A Survey of Dynamic Voltage and Frequency Scaling for High-Performance Low-Power Systems

Shi Liu

Perimeter College, Georgia State University, Alpharetta, Georgia, 30005, the United State

sliu53@student.gsu.edu

Keywords: Dynamic Voltage Frequency Scaling, High Performance Computing, Low-Power, Data Center, Energy Efficiency

Abstract: Dynamic Voltage Frequency Scaling (DVFS) is a critical power management technique employed in modern computing systems to optimize energy efficiency while maintaining performance. This paper explores the principles and methodologies of DVFS, highlighting its role in balancing computational demands with power consumption. By adjusting voltage and frequency according to workload requirements, DVFS significantly reduces energy usage, which is crucial for battery-operated devices and large-scale data centers. This paper begins with a comprehensive overview of DVFS applications, followed by the responsiveness and predictive capabilities of DVFS. The paper also discusses the DVFS algorithm and implementation to real-time dynamically adjust settings in response to immediate workload changes. Finally, the paper also examines the trade-offs with experiments, demonstrating the effectiveness of DVFS in various applications.

1. Introduction

Energy utilization efficiency is extremely important in scientific and technological industry. The shortage of resources and the deterioration of the environment and other factors all promote the exploration of low power technology.

1.1 Low Power Technology in data centers

In the digital age, data centers which hosting everything from network services to cloud storage are the backbone of the Internet, and everyone from individuals to enterprises relies on it. As the demand for data processing and storage surges, there is increasing pressure on energy efficiency, especially power consumption, and low-power technology is an extremely important solution to meet these challenges.

One of the primary drivers for adopting low power technology in data centers is the environmental impact. Data centers consume a substantial amount of electricity, often sourced from non-renewable resources. According to the International Energy Agency (IEA), data centers account for about 1% of global electricity demand [1] and all this consumption of greenhouse gases such as carbon dioxide will exacerbate the greenhouse effect. Low power technologies, such as utilizing built-in server power management features, reducing energy losses from power distribution units (PDUs), and Reducing energy losses from uninterruptable power supply (UPS) systems, help mitigate the environmental impact [2].

Power costs is also a critical factor why the low power techniques should be applied, which represent a major portion of data center operating expenses. By adopting low power technologies, data centers can achieve significant cost savings. For example, since most servers are not operating at maximum capacity, economic losses can be greatly reduced by consolidating and improving server utilization, and the energy savings can also be considerable. A 2014 report by NRDC estimated average server utilization at 12 to 18 percent and noted that it had largely remained static from 2006 through 2012 [3]. Advancements in power management and virtualization technologies enable better utilization of hardware, allowing data centers to consolidate workloads and reduce the number of

  DOI: 10.25236/icetmr.2024.011

physical machines required. This consolidation results in further savings on power and cooling costs [4].

Operational efficiency is equally significant to drive the adoption of low power technologies. Data centers aim to maximize uptime and reliability while minimizing operational complexity and costs. Energy-efficient components often have better cooling and thermal management, which can lead to increased reliability and longer lifespans. For example, newer server designs that incorporate low power consumption technology tend to have lower failure rates. Efficient power usage allows data centers to scale operations more effectively. By using less energy per unit of computing power, data centers can handle more workloads without a proportional increase in energy consumption. This scalability is crucial for accommodating the growing demand for data services without a corresponding rise in energy requirements. Low power technologies often come with advanced monitoring and management tools that help data center operators optimize performance and preemptively address issues. This reduces the need for frequent maintenance and lowers operational costs, as shown in Fig. 1.

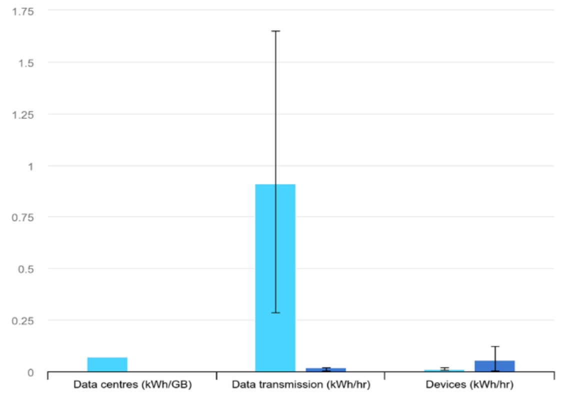

Fig. 1 Assumptions for energy intensity of data centres, data transmission networks and devices in 2019

1.2 Low Power Technology in Automotive Industry

Likewise, the automotive industry has always been at the forefront of technological innovation. In recent years, the design and production trend of vehicles has been moving toward low-power technology, and this shift is driven by market regulatory requirements and consumer demand.

The pressure of market regulation is still the pursuit of environmental protection. More than 20% of the EU's greenhouse gas emissions are caused by transport activities, and emissions from transport activities have been growing over the past few years compared to other energy-intensive sectors such as power generation and industry [5]. Therefore, the European Union's CO2 emission standards require new cars to emit no more than 95 grams of CO2 per kilometer by 2021. Similar regulations exist in the United States, China, and other major markets. These regulations have forced automakers to explore ways to improve the energy efficiency of vehicles, including low-power technologies. By reducing the power consumption of each part of the car, the overall energy consumption is reduced, and the emission is reduced. The most representative example is the use of low-power electronics in control systems, as illustrated in Fig. 2.

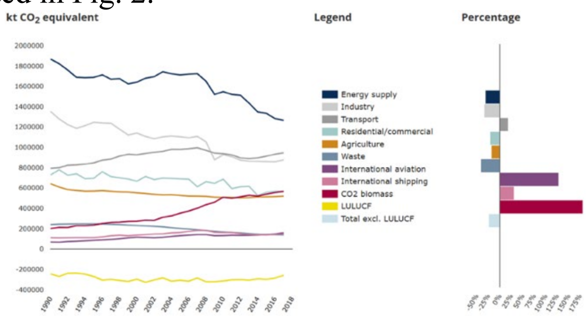

Fig. 2 Greenhouse gas emissions by aggregated sector

Consumer demand for fuel efficiency is also driving advances in low-power technologies. With wild fluctuations in fuel prices, among other factors, consumers tend to choose vehicles with better fuel economy.

In fact, thanks to the maintenance costs of fuel vehicles, the trolley industry has developed rapidly in recent years, and it has also become an important driving force for low-power technology. By 2030, global car ownership is expected to increase from 130 million to 2 billion [6]. Because the tram is completely dependent on battery power supply and the limited capacity of the battery, the power supply efficiency of the battery will directly affect the driving range of the car, which is one of the important factors that EVs consumers need to consider. Therefore, whether it is low-power processors for vehicle control systems or optimized thermal management systems all contribute to extending the range of Advances in EVs battery management systems are also minimizing energy loss during charge and discharge cycles.

1.3 Low Power Technology in Medical Industry

The medical industry has also been increasingly adopting low power technology across various devices and systems, driven by the need for energy efficiency, portability, patient safety, and cost-effectiveness.

Medical imaging such as CT scanners and MRI machines also benefit from low-power technology. "CT and MRI energy consumption is substantial. Considerable energy- and cost-saving potential is present during nonproductive idle and system-off modes, and this realization could decrease total cost of ownership while increasing energy efficiency." [7] However, "Awareness of energy efficiency has been rising in the industrial and residential sectors but only recently in the health care sector. "[7]. In recent years, there has been a push to develop more energy-efficient imaging technologies that do not compromise image quality. For example, new families of power ultrasonic transducers have been developed in recent years to meet the needs of various industrial fields [8].

Safety is particularly important in the medical industry. Failure of medical devices can lead to life-threatening consequences, and low-power technologies play an unparalleled role in improving the safety of medical devices. Reduced power consumption will greatly reduce the heat production of the device, thus extending the service life and reducing the risk of overheating of the device, which is the common case of device failure. At the same time, the monitoring and regulation of power consumption by low-power technology makes the key functions of the equipment always get priority. The advanced power management system enables the equipment to operate normally under low power consumption conditions in emergency situations.

2. The Reason for Using Dynamic Voltage and Frequency Scaling

The increasing complexity and performance requirements of modern computing systems have led to an increase in power consumption and heat output. These issues are particularly challenging in mobile devices, data centers, and embedded systems. If the circuit system still provides high power output when the hardware module in the case of low load, there will be a large amount of energy loss. High power will directly lead to system heat production and significantly reduce the span life of the device, as shown in Fig. 3.

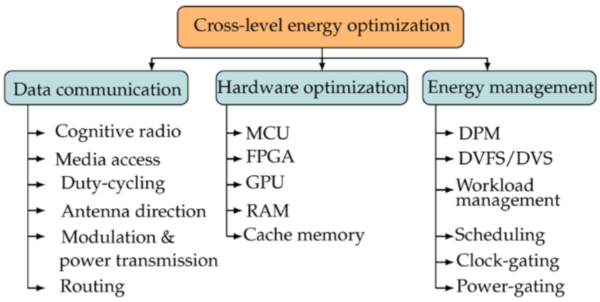

Fig. 3 Cross-level Energy Optimization Taxonomy

Dynamic Voltage and Frequency Scaling (DVFS) is a power management technique that addresses these challenges by adjusting the voltage and frequency of the processor in real time, depending on the workload. DVFS are essential for improving energy efficiency, managing heat output and extending the service life of equipment.

The most significant advantage of DVFS is its ability to enhance power efficiency. The workload can vary considerably when the processor executes the task. Generally, complex algorithms and large games are fully loaded, and simple tasks such as idle time or editing text are minimally loaded. However, if the output is still at high voltage and frequency during low-load tasks, it will result in a huge power loss. DVFS saves a lot of energy by dynamically adjusting the output during this period. In powered devices, which directly extends battery life.

Another advantage of DVSF is effective thermal management. The heat generated by a processor while running is proportional to the power consumed. Overheating can cause a range of problems including reduced performance and even permanent damage to hardware and equipment failure. DVFS reduces the heat generated by the processor by dynamically adjusting the output during periods of low demand to help keep the system temperature within safe limits.

While reducing power consumption, DVFS also plays a crucial role in optimizing performance. The rapid iteration of modern computers has led to continuous optimization of their performance, and DVFS helps them achieve a delicate balance with energy efficiency. In demanding tasks, DVFS allows the processor to operate at higher frequencies and voltages, providing the required computing power. Conversely, during low-demand tasks, the processor can reduce its performance to save energy. This ability to balance performance and power consumption is especially valuable in data centers, where servers may handle different workloads throughout the day. By using DVFS, data centers can reduce energy consumption during off-peak hours without compromising their ability to handle peak workloads when necessary. This not only reduces operating costs, but also contributes to the sustainability of large-scale computing operations. Adding DVFS to workload management can even help. The other research demonstrates that “including DVFS awareness in workload management provides substantial energy savings of up to 41.62% for scenarios under dynamic workload conditions”, and we can see in Fig. 4. [9]

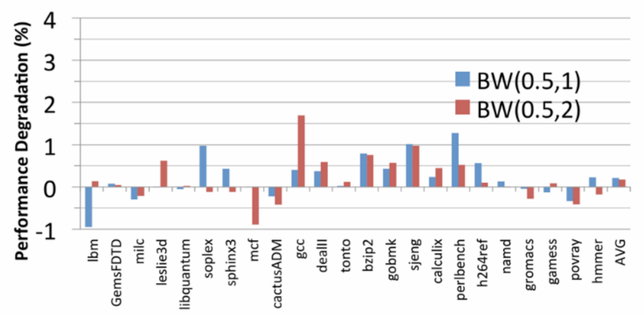

Fig. 4 Performance Impact of Memory DVFS, minimal performance degradation: 0.2% (avg), 1.7% (max)

The effects of DVFS also extend to the environment. With energy demand continuing to grow, significant energy consumption will significantly reduce the carbon footprint of equipment. Taking data centers as an example, while reducing operating costs, sustainability factors can be taken into account, in line with the global trend to reduce greenhouse gas emissions.

Today, DVFS has become a standard feature in a variety of computer environments. In mobile devices, the dynamic adjustment function of DVFS can extend battery life and span life, ensuring the user experience. In the data center, DVFS 'ability to optimize server energy efficiency can significantly save energy and maintenance costs. In embedded systems, where the device typically operates in a power-limited environment, DVFS enables the system to operate efficiently for a long time, even in a power-limited environment, as illustrated in Fig. 5.

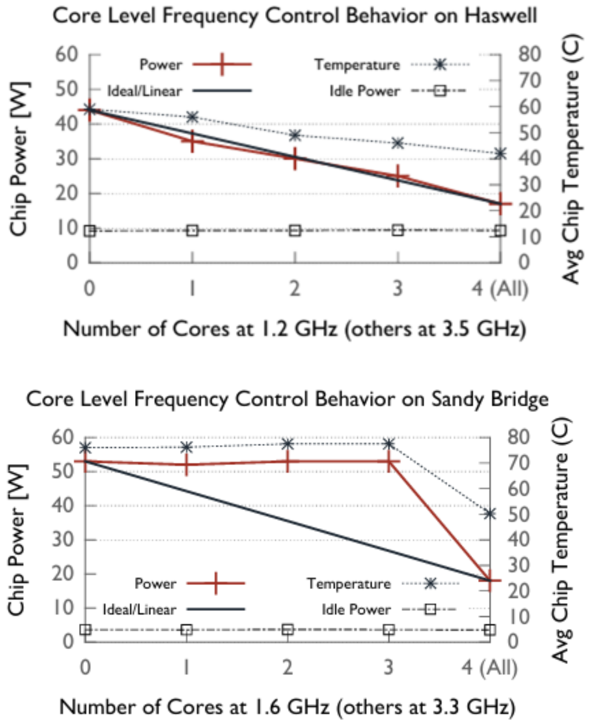

Fig. 5 Core level DVFS on Haswell architecture shows proportional/linear decrease in power when core frequencies are dropped on by one. On the other hand, since Sandy Bridge do not have per core voltage regulators, all core frequencies need to be dropped together to for a reduction in power and temperature, as shown in Fig. 6. [10]

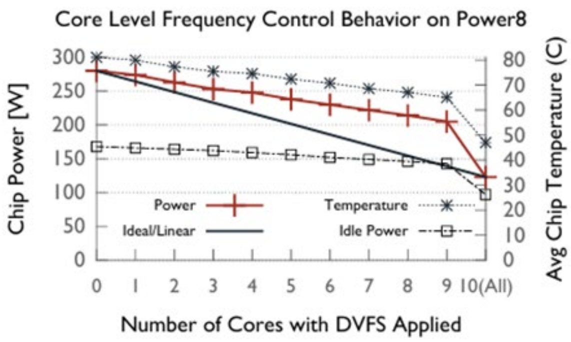

Fig. 6. Core level frequency scaling on IBM POWER8 pro-cessor.[10]

3. The Principle of DVFS

The dynamic power consumption of CMOS circuit can be expressed by the following formula:

$$P_{dynamic} = ACV^2 f \tag{1}$$

Where C represents the capacity of the load capacitance, V is the operating voltage, A is the average turn-over rate of the circuit at the current frequency, and f is the operating frequency. It can be seen from the formula that C, V, A and f determine the power consumption of the entire CMOS circuit, and DVFS technology is mainly to adjust the power consumption of the system by changing the value of frequency f and voltage V.

Voltage scaling is a critical aspect of DVFS. By lowering the supply voltage V, DVFS can reduce dynamic power consumption significantly. However, decreasing voltage also impacts the circuit's delay t_d, which affects the maximum operating frequency of the processor. The relationship between delay and voltage can be expressed as

$$t_d = \frac{V}{(V - V_{th})^2} \quad (2)$$

where V_{th} is the threshold voltage of the transistor. As the supply voltage nears the threshold voltage, the delay increases sharply, limiting the maximum operating frequency. This trade-off is a crucial consideration in DVFS, as it requires balancing power savings with the need to maintain adequate performance levels. Voltage scaling is often combined with frequency scaling in DVFS implementations to maximize power efficiency without compromising performance.

Under the current chip implementation process, higher frequencies require higher supply voltages. Because the higher frequency means that the dynamic power consumption increases, and under the premise that the overall capacitor resistance value of the chip is unchanged, the higher supply voltage can provide higher power supply and then the clock frequency. When the power needs to be reduced, reduce the clock frequency first, and then reduce the supply voltage. Blindly reducing frequency and voltage does not necessarily reduce power consumption, because the system processing task time will increase at low frequencies, and in some cases, it will lead to increased power consumption.

The core of DVFS is the dynamic adjustment of the task load to meet the minimum power requirements of performance. The specific strategy is to first find out the high-power components such as the CPU, calculate the load, and increase the frequency to increase the voltage, and vice versa. By counting the working hours of the module, different thresholds are set to be used as the conditions for DVFS conversion.

4. DVFS Implementation

Implementing DVFS involves adjusting the processor's voltage and frequency dynamically, depending on the workload, to optimize power usage while maintaining adequate performance, which requires specific hardware, software frameworks, and powerful algorithms.

Implementing DVFS requires specific hardware support to allow the dynamic adjustment of voltage and frequency.

Voltage regulators are critical in implementing DVFS, as they control the supply voltage to the processor. Modern processors are typically equipped with integrated voltage regulators (IVRs) that enable rapid and fine-grained control of voltage levels. These IVRs must support a wide range of voltages and be capable of switching between them with minimal latency to ensure that performance is not unduly impacted during transitions.

PLLs are used to generate and control the clock signal frequency in processors as shown in Fig. 7. To implement DVFS, the PLL must be capable of adjusting the clock frequency dynamically. This capability is essential because changing the frequency allows the system to scale performance up or down as needed. The design of PLLs must consider the trade-off between stability and response time, as the frequency adjustments must be stable yet quick enough to respond to real-time changes in workload, as shown in Fig. 7.

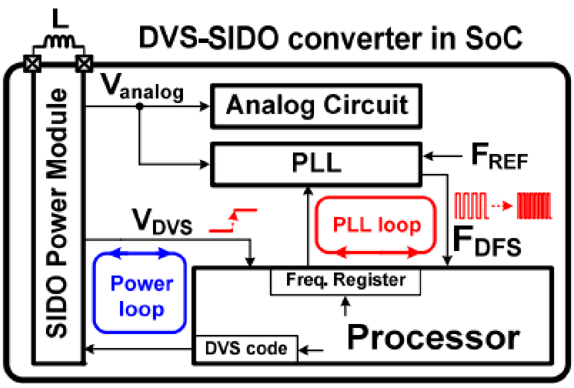

Fig. 7. SoC system with the DVFS implementation with PLL and voltage regulator.[11]

DVFS is often used in conjunction with thermal management strategies to prevent overheating. Thermal sensors embedded within the processor monitor the temperature and provide feedback to the

DVFS controller. If the temperature exceeds a predefined threshold, the DVFS system can lower the voltage and frequency to reduce heat generation, thereby protecting the processor from thermal damage.

Power gating circuits are used to reduce leakage power in inactive portions of the processor. In a DVFS-enabled system, power gating can be combined with voltage scaling to further reduce power consumption when certain cores or functional units are not in use. This approach is particularly useful in multi-core processors, where different cores can operate at different voltage and frequency levels depending on the workload.

While hardware components provide the foundation for DVFS, software frameworks are necessary to manage the dynamic adjustment of voltage and frequency.

The operating system (OS) plays a crucial role in DVFS implementation. The OS is responsible for monitoring the system's workload and making decisions about when to scale the voltage and frequency. Many modern operating systems, including Linux and Windows, have built-in DVFS support, often referred to as "CPU frequency scaling" or "power management" features. The OS uses a governor, which is an algorithm that decides the appropriate P-state (a combination of voltage and frequency) based on the current workload. Common governors include "ondemand," which increases the frequency rapidly when demand increases, and "conservative," which scales the frequency more gradually.

Application Programming Interfaces (APIs) provided by the OS or hardware vendors allow software developers to interact with the DVFS mechanism. These APIs can be used to query the current P-state, set specific P-states, or customize the behavior of the DVFS governor. For example, the Linux kernel provides the subsystem, which offers a range of functions for controlling CPU frequency and voltage settings. Developers can write custom governors or use existing ones to optimize the performance and power consumption of their applications.

In embedded systems and real-time applications, where timing constraints are critical, DVFS must be implemented in a way that does not compromise the system's ability to meet deadlines. Real-time operating systems (RTOS) like FreeRTOS and VxWorks provide mechanisms to integrate DVFS while ensuring that the timing requirements are met. This integration is achieved by carefully balancing the trade-offs between power savings and performance, often through the use of deadline-aware DVFS algorithms.

In virtualized environments, DVFS implementation becomes more complex because multiple virtual machines (VMs) may share the same physical resources. The hypervisor, which manages these resources, must ensure that DVFS adjustments do not negatively impact the performance of any VM. This challenge is addressed through techniques such as coordinated DVFS, where the hypervisor monitors the workload of all VMs and adjusts the voltage and frequency accordingly. Additionally, power-aware scheduling can be used to allocate VMs to physical cores based on their power consumption profiles, optimizing overall energy efficiency.

Effective DVFS implementation requires robust algorithms as well to decide when and how to adjust the voltage and frequency.

The first step in any DVFS algorithm is to accurately characterize the workload. Workload characterization involves analyzing the computational demands of the running tasks, which can be done using performance counters, power models, or direct measurements of CPU utilization. Based on this characterization, the DVFS algorithm can predict the future workload and adjust the P-state accordingly. For example, if the workload is expected to remain low, the algorithm might lower the frequency and voltage to save power.

Predictive algorithms use historical data and workload patterns to forecast future CPU demand. These algorithms typically employ machine learning techniques or statistical models to make predictions. For instance, a predictive DVFS algorithm might use regression analysis to predict the required CPU frequency based on recent workload trends. The advantage of predictive algorithms is that they can preemptively adjust the voltage and frequency before the workload changes, minimizing response time and improving energy efficiency.

Reactive algorithms adjust the voltage and frequency in response to immediate changes in workload.

These algorithms typically use feedback control systems, where the current performance metrics (e.g. CPU utilization, power consumption) are compared against predefined thresholds. If the metrics exceed the thresholds, the DVFS system increases the frequency and voltage; if they fall below, the system decreases them. Reactive algorithms are simpler to implement than predictive ones but may suffer from lag, as adjustments are made only after the workload changes.

Hybrid algorithms combine predictive and reactive approaches to leverage the strengths of both. For example, a hybrid DVFS algorithm might use predictive modeling to set a baseline frequency and then apply reactive adjustments based on real-time feedback. This combination can provide a good balance between responsiveness and stability, making it suitable for a wide range of applications, from mobile devices to data centers.

A decrease in the supply voltage results in an increase in the gate propagation delay (td) :

$$t_d \approx \frac{CV_{dd}}{V_{dd} - V_t} \alpha$$

where V_t is the threshold voltage and α is the velocity saturation index (≈1 in the nanometers). In order to ensure the normal operation of the synchronous system, the frequency must usually change with the voltage. In multi-processor architectures, the performance overhead of frequency and voltage scaling can be mitigated by taking advantage of workload variations across processor arrays. The processor can operate at higher voltages during high workloads and lower voltages during low workloads to minimize energy consumption. The effectiveness of a DVFS circuit can be estimated by multiplying the transmission delay by the processor core voltage, known as the energy delay product (EDP), as shown in Fig. 8.

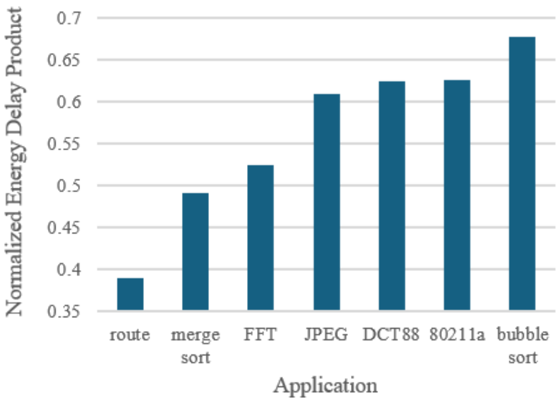

Fig. 8. Relative energy delay product with DVFS compared to a non-DVFS design for various applications. [12]

Figure 8 analyses the effectiveness of the DVFS architecture (with the same voltage configuration Settings) in different applications[12]. The processor uses 65nm CMOS technology and has a total of 48 power gates per processor. The simulation was carried out at a voltage supply with a V_{ddhigh} of 1.3V, a V_{ddlow} of 0.8V and a step transition of 0.1V. The behavior of each processor workload, represented by FIFO utilization and stall duration, directly affects the performance overhead and EDP. In merge sort applications, where the workload per processor is constant (constant large or small), DVFS results in a lower performance overhead and therefore a lower EDP. In contrast, in bubbling sort applications, where the workload per processor fluctuates between large and small, DVFS can lead to a greater performance overhead, leading to a larger EDP. The average of the results in Figure 8 is also 0.56.

5. Conclusion

Dynamic Voltage and Frequency Scaling (DVFS) is a powerful technique for optimizing power consumption and performance in modern computing systems. By dynamically adjusting the processor's voltage and frequency based on real-time workload demands, DVFS provides a flexible approach to managing energy efficiency, thermal output, and performance. The principles of voltage scaling and frequency scaling form the foundation of DVFS, enabling its effective application across various computing environments, from mobile devices to data centers and embedded systems. As the demand for energy-efficient computing continues to grow, the role of DVFS in managing power consumption and thermal output will become increasingly important.

References

[1] IEA (2020), Assumptions for energy intensity of data centres, data transmission networks and devices in 2019, IEA, Paris https://www.iea.org/data-and-statistics/charts/assumptions-for-energy-intensity-of-data-centres-data-transmission-networks-and-devices-in-2019, Licence: CC BY 4.0

[2] U.S. Environmental Protection Agency (EPA). IT Solutions and Power Infrastructure. (2021) https://www.energystar.gov/products/data_center_equipment/it-power-infrastructure

[3] NRDC. 2 Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers. (2024) https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf

[4] Bertoldi, Paolo. "A market transformation programme for improving energy efficiency in data centres." *ACEEE Summer Study on Energy Efficiency in Buildings*. 2014.

[5] European Environment Agency. Greenhouse gas emissions by aggregated sector. (2021) https://www.eea.europa.eu/en

[6] Martins, Lívia Salles, et al. "Electric car battery: An overview on global demand, recycling and future approaches towards sustainability." *Journal of environmental management* 295 (2021): 113091.

[7] Heye, Tobias, et al. "The energy consumption of radiology: energy-and cost-saving opportunities for CT and MRI operation." Radiology 295.3 (2020): 593-605.

[8] Yao, Ye, Yue Pan, and Shiqing Liu. "Power ultrasound and its applications: A state-of-the-art review." Ultrasonics sonochemistry 62 (2020): 104722.

[9] Arroba, Patricia, et al. "Dynamic voltage and frequency scaling-aware dynamic consolidation of virtual machines for energy efficient cloud data centers." *Concurrency and Computation: Practice and Experience* 29.10 (2017): e4067.

[10] Acun, Bilge, Kavitha Chandrasekar, and Laxmikant V. Kale. "Fine-grained energy efficiency using per-core dvfs with an adaptive runtime system." 2019 Tenth International Green and Sustainable Computing Conference (IGSC). IEEE, 2019.

[11] Lee, Yu-Huei, et al. "On-the-fly dynamic voltage scaling (DVS) in 65nm energy-efficient power management with frequency-based control (FBC) for SoC system." *IEEE Asian Solid-State Circuits Conference 2011*. IEEE, 2011.

[12] Cheng, Wayne H., and Bevan M. Baas. "Dynamic voltage and frequency scaling circuits with two supply voltages." *2008 IEEE International Symposium on Circuits and Systems*. IEEE, 2008.