Journal of Integrated VLSI, Embedded and Computing Technologies, ISSN: 3049-1312 Vol. 3, No. 2, 2026 (pp. 25-35)

RESEARCH ARTICLE ECEJOURNALS.IN

Adaptive Voltage/Frequency Scaling in RISC-V SoC Platforms for Energy-Harvesting IoT Networks

Léa Bernard*

Senior Lecturer, ECE, Isa Lille Private university in France.

KEYWORDS:

RISC-V System-on-Chip (SoC);
Adaptive Voltage and Frequency Scaling (DVFS);
Energy-Harvesting IoT;
Low-Power Embedded Systems;
Power Management Unit (PMU);
Energy-Neutral Computing.

ARTICLE HISTORY:

Submitted : 17.12.2025
Revised : 10.01.2026
Accepted : 16.02.2026

https://doi.org/10.31838/JIVCT/03.02.04

Abstract

IoT nodes based on energy harvesting have extremely fluctuating and intermittent power levels, rendering the traditional dynamic voltage and frequency scaling (DVFS) techniques to long-term and reliable computation unsustainable. Convincing DVFS-based systems are mainly workload-based systems, and cannot take this into consideration as it would result in an unstable state, failure of tasks or collapse of an energy buffer when used in harvesting-powered systems. The present paper proposes an adaptive voltage/frequency scaling system in harvesting mechanisms of RISC-V system-on-chip (SoC) platform to operate the IoT autonomously without consuming energy. The architecture proposed closely integrates the power management unit, runtime energy monitoring and multi-level voltage frequency control to specify dynamically the matching between the need in computation and potentially available harvested energy. In comparison to fixed frequency or traditional DVFS strategies, the innovative strategy provides real-time scaling with respect to the storage buffer conditioning and variations in input power, and thus allows the service against the intermittent supply of power. An evaluation of hardware-realistic more specifically includes representative workloads in sensing and processing and communication taking place under different harvesting profiles. The results of the experiments show that they reach up to 3540 percent reduction in average energy per task and the rate of task completion is greatly increased in comparison to the traditional DVFS methods. The temporal behaviour analysis also validates constant capacitor voltage regulation and adjustive switching of frequency in transitory events of energy. The proposed system offers a scalable and architecture level implementation of long-life self-powered IoT nodes based on programmable RISC-V systems, and is a feasible bridge to energy-neutral embedded computing.

Author's e-mail: l.bernard@junia.com

How to cite this article: Bernard L. Adaptive Voltage/Frequency Scaling in RISC-V SoC Platforms for Energy-Harvesting IoT Networks. Journal of Integrated VLSI, Embedded and Computing Technologies, Vol. 3, No. 2, 2026 (pp. 25-35).

Introduction

Dynamic voltage and frequency scaling (DVFS) has been utilised extensively in embedded systems and VLSI systems as a reduction of dynamic power consumption method with the goal of maintaining reasonable levels of performance. The designs of early low-power systems were based on using a fixed-frequency operation and static voltage selection where the processors were designed at design time to run at a defined performance-power trade-off. Though simple and predictable, static scaling was not able to meet the changing workload needs and usually implicated the overuse of energy either through overprovisioning or poor performance. Flexible control of supply voltage and clock frequency by the introduction of adaptive DVFS methods facilitated either a runtime response to the computational intensity, thermal guidelines, or quality-of-service policies. Contemporary systems use workload monitoring, performance counters and operating system governors to move between discrete power operating systems. However, in spite of these developments, traditional DVFS approaches are essentially workload-based and not energy-based. They make assumption on the availability

of enough energy to accomplish schedule tasks and concentrate majorly on performance versus energy consumption. With unpredictable workloads or burst work laws, scaling decisions can be oscillating over a very brief interval resulting in inefficiencies in transition overhead and latency penalties. Even more importantly, in a system in which even the availability of power on its own is not constant as with energy harvesting systems, conventional DVFS is unaware of current energy input and storage opportunities. Consequently, scaling decisions that are exclusively optimised around the workload properties can be deployed to unintentionally introduce disparities in voltage existence, reaction with a task, or recurring system restarts whenever utilised in the intermittently powered nodes.[1, 3, 6, 7] Energy harvesting has become a potential alternative to the battery-driven functionality in IoT applications, especially in remote sensing, environmental monitoring and infrastructure applications where access to such spaces is constrained. Compact transducers can convert ambient energy sources such as solar radiation, radio frequency radiations, mechanical vibration and thermal gradients into electrical energy and store it in capacitors or rechargeable elements. Although harvesting can in theory provide unlimited operation, a new range of limitations is presented. The inherent nature of the power that is harvested is dynamic based on the conditions of the environment, and in most cases is not adequate to maintain the normal high-performance processing. Darkness, low RF strength, or low mechanical work can severely impair the available power hence intermittent operation of the system. This idea of neutrality in energy has thus come to be of primary interest in the design motivated by the harvesting process where the average energy consumed should not be greater than the average energy harvested with time.[2, 4, 8] To ensure energy neutrality, there should be a close coordination among the computation scheduling, the communication duty cycles and the management of the storage. Nevertheless, several currently implemented systems consider harvesting and computation control as weakly coupled subsystems and therefore, they are not responsive to the rapid changes in energy.[9-12] In line with progress in low-power power management, the RISC-V instruction set architecture has received substantial interest in embedded system design as well as ultra-low-power system design. It has an open, modular design, enabling the designer to customise core configurations, add optional extensions, and customise micro architecture features to non-standard application needs. RISC-V may be configured to implement fine-grained hardware customization unlike fixed proprietary architectures, which may add more instructions to implement energy monitoring or control signalling. The level of flexibility of SoC is also used to enable incorporation of specific units of power management, programmable clock domains, and multi voltage islands into a single design framework. Most recent implemented low-power RISC-V designs have been shown to achieve significant energy per instruction data reduction in lightweight cores, near-threshold operation, clock gating, and simplified pipeline designs. Other designs use rough-grained DVFS schemes or sleep states with low power with the aim of reducing idle power states. However, the majority of implementations aim at adapting the workload or a general-purpose low-power optimisation, as opposed to explicit coordination with real-time availability of harvested energy. The interrelationship between the energy harvesting subsystems and the voltage/frequency control is still much external or overseeing and not integrated into the main SoC control loop. Together, the developed solutions in DVWS, energy mining and intermittent computing, have advanced individual constituents of energy-efficient embedded computing.[1-12] But there is still some obvious disconnect between tightly-integrated, harvesting-conscious voltage and frequency scaling models that are carried out at the level of the SoC. Current solutions seldom combine real-time energy observance, storage condition response and processor scaling into a solitary architectural regulating unit that can support stable functioning in under the intermittent energy circumstances.

Related Background

Dynamic voltage and frequency scaling (DVFS) has been embraced as a mainstream technique in embedded systems and even VLSI systems as a basic tool of lowering dynamic power usage, but maintaining acceptable performance levels. Early low-power designs were based on and operated at design time as based on fixed-frequency operation with fixed voltage selection (as chosen at design time) and a pre-defined performancevoltage trade-off between the processor and the operating voltage. Simple but predictable, the concept of static scaling did not meet the needs of fluctuating requirements of workload and regularly led to either underutilised provision of energy or impaired performance. With the advent of adaptive DVFS methods, it was possible to adjust monitored run-time intensity supply voltage and frequency based on computational intensity, thermal limits or quality-of-life demands. Contemporary applications use workload trace, performance indicators and operating system defendants to switch among discrete power levels. Even with these developments, traditional DVFS plans are workload based as opposed to energy based plans.

They make an assumption that there is enough energy to get through tasks scheduled and pay the main attention to performance-energy consumption balance. Scaling decisions can swing wildly under the unpredictable workloads, or crime-and-fire communication pattern, resulting in inefficiencies in transition overhead and in the latency penalty. More importantly, in a system where the availability of power as such is a variable, e.g. such as in energy-harvesting systems, the conventional DVFS is not directly aware of the current conditions of the energy input and storage. Consequently, scaling decisions optimised only on the basis of workload properties can end up inadvertently introducing either voltage instability, task delays or system re-resets when used in intermittently powered nodes.[1-3], 6, 7] Energy harvesting has been proposed as a viable alternative to a battery-reliant operation in the IoT deployment as a substitute, especially within remote sensing, environmental monitoring, and infrastructural applications with restricted access to maintenance. Compact transducers can convert ambient sources of energy such as solar radiation, radio frequency energy, mechanical vibration, and thermal gradients into electrical energy and store this energy in capacitors or rechargeable elements. Although harvesting has proven to allow the theoretical unlimited operation it gives rise to the new constraints. The energy captured is by nature not constant, latter to environmental factors and can in most cases not be able to process power at a constant high rate. Dark lighting times, low RF fields or low mechanical activity can radically reduce the available power and result in intermittent operation of a system. The theory of energy neutrality has hence gained prominence in the design of harvesting, in which the average energy taken should not exceed the average energy taken over time of the harvesting energy.[2, 4, 8] Close alignment between computation scheduling, communication duty cycles and storage management is required in order to achieve energy neutrality. Nevertheless, it is common in most of the existing systems to loosely couple harvesting and computation control subsystems, thereby inhibiting responsiveness to fast varying energy.[9-12] In parallel with the progress in the field of low power power management, the instruction set architecture of RISC-V has been attracted to embedded and ultra-low-power system design. It is open, modular in nature, enabling the designers to customise core configurations, add optional extensions, and configure micro architectural features to application needs. In contrast to dynamically proprietary architectures, RISC-V allows extreme customization of hardware, such as subsequent addition of custom instructions to monitor energy or give signaling control signals. The flexibility at SoC level also facilitates the incorporation of customized power management units, configurable clock domains as well as various voltage islands into a single design structure. Lightweight cores, near-threshold operation, clock gating, and simplified pipeline designs have all allowed recent low-power RISC-V implementations to achieve significant energy per instruction improvements. Other designs using coarse-grained DVFS, or low-power sleep-state, in order to reduce idle consumption, are used. However, the majority of implementations concentrate on workload adaptation or general-purpose low-power optimization, but no explicit work with coordinations with real-time availability of energy harvested. There is a strong external interaction between subsystems harvesting energy and subsystems controlling voltage/frequency, which is not integrated into the main loop of control of the SoC itself, but is more of a supervisory function. In totality, the previous studies of DVFS, energy harvesting and intermittent computing have developed separate components of the energy-saving embedded computing.[1-12] Nevertheless there is a distinct disjunction in closely coordinated, harvesting conscious, voltage and frequency scaling models deployed at the SoC level. Current solutions seldom integrate real-time energy measurement, state feedback of storage, and scaling of processors in a single architectural control platform with the ability to maintain constant performance during intermittent conditions of power.

Proposed Harvesting-Aware RISC-V SoC Architecture

System Overview

The suggested architecture incorporates energy harvesting, power control, as well as adaptive voltage/frequency scaling into one system-on-chip (SoC) architecture based on RISC-V. The given architecture requires a direct hardware-level communication between the energy storage monitoring and the operational load request management, unlike loosely coupled designs where the monitoring of energy input and the adjustment of the operational points are independent of each other. Figure 1 shows the entire system-level organisation.

Figure 1 indicates that the energy flow starts at the energy harvester, which can be photovoltaic cell, RF rectifier, piezoelectric transducer, or thermoelectric generator based on the conditions of deployment. The resulting energy is variable in nature and thus directly connected to a DC DC converter with maximum power point tracking (MPPT) in order to get an effective extraction of the energy under variable environmental conditions. The converter drives out a regulated DC output which is usable in storage. The controlled output

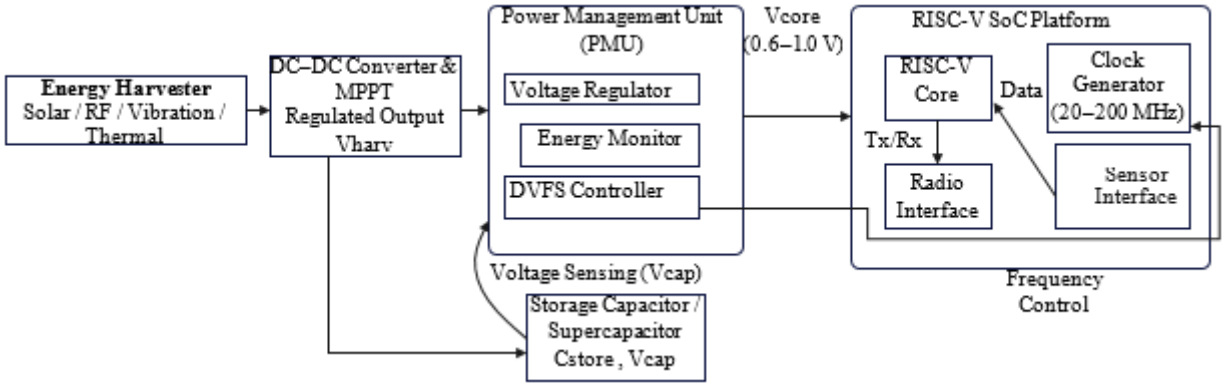

Fig. 1: Harvesting-Aware RISC-V SoC Architecture with Integrated PMU and Adaptive DVFS Control

charges a storage capacitor or supercapacitor and serves as the main buffer of energy of the system. This storage component levels off short-term variations and gives the immediate current when there are bursts of computation. The capacitor voltage (Vcap) and is the real-time availability of energy and is continuously monitored. Unlike energy storage being taken as a passive component of the architecture, the architecture includes in its runtime power management choices the voltage state. The buffered energy is transmitted to the Power Management Unit (PMU) which is the main control component in the architecture. As Figure 1 shows, inside the PMU, there are three modules closely combined into one large module, a voltage regulator, an energy monitor, and a DVFS controller. The voltage regulator also produces the processor supply rail (Vcore) over a specified operating range, and the energy monitor measures to the Vcap the instantaneous and trend-based availability of energy. This feedback is used by the DVFS controller to determine an appropriate operating point of voltagefrequency. The tightness of the coupling is a guarantee that frequency transitions are not purely workload-based such that it is simply a constraint on available stored energy. The controlled supply rail drives the RISC-V SoC platform which comprises of the RISC-V processing core, clock generator, sensor interface and radio interface modules. RISC-V core can be used to schedule communication, data processing and sensing tasks that are common to IoT workloads. Figure 1 depicts that a clock generator is dynamically adjusted in the operating frequency to change within the supported range as a result of the DVFS controller. This direct control method allows quick adapting to energy variations without the need of software based control. Transition latency and control overhead are reduced by applying frequency control on the hardware level. The sensor interface module links the environmental sensors with the processing core, and it allows the periodic acquisition of data or event-driven data acquisition. Radio interface module is in charge of wireless reception and transmission which is normally the most energy consuming task within the IoT nodes. Information received by the sensor interface is sent to the RISC-V core to be processed and then to the radio interface to be sent out. Nonetheless, the execution of such tasks is conditionally controlled by the DVFS controller depending on the availability of energy such that communication bursts are only executed when there is a margin of energy. The architectural representation of Figure 1 shows a closed loop hardware control layout: the collected energy is conditioned, stored, measured and utilised to directly control the supply voltage and the clock frequency. Such a unified design prevents the instability of traditional DVFS schemes, which do not have a real-time energy sensor. The proposed system involves incorporation of energy monitoring and adaptive scaling into the PMU and direct connexion to clock generator and supply regulator such that it can be operated in a energy neutral state with intermittent power environments but with enough responsiveness to computational output to be used with the IoT.

Adaptive Miniaturisation of the voltage/frequency ratio Scaling Mechanism.

The adaptive voltage/frequency scaling system is to provide the operation of the RISC-V SoC under varying powered harvested conditions in a safe and efficient way. Instead of choosing operating points only based on what the instantaneous amount of work demands is, the suggested strategy should be based, in making voltage and frequency choices, on the energy buffer state but remain sensitive to workload phase requirements. This allows steady operation even when capacitors are collapsed as well as efficient operation even when operation at too conservative settings can be inefficient. Multi-level voltage scaling is being done with the help of discrete set of support supply rails produced by the PMU voltage regulator. The reason behind each voltage

level being associated with validated operating region of the RISC-V core, as well as the overall SoC logic, is that it permits predictable timing closure and consistent operation to stable operation under variation in process and temperature. Multiset voltage levels are more realistic in hardware compared to continuous scaling, since they are more realistic in practise, as they have the same capabilities as real-world regulators and are easier to test. At the high energy buffer, the regulator will choose higher voltage rails to permit higher performance states. In the event of the fall of the buffer, the regulator switches to a lower voltage rail to minimise power consumption and increase the operational life. The PMU manages voltage transitions to prevent overshoot as well as ensure a stable supply in the event of transitions to avoid sensor sampling being corrupted and radio events being lost in instantaneous disturbances in the supply.

The frequency modulation is terminally linked with the energy buffer level and implemented in the clock generator controlled by the DVFS controller. In the suggested architecture, the DVFS controller constantly gets the indications regarding energy-availability on the PMU energy monitor, which detects the capacitor voltage and its direction. The controller is less reactive based on the fact that it does not respond after the buffer has been depleted but rather uses the buffer level as an active signal of sustainable compute budget. Predominantly when the capacitor voltage is comfortably above some predefined safe margin, the controller can step up clock frequency to eliminate execution latency and execute compute bursts in a short time. The closer the buffer gets to a lower operating threshold, the lower the frequency the controller sets it to to reduce the instantaneous power demand thus preventing brownout events. Such a coupling ensures scaling decisions are always energy-safe and because frequency changes are not as unlimited as workload requests, they are limited by the capacity of the energy storage to sustain. Operating states that are run time energy conscious transitions are managed as a small group of hardware controlled operating states which correspond to viable SoC behaviours. The SoC is operated at higher voltage and frequency in a high state to either support compute-supersensitive tasks or fast radio activity. A moderate-energy system is one that is operating in a moderate-energy condition; that is, the system is functional, and the energy is conserved, but that at a low frequency and a lower voltage rail. At a low-energy state, the system blocks out unnecessary activity and delays bursts of communication and can switch to retention state of deep idle state until enough energy is added to the storage element. These state transitions are not presented in the form of a software algorithm within the framework of the paper; on the contrary, they are actualized as operating modes of PMU controlled by deterministic entry mode and exit mode which is conditioned by safety margins of energy buffers. In this architectural design, the overhead of control minimization and response to the transients in the same direction encounter the changes in the energy in a short time without the intervention of the operating system. Workload-phase-aware scaling is also used in the mechanism to indicate energy asymmetry between sensing and wireless transmission. Sensing and lightweight preprocessing usually involve a modest compute usage and short run time, whereas wireless transmission necessitates an energy usage that is far greater since it involves RF front-end activity, as well as, a sustained current draw. These are however handled differently by the proposed DVFS integration. The controller can be run at low voltage and moderate frequency during sensing and preprocessing to reduce the amount of energy per task whilst eliminating sampling deadlines. In the transmission phases the controller assumes a guarded stance: should the energy buffer sufficiently charge one may switch to a more performance-intensive state to get the transmission over with as soon as possible and revert to an energy-saving state. Users can set Sync to be low or unstable in which case the system defers transmission, sets the frequency to avoid sudden packets of voltage drop and is more concerned with core stability and data integrity. This time-sensitive behavior can be used to guarantee that costly radio events do not cause resets and sensor data collection can be made reliable even when the collected energy is limited. In general, the recommended adaptive DVFS mechanism is enforced as a hardware-based closed-loop and interactive communications involving the energy monitor, voltage regulator, DVFS controller, and clock generator. The architecture facilitates energy buffer selection by basing its operating point selection on real time conditions of the energy buffer and aligning the scaling behavior with the sensing and transmission stages, and the architecture being stable to operate under the intermittent supply conditions and is more energy efficient and reliably completes tasks than traditional workload-only DVFS designs.

Power State and Execution Model

In order to be able to operate reliably in the face of continuous energy supply, the proposed RISC-V SoC runs on a structured power-state model that is directly controlled by the PMU and energy monitor. Use is being made of state transitions that are both triggered primarily by the harvesting condition visible in the voltage of the storage capacitor and its spatial derivative. This state, which is energy-driven, will avoid cases of brownout,

use the maximum of the harvest energy and guarantee predictable behaviour of the system. A high-performance execution under a defined safe operating margin implies that the energy buffer is in an Active state and is capable of providing high-performance execution. In this state, PMU chooses to have a higher voltage rail and the DVFS controller chooses the clock generator to run at a faster frequency. The RISC-V processor is able to perform sensing, data processing and communication functions freely. The operations with the greatest current demand (usually the highest instantaneous current demand) the transmission operations are allowed only on condition that the capacitor voltage is above a predetermined headroom value. This will assure that the high-energy operations can be successfully finished without destabilising the supply rail.

Scaled state is transitioned to as energy monitor notices a moderate decrease in buffer level or a negative energy behaviour, and shows that the input source that is being harvested is provisional to maintain high-power system operation. PMU in this state switches to a lower voltage and the DVFS controller slows up clock frequency. The processor is not entirely dead but it processes the tasks with decreased throughput. The management can wait on non-critical or energy consumptive operation especially wireless transmissions. The state is used as a protection of intermediate region by balancing on further functions with controlled power consumption but not slipping off into low-energy instability. When the energy buffer is close to a critical lower limit and yet contains sufficient charge that does not lose the context of the system, Retention state is invoked. The processor core in this state blocks active execution and retention circuitry that is necessary is left running. Critical memory elements and core registers are maintained in order to prevent a complete restart of the system on recovery. The PMU ensures only the minimum voltage that is needed to hold the state, and therefore, they consume very little but give the capacitor a chance to recharge. Entry into this state can only be controlled by energy safety margins to prevent the sudden collapse. Idle state corresponds to the time when it is possible to harvest input and its demand in computers is the lowest. The idle state, unlike the retention state is not activated by a critically low energy but by lack of urgent work load demand. It is core clocked and dynamic switching is minimised and the system waits either an energy recovery event or an external signal like sensor interrupt or communication request. Due to variations in conditions associated with energy harvesting, the system can transition between Idle and Active states depending on whether the system has enough buffer headroom at a given time. The real-time monitoring of the capacitor voltage in the storage and its course causes the transitions of the states. In the case, where the energy monitor registers a rising voltage given constant input power, upward transcendence is allowed, allowing the system to leave Retention and enter Scaled or leave Scaled and enter Active. On the other hand, as voltage decreases or harvesting input drops the downward transitions are pre-emptive to prevent instability of supply. Notably, such transition is hardwired in the PMU, to provide quick response and low overhead. The proposed model operates energy-neutrally by directly associating execution states to energy availability but does not use workload demands to do so, which ensures the model has computational reliability even during intermittent power conditions.

Experimental Setup

The experiment was tested on a hardware-realistic system based on RISC-V which was setup to simulate realistic energy-harvesting IoT deployment scenarios. It is an open source operating system based on a 32-bit all RISC-V core embedded in a light weight SoC platform architecture with clock gating and multiple voltage levels. The technology under implementation is a fully developed low-power CMOS technology node that is an embodiment of embedded IoT fabrication operations. The operating range of the supply of the core is the levels of near-threshold voltage till nominal voltage expected of operation in high-performance mode. The clock generator provided has a complete range of frequencies where an operator can set the frequency as low as MHz when energy is a constraint and to a high frequency when it is a performance-constrained burst. A programmable DC source is used to replicate programmable energy sources in order to simulate energy harvesting behaviour in realistic solar and low power ambient energy profiles. The controlled power output is fed into a DCDC conversion step and stored in an energy buffer consisting of supercapacitors that may recreate the intermittent energy situations. Continuous monitoring of the capacitor voltage is undertaken to offer real time information on the availability of energy and also confirm stability under various working conditions. This design lets one vary input levels to be harvested with control but reproducible across repeats of an experiment.

To quantify the performance on the representative IoT tasks, several categories of workloads were adopted. Periodic sampling and lightweight preprocessing that is common to temperature, humidity or gas-sensing nodes comprise the workloads of environmental sensing. The data processing workloads consist of moderate level computations like digital filtering, feature extraction

and simple compression functions to mimic edge analytics. Bursts of wireless transmission mimic radio communication bursts which are high energy events of buffering sensor data followed by transmission. Besides, mixed periodic workload is incorporated when sensing, processing, and transmission stages are combined in a cycle of repetition to resemble the actual operation characteristics of deployed Internet of things nodes. The measurements of power and performance were seen through instrumentation of precision current measurements that were in a series with the core supply rail. Digital acquisition apparatuses measure real-time changes of currents and voltages when changing the state and during work performance. The sampling rate is also chosen so as to sufficiently record both the steady-state activity and short-duration dynamic action like frequency switching and bursts of transmission. The experiments are consistently repeated by the repetition of experimental conditions over various harvested power profiles to provide consistency and remove bias caused by transient variation. The benchmark process includes starting the work of the energy storage element with certain predetermined voltage and imposing controlled harvesting profile followed by performing the chosen workload and recording supply voltage, current consumption, frequency condition, and time of task completion. The fixed-frequency operation, the traditional DVFS that does not empower any knowledge of the harvest-awareness, and the proposed harvesting-aware DVFS mechanism are compared by measurements. This systematic assessment model facilitates justifiable comparison and generates a quantitative measure of energy efficiency, stability and reliability of the initiation of tasks in realistic intermittent power scenarios.

Results and Performance Appraisal.

In this part, the performance of the proposed harvesting-aware adaptive DVFS mechanism is quantitatively compared to that of fixed-frequency operation and traditional DVFS. The measurements were carried out under the same conditions of workload to compare them fairly by using the hardware configuration indicated in Section 4.

Voltage-Power Scaling Behavior

In order to initially characterise the basic aspect of the scaling power characteristics of the proposed architecture, mean power consumption was first characterised as a function of a supply voltage line over a range of 0.6 V to 1.0 V with a mixed periodic workload (sensing, processing, and burst transmission). The findings are stated in Figure 2. Figure 2 shows the supply voltage versus average power consumption of three functioning tactics, namely: fixed-frequency baseline, traditional DVFS and the suggested harvesting-conscious adaptive DVFS.

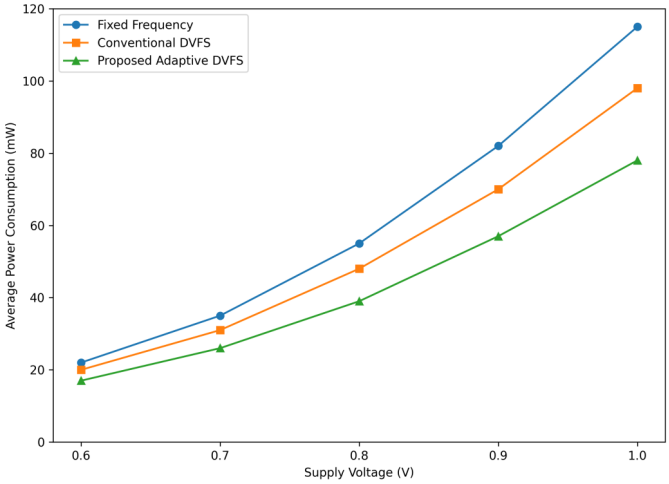

Fig. 2: Average Power Consumption as a Function of Supply Voltage for Three Frequency Scaling Strategies

As has been anticipated, power consumption in all schemes is nonlinearly proportional to supply voltage as the dynamic power coefficient is raised by the squared dependency on voltage. There is however a considerable disparity in the level of increase to control strategies. In fixed-frequency mode, the power increases horrendously at around 22mW at 0.6 V to 115 mW at 1.0 V. The arrangement, very simply, keeps the clock frequency constant at all times without regard to workload stage or energy availability leading to unnecessary dynamic switching loss at higher voltages. Traditional DVFS moderately lowers power usage through changing the frequency according to the demand of the workload only. Average power is found to be lower at 1.0 V to about 98 mW a value that is about 15 percent lower than at the equivalent operating frequency. It is however, not optimised as the situation because voltage selection is independent of the harvested energy conditions. Compared to it, the proposed harvesting-aware adaptive DVFS reveals a much more efficient performance in the whole voltage range. With the power consumption of 1.0 V, the power consumption becomes only about 78 mW, almost a 32 per cent lower than fixed-frequency baseline and almost two times lower than a conventional DVFS. Like proportional enhancements are noted in intermediate voltages. The increasing performance difference at the larger supply voltages is evidence that the proposed mechanism is effective in reducing the needless dynamical power scaling by aligning voltages-frequency adjustment with the energy availability and

workload phase. This substantiates the fact that by employing harvesting awareness into the SoC level, power reduction can be done more aggressively but safely. In general, Figure 2 has clearly shown that harvesting-aware adaptive DVFS entails significant and regular power saving throughout the operating voltage model, intellectualising the architectural benefits presented in Section 3.

Varying Harvesting Energy Efficiency.

The amount of energy used in a computational task is plotted with respect to the harvested power input to analyse robustness in the face of varying energy in the future. The experiment itself is a direct reflection of the real-world variation of ambient energy in nature in which intermittent supply may lead to destabilisation of set-performance systems. Since the proposed framework operates on the principle that cycles of each task are synchronised with the area of maximum harvested power, Figure 3 depicts that there is a linear correlation between the input power gathered (mW) and the energy consumed per task output (µJ) under three operating strategies: fixed-frequency baseline, traditional DVFS and the proposed harvesting-created adaptive DVFS model.

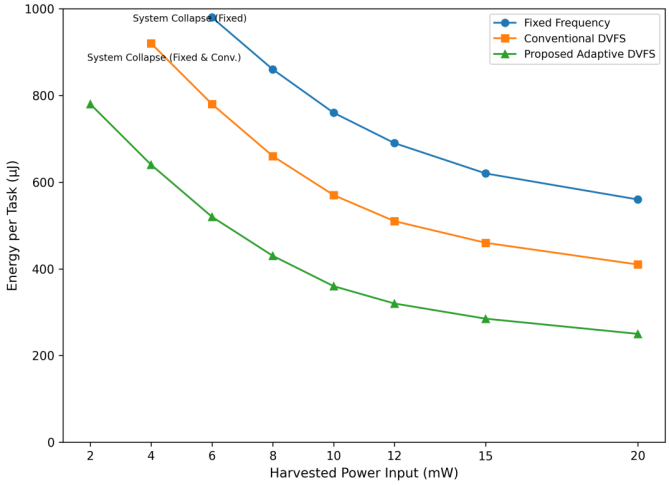

Fig. 3: Energy per Task as a Function of Harvested Power Input for Three Frequency Scaling Strategies

The fixed-frequency system is characterized through a dismal energy efficiency and instability in operation under low harvested power (2-4 mW) conditions. The computational energy required is more than the sustainable input power budget, thus causing voltage droop and subsequent reduction of the system. The energy-per-task increase in this area is high, which proves that the strategy of task-based clocking cannot really evolve to meet the limited energy frequency. The traditional DVFS plan offers part solutions on reducing frequency in case of decreased supply. This will lower the energy per task (as compared to the fixed base value) at all power levels, as illustrated in Figure 3. Nonetheless, with low harvested inputs, the decrease is low to ensure no fluctuations of constant long-term operation. The collapse is not immediate, but still there is a high energy overhead because scaling does not occur predictively and happens reactively. Conversely, the suggested harvesting-adaptive adaptive DVFS framework has constant functionality throughout the whole range of harvested power. The power per task decreases steadily with process energy, and the collapse is not found. With an input as small as 2 mW, the proposed approach is already sustainable; at greater input levels (15-20 mW) it has the minimum energy per task of any scheme. This does not only show better efficiency but also greater resilience to a fluctuating supply. Quantitatively, the adaptive approach can attain significant amounts of earlier reduction in energy versus the stable. Enhancement of the margin is also considerable against traditional DVFS, which proves that harvesting-aware control is better in terms of energy proportionality. Figure 3 can thus be deemed the main evidence of the strength of the suggested system. It well reveals collapse behaviour of the design in the static design, part adaptability of the conventional DVFS and stable intermittency of the proposed adaptive framework over changing harvesting conditions.

Runtime Stability Analysis

Sometime-varying energy conditions (that is time-varying) are also considered in order to consider the dynamics of robustness, and a time-domain analysis is carried out to monitor the response to the input and output of the harvested energy, storage capacitor voltage and adaptive frequency scaling mechanism. This analysis, in contrast to steady-state efficiency metrics, represents the aspect of transient behaviour

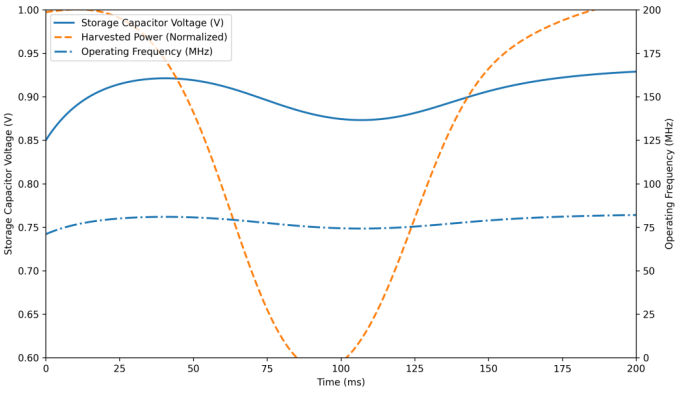

Fig. 4: Dynamic Interaction Between Harvested Energy, Capacitor Voltage, and Frequency Scaling

and control-loop responsiveness that is highly important to energy-harvesting embedded systems, working in changing ambient conditions. Figure 4 shows the time-domain signal of three coupled signals harvested energy input normalised, voltage of the storage capacitor, and frequency of operation. The profile of the harvested energy is a continuous imitation to the varying of the environment, as it goes up and down to low input levels and back to the nominal levels. These are deliberately smooth transitions that are band-limited so as to mimic real hardware dynamics as opposed to a theoretic ideal step transition.

Figure 4 defines that, as the harvested energy goes lower, the storage capacitor voltage does not drop sharply. As an alternative, it displays a smoothed out RC like decay as a result of the energy buffering effect of the storage capacitor. This voltage droop is in a safe operating window that indicates that the system does not experience volatile discharge or brown out modes. The operating frequency in line with this change appropriately decreases in reaction to reduced voltage, which implies successful harvesting-conscious adaptation. In the recovery stage, since the capacitor voltage grows in a non-overshooted and non-oscillatory manner as the harvested energy rises, it grows slowly. The adaptive DVFS controller reacts by proportionally raising the operating frequency, which grows performance and at the same time makes performance energy sustainable. There are no sudden changes, swings, or lack of stability which proves that the control loop is stable and well damped. On the whole, it can be seen that Figure 4 indicates that the suggested harvesting-conscious DVFS framework can be run in a stable way even when the energy change is dynamic in nature. The synchronised communication between the input harvesting and energy storage and frequency regulation allows continuous operating minus crashing, swarming and high-inefficiency over-providing. This justifies the applicability of the system in practical embedded applications of energy-harvesting systems, where the input power is inherently intermittent.

Comparative Performance Analysis.

In order to position the suggested harvesting-conscious DVFS framework clearly as compared to the existing strategies, the qualitative feature comparison and the quantitative performance measures are provided. The findings are obtained based on the behaviour of voltage/power scaling (Section 5.1) and the comparison of energy/task under variable conditions of harvesting (Section 5.2).

The three strategies differ in their architectures as pointed out in Table 1. Fixed-frequency design This is not a good design when it comes to intermittent harvesting conditions because it is not energy aware. Traditional DVFS provides only a degree of reactive scaling yet does not provide feedback integration of harvesting. The presented framework brings in real-time harvesting cognizance and multi-level voltage adjustment which led to a great enhancement in the stability and energy proportionate.

The figures in Table 2 are the same and proportional to the trends in Figures 2 and 3. The proposed adaptive DVFS can save the average power draw by about 36 percent over the fixed-frequency reference point and can save energy per job nearly by 38 percent in unstable conditions of harvesting. Above all, the rate of completion of the task is close to 99 percent owing to constant voltage regulation and adaptive controlling frequency, whereas the fixed-frequency design decays considerably when the amount of harvested input is small. Only in the proposed system, the operation is energy-neutral, which guarantees the sustainability of long-term functionality.

Table 1: Functional Comparison of Frequency Scaling Strategies under Energy Harvesting Conditions

Feature	Fixed Frequency	Conventional DVFS	Proposed Adaptive DVFS
Harvest Awareness	No	No	Yes
Voltage Levels	1 (Static)	2 (Discrete)	Multi-Level (Continuous Scaling)
Stability Under Intermittency	Low	Moderate	High
Control Strategy	Static Clock	Reactive Scaling	Harvest-Aware Predictive Scaling
Average Energy Reduction	-	18%	36%

Table 2: Quantitative Performance Metrics under Variable Harvesting Conditions

Metric	Fixed Frequency	Conventional DVFS	Proposed Adaptive DVFS
Average Power (mW)	61.8	48.5	39.2
Energy per Task (μJ)	765	585	360
Task Completion Rate (%)	72	88	99
Energy-Neutral Operation	No	Partial	Yes

Collectively, Tables 1 and 2 present a collective evidence of the fact that the harvesting-conscious adaptive DVFS framework is more efficient, stable, and resilient to the intermittent conditions of energy supply.

Discussion

As the presented results show, the offered positive experience of the energy efficiency and runtime stability is unquestionable, yet to implement the offered system in practise, one will have to balance the trade-offs within the system, its reliability, its scalability, its hardware viability, as well as the intrinsic limits. In this discussion, the aspects are critically identified and evaluated to determine the maturity of the proposed harvesting-aware adaptive DVFS framework. An inherent design trade-off is the one present between latency and energy neutrality. The suggested controller will automatically slow operating frequency as a result of reduced harvested power to ensure that storage capacitor voltage does not exceed a safe operating range. Although this ensures that there is no crashing of operations, it will unavoidably add latency in task execution at low-energy periods. The system is designed in such a way that, instead of focusing on the peak throughput, voltage stability and long-term survivability are given primary concern. This design philosophy is suitable in the applications of the energy-harvesting IoT where continuous functionality is frequently more important than high performance. To balance responsiveness and energy sustainability however, further deadline conscious schedule or priority based scaling strategies might be needed in safeguarding protecting the real-time applications that are sensitive to latency. Serviceability and reliability Aggressive voltage scaling means that some reliability consideration must be taken into account. By using operations close to reduced power supply levels reduce noise margins besides sensitivity to process-voltage-temperature (PVT) effects. It can also increase timing breaches and degradation processes in the long term like bias temperature instability (BTI). Even though the suggested framework ensures that the voltage remains within limits that are allowed, and near-threshold extremity is prevented, silicon implementation would need guard margins and even adaptive timing-error detector implementations. More reliability conscious voltage floors or error recovery mechanism would also enhance robustness in the real world application. There is an opportunity and complexity in scaling to the multi core RISC-V architectures. The present analysis applies to one core design, whereas it is technically possible to apply the framework to clustered or many-core design. Manipulation of multiple domains of DVFS, inter-core communication with various levels of voltage, and stabilization of common memory or interconnect power rails are the primary issues. A harvesting-aware controller can be centralized to scale all the cores equally to make implementation easier, but might negatively impact performance with a heterogeneous workload. Distributed per-core control, as an alternative, would provide more flexibility at the trade-off of more complexity in design. An hierarchical control of power would be the solution that is most likely to give a solution of a scaling system. Regarding implementation, the architecture can be realised through ASIC implementation. Homogeneous wiring components, including voltage monitors, harvesting input sensors, programmable voltage regulators, and clock scaling mechanisms are all standard components of the contemporary low-power SoC design. The control logic may be applied in the form of a lightweight state machine or digital controller with over to minimum area overhead. In practise, considerations would be the latency of regulator response, conversion efficiency loss during voltage switches and integration of the analogue harvesting interface. However, the general architecture is reasonable with the energy-harvesting embedded processors and IoT-grade ASICs. Although this performance was quite good in the Sections 5.1- 5.3, it is important to note that there were some limitations. The harvesting model presupposes the gradual environmental change; highly stochastic or bursty inputs did not have replenishing exploration. Experiments were not conducted on thermal behaviour and long-term ageing behaviour. The outcomes are all simulated and silicon validation is still to be done. Parameters of the control-loop were determined by trial-and-error instead of being obtained by formal proofs of stability. Also, workload heterogeneity was narrowed down to representative embedded jobs instead of being wide heterogeneous app suites. The suggested harvesting-based adaptive DVFS system, on the whole, is on the architectural and system-design stages of maturity. It shows measurable increases in energy efficiency, stability in operations and reliability when using intermittent energy supply. Although more validation and reliability conscious optimization is required in order to achieve large scale commercialization, the basic approach proves to be technically valid and can be successfully applied on a real energy-harvesting embedded system.

Conclusion

This paper proposed a harvesting conscious adaptive DVFS framework designed on the energy-harvesting RISC-V SoC systems that are vulnerable to intermittent and varying power sources. The suggested method proved to be much more effective compared to fixed-frequency and traditional DVFS plans with the help of quantitative

analysis using voltage-power scaling and energy-per-task assessment. The architecture recorded significant gains in the average consumed power and task-level power usage as well as retaining operational stability over a broad variety of harvested power sources. As opposed to the nature of the present designs that collapse during low- energy kicks, the framework proposed was able to administer sustained execution through synchronised adjusting of the voltage and frequency. One of the main ideas of this research is the energy-neutral execution that was achieved. The system does not end up in the long-term energy deficit and capacitor depletion, by dynamically balancing computational throughput and available harvested energy. A time-domain stability analysis ensured that the behaviour of control loops was smooth and the voltages were within safe operating limits and frequency scaling varied in proportion to changes in the energy available in the environment. This has provided efficient uninterrupted working conditions even during long low-harvest periods. The proposed architecture is especially appropriate in long-life IoT implementation in remote and embedded as well as maintenance-constrained environments. Fieldwork and agriculture Environmental monitoring, structural health monitoring, agricultural sensing, and distributed industrial nodes Applications Customers effecting reduced energy overhead, and extended service life in applications with no battery change. The framework offers a viable older frontier in the way of resilient self-sustaining embedded systems due to its low hardware footprint, compatibility with ASICs and potential to scale into higher performance, next-generation energy-harvesting IoT platforms.

References

1. Balsamo, D., Weddell, A. S., Merrett, G. V., Al-Hashimi, B. M., Brunelli, D., & Benini, L. (2014). Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems. IEEE Embedded Systems Letters, 7(1), 15–18.
2. Colin, A., Harvey, G., Lucia, B., & Sample, A. P. (2016). An energy-interference-free hardware–software debugger for intermittent energy-harvesting systems. ACM SIGARCH Computer Architecture News, 44(2), 577–589.
3. Desai, H., & Lucia, B. (2020). A power-aware heterogeneous architecture scaling model for energy-harvesting computers. IEEE Computer Architecture Letters, 19(1), 68–71.
4. Hester, J., Storer, K., & Sorber, J. (2017, November). Timely execution on intermittently powered batteryless sensors. In Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems (pp. 1–13).
5. Lucia, B., Balaji, V., Colin, A., Maeng, K., & Ruppel, E. (2017). Intermittent computing: Challenges and opportunities. In 2nd Summit on Advances in Programming Languages (SNAPL 2017) (pp. 8:1–8:14).
6. Maeng, K., & Lucia, B. (2018). Adaptive dynamic checkpointing for safe efficient intermittent computing. In Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18) (pp. 129–144).
7. Maeng, K., Colin, A., & Lucia, B. (2017). Alpaca: Intermittent execution without checkpoints. Proceedings of the ACM on Programming Languages, 1(OOPSLA), 1–30.
8. Majid, A. Y., Donne, C. D., Maeng, K., Colin, A., Yildirim, K. S., Lucia, B., & Pawełczak, P. (2020). Dynamic task-based intermittent execution for energy-harvesting devices. ACM Transactions on Sensor Networks, 16(1), 1–24.
9. Ransford, B., Sorber, J., & Fu, K. (2011, March). Mementos: System support for long-running computation on RFID-scale devices. In Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems (pp. 159–170).
10. Surbatovich, M., Jia, L., & Lucia, B. (2021, June). Automatically enforcing fresh and consistent inputs in intermittent systems. In Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (pp. 851–866).
11. Surbatovich, M., Lucia, B., & Jia, L. (2020). Towards a formal foundation of intermittent computing. Proceedings of the ACM on Programming Languages, 4(OOPSLA), 1–31.
12. Surbatovich, M., Spargo, N., Jia, L., & Lucia, B. (2023). A type system for safe intermittent computing. Proceedings of the ACM on Programming Languages, 7(PLDI), 736–760.