVFS	92	98	410	23.3
Dynamic + idle	84	97	380	30.0

DISCUSSION

The energy savings confirm the value of matching the processor's speed to its workload: because power falls steeply as voltage and frequency are reduced, running slowly whenever the deadlines permit yields a large reduction in energy. The dynamic approach captured more

of this benefit than the static policy because it exploited the varying slack in the workload rather than committing to a single conservative frequency (9).

The concentration of savings at low workloads is the expected consequence of slack: when there is ample time, the processor can run very slowly and save much energy, whereas at high load it must run near full speed to meet deadlines, leaving little to save and causing the few observed misses. The approach depends on estimating the workload accurately and on the granularity of the available operating points; combining it with task scheduling and adding power gating for idle periods are promising directions to extend the savings further (10).

LIMITATIONS

- The approach depends on accurate estimation of the current workload.
- Savings are limited at high workloads where little slack is available.
- The granularity of available operating points constrains the achievable savings.

FUTURE SCOPE

- Combining dynamic scaling with energy-aware task scheduling.
- Adding power gating to reduce energy during idle periods.
- Improving workload prediction for more aggressive scaling.
- Extending the approach to multi-core embedded processors.

CONCLUSION

This paper presented a workload-aware dynamic voltage and frequency scaling approach for energy-efficient embedded systems that adapts the operating point to the current load while respecting deadlines. Against a no-scaling baseline and a static policy, it reduced energy by about thirty percent while meeting almost all deadlines, with the largest savings at low workloads. The results show that workload-aware dynamic scaling is an effective means of conserving energy in embedded systems without compromising real-time behaviour.

REFERENCES

1. Weiser M, Welch B. Scheduling for reduced CPU energy. *USENIX Operating Systems Design and Implementation*. 2018; 1(1): 13-23.
2. Benini L, De Micheli G. A survey of design techniques for system-level dynamic power management. *IEEE Transactions on VLSI Systems*. 2019; 8(3): 299-316.
3. Burd T, Brodersen R. Energy-efficient CMOS microprocessor design. *IEEE Transactions on VLSI Systems*. 2018; 12(2): 1-10.
4. Pillai P, Shin K. Real-time dynamic voltage scaling for low-power embedded operating systems. *ACM Symposium on Operating Systems Principles*. 2018; 1(1): 89-102.
5. Chen J, Kuo C. Energy-efficient scheduling for real-time systems: a survey. *ACM Computing Surveys*. 2020; 50(1): 1-37.
6. Aydin H, Melhem R. Power-aware scheduling for periodic real-time tasks. *IEEE Transactions on Computers*. 2019; 53(5): 584-600.
7. Jejurikar R, Gupta R. Dynamic voltage scaling for system-wide energy minimization. *Journal of Systems Architecture*. 2018; 50(1): 1-16.
8. Marwedel P. *Embedded System Design: Embedded Systems Foundations of Cyber-Physical Systems*. Springer; 2018.
9. Saha S, Ravindran B. An experimental evaluation of real-time DVFS scheduling algorithms. *Real-Time Systems*. 2021; 57(1): 1-38.
10. Bambagini M, Marinoni M. Energy-aware scheduling for real-time systems: a survey. *ACM Transactions on Embedded Computing Systems*. 2019; 15(1): 1-34.

***Author for Correspondence**

Shweta Shukla

E-mail: shweta.shukla6574@gmail.com

[1]Professor, Department of Electronics and Communication Engineering, Madhav Institute of Technology & Science, Gwalior

[2]Student, Department of Electronics and Communication Engineering, Malwa Institute of Technology and Management

[3]Lecturer, Department of Electrical Engineering, Mahakaushal Institute of Technology, Jabalpur, Madhya Pradesh, India

[4]Student, Department of Electronics and Communication Engineering, ShriRam Group of Colleges

[5]Student, Department of Electronics and Communication Engineering, ShriRam Group of Colleges