# Power Neutral Performance Scaling for Energy Harvesting MP-SoCs

Benjamin J. Fletcher, Domenico Balsamo and Geoff V. Merrett

Department of Electronics and Computer Science

University of Southampton, UK

Email: {bjf1g13, d.balsamo, gvm}@ecs.soton.ac.uk

*Abstract*—Using energy 'harvested' from the environment to power autonomous embedded systems is an attractive ideal, alleviating the burden of periodic battery replacement. However, such energy sources are typically low-current and transient, with high temporal and spatial variability. To overcome this, large energy buffers such as supercapacitors or batteries are typically incorporated to achieve *energy neutral* operation, where the energy consumed over a certain period of time is equal to the energy harvested. Large energy buffers, however, pose environmental issues in addition to increasing the size and cost of systems. In this paper we propose a novel *power neutral* performance scaling approach for multiprocessor system-on-chips (MP-SoCs) powered by energy harvesting. Under power neutral operation, the system's performance is dynamically scaled through DVFS and DPM such that the instantaneous power consumption is approximately equal to the instantaneous harvested power. Power neutrality means that large energy buffers are no longer required, while performance scaling ensures that available power is effectively utilised. The approach is experimentally validated using the Samsung Exynos5422 big.LITTLE SoC directly coupled to a monocrystalline photovoltaic array, with only 47mF of intermediate energy storage. Results show that the proposed approach is successful in tracking harvested power, stabilising the supply voltage to within 5% of the target value for over 93% of the test duration, resulting in the execution of 69% more instructions compared to existing static approaches.

## I. Introduction

With recent developments in compiler and runtime optimisation, multicore systems are able to outperform single-core systems in both performance and power consumption [1]. As such, heterogeneous multiprocessor system-on-chips (MP-SoCs) are rapidly becoming the de-facto technology for powering modern, high-performance embedded devices. When considering autonomous embedded devices, such as typical Internet of Things (IoT) end-devices, maintaining a low power-budget is essential. This is due to the fact that they are often battery powered, meaning that their lifetimes becomes restricted by battery capacity and discharge rate [2]. Motivated by the limited, finite lifetimes achievable using batteries, research has more recently looked to supplement, or even replace batteries by means of energy harvesting (EH). This is where devices scavenge ambient energy from their environments [3] (e.g. from vibrations, light, thermal gradients etc) in order to be self-sufficient. Although this provides a sustainable solution whereby lifetimes become potentially infinite [4], the
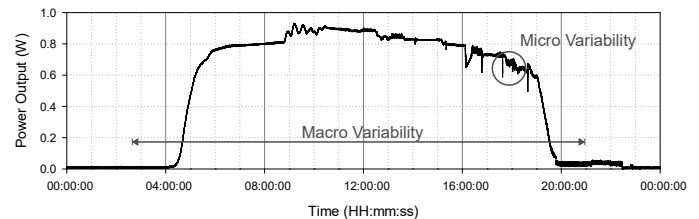
Fig. 1. Experimentally obtained data showing the varying power output of a $250cm^2$ solar cell over the course of a day.

energy harvested from these sources is inherently transient and unpredictable.

Consider the power output of a solar cell array, shown in Fig. 1. Here, the harvested power exhibits both 'micro' variability (where the power changes rapidly due to shadowing) and 'macro' variability (where the power changes more slowly over the course of the day). Conventionally, in order to deal with this unpredictable energy harvest, an energy buffer is employed between the harvester and the device to smooth $V_{CC}$ and make it appear relatively constant, similar to a battery-powered system. Typically, in EH systems this energy buffer takes the form of a supercapacitor due to their increased power densities and enhanced operational lifetimes [5]. Considering a simple EH system (Fig. 2), if the supercapacitor is sufficiently large that the microcontroller operates perpetually, this is known as *energy neutral* operation. Under energy neutrality, the energy harvested in times of copious harvest compensates for the deficit when the harvest is diminished [6]. As such the energy consumed is equal to the energy harvested over a period of time, $T$ (e.g. typically 24 hours for solar powered devices).

Whilst energy neutral operation offers a viable method of overcoming variability in EH supplies, it also has weaknesses. The addition of energy storage introduces energy losses due to parasitic leakage currents, causes deterioration in performance over time [7], increases device size and cost [8] and requires the addition of maximum power point tracking (MPPT) circuits and algorithms to operate efficiently.

To address these issues, research has targeted *storageless* systems powered directly by energy harvesting, without additional energy buffers. 'SolarTune' [9], demonstrates such a system where a multicore CPU platform is coupled directly to a photovoltaic (PV) EH source. SolarTune uses harvesting-
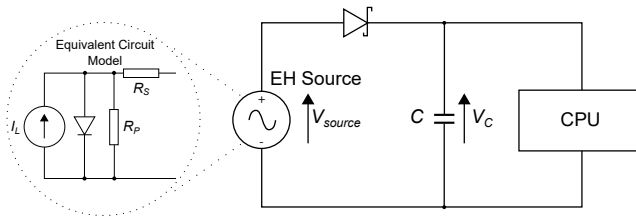
Fig. 2. Schematic illustrating the circuit topology of a typical simple system powered by solar energy harvesting.
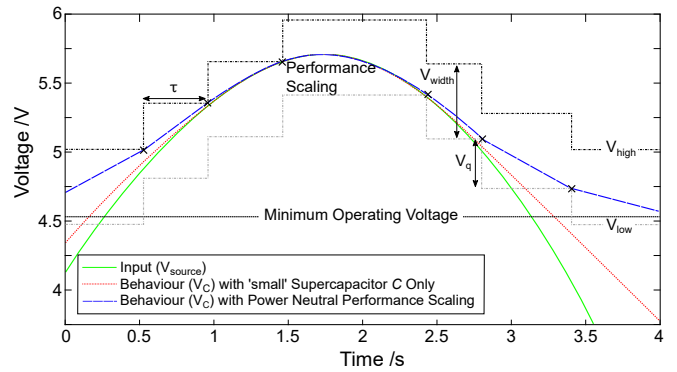


Fig. 3. Behaviour of an EH system to a transient input (green), with (blue) and without (red) power neutral performance scaling.
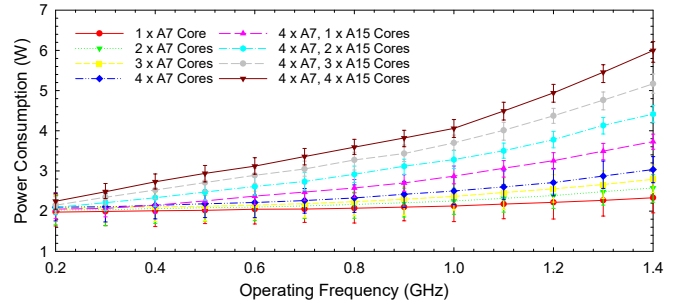


Fig. 4. Board power consumption vs operating frequency for multiple core configurations, experimentally obtained for the ODROID XU4 embedded SoC platform whilst running CPU intensive ray tracing software [12].

aware runtime task scheduling (similar to [10]) to adapt performance with respect to the predicted availability of harvested solar energy such that no significant energy buffer is required [9]. These schemes, however, rely heavily upon accurate prediction of future availability of harvested power, making them unsuitable for use with sources exhibiting significant 'micro' variability which is almost impossible to predict.

The *power neutral* operating paradigm [11], where a device's performance is dynamically modulated such that instantaneous power consumption matches the instantaneous harvested power, negates the need for energy storage. The authors demonstrate a practical implementation on an ultra-low power single-core MCU where power consumption is controlled through dynamic frequency scaling (DFS). This instantaneous performance scaling approach has the advantage that it does not rely upon the prediction of available harvested power, making it much more suitable for EH sources with significant 'micro' variability.

In this paper, we extend the concept of power neutral operation to a heterogeneous multi-core applications processor, and propose a novel performance scaling approach (Section II) by controlling both DVFS (dynamic frequency and voltage scaling) and DPM (dynamic power management). The approach is modelled (Section III) and experimentally validated (Sections IV and V) using energy harvested by photovoltaic cells.

## II. POWER NEUTRAL PERFORMANCE SCALING USING DVFS AND DPM

Fig. 3 illustrates the concept of power neutrality. Here, the EH system depicted in Fig. 2 is considered, where the harvested power varies sinusoidally such that the source is transient (shown in green). The addition of a tiny capacitance results in a marginal lifetime increase (shown in red), however, if the device's performance (and hence power consumption) is gracefully reduced, a much more significant increase can be realised (shown in blue), potentially alleviating the burden of system hibernation.

In MP-SoCs, this 'on-the-fly' performance scaling can be best achieved through DVFS and DPM (in this work, we provide DPM through the enabling and disabling of CPU cores at runtime, also known as *core hot-plugging*). The combination of these two power management techniques results in a variety of operating performance points (OPPs), each with associated power consumptions. Fig. 4 shows this for a typical

heterogeneous MP-SoC platform (the Samsung Exynos5422 big.LITTLE SoC used for experimental validation in Section IV, featuring four 'big' high performance ARM A15 cores, and four 'LITTLE' low-performance ARM A7 cores).

The challenge when devising an effective power neutral control algorithm for heterogeneous MP-SoCs is quickly, and accurately, identifying the correct OPP at any given instant, such that the consumed power is as close as possible to, without exceeding, the available harvested power.

### A. Proposed Control Approach

The proposed approach uses two voltage thresholds, $V_{high}$ and $V_{low}$, which are dynamically adjusted as shown in Fig. 3. These thresholds monitor the input voltage $V_C$ across the small capacitor, $C$, which buffers any latency in the control system (detailed in Section IV-A). Fig. 3 also introduces two algorithmic parameters, $V_{width}$ and $V_q$ for which optimal values are selected through simulation in Section III.

Initially $V_{high}$ and $V_{low}$ are calibrated such that they bound $V_C$, with a voltage $V_{width}$ between them, explicitly:

$$V_{high}(t=0) = V_C + \frac{V_{width}}{2} \text{ and } V_{low}(t=0) = V_C - \frac{V_{width}}{2} \quad (1)$$

In the event of a reduction in harvested power, $V_C$ crosses the $V_{low}$ threshold and the device moves to a lower OPP (as determined below). Following this, the thresholds are reduced by $V_q$. If the harvested power continues to fall, this process is repeated such that the thresholds follow $V_C$, 'tracking' the

harvested power supply. A similar but opposite process occurs as harvested power increases and $V_C$ crosses $V_{high}$.

### B. Operating Performance Point Selection

Fig. 5 illustrates how the performance scaling response is determined in the event of a threshold being reached. Initially, as the latency associated with frequency scaling is typically lower than that associated with core hot-plugging [13], DVFS is performed to deal with 'micro' variation in the harvested power supply. More specifically, linear control is applied, and the system's operating frequency is migrated to the next lowest of $N$ predefined operating frequency levels $(f_0, f_1 \ldots f_{N-1})$.

Secondly, to deal with the 'macro' variation in the harvested supply, derivative control is applied to calculate a core hot-plugging response so that the number of active cores is proportional to $dV_C/dt$.

To explain the way in which the core hot-plugging response is determined, given that we are working within a heterogeneous multicore architecture, it is useful to define two ternary 'core scaling factors': $S_b$ for 'big' cores, $S_L$ for 'LITTLE' cores, where:

$$S_x = \begin{cases} 1 & \text{denotes the addition of a core of type } x \\ 0 & \text{denotes no alteration} \\ -1 & \text{denotes the removal of a core of type } x \end{cases}$$

Two constant gradient threshold parameters $\alpha$ and $\beta$ are also defined for 'LITTLE' and 'big' cores respectively, which represent the minimum gradient required to warrant a change in the existing core configuration.

$$S_b = \begin{cases} 1 & \text{if } \frac{dV_C}{dt} > \beta \\ -1 & \text{if } \frac{dV_C}{dt} < -\beta \\ 0 & \text{otherwise} \end{cases}, S_L = \begin{cases} 1 & \text{if } \frac{dV_C}{dt} > \alpha \\ -1 & \text{if } \frac{dV_C}{dt} < -\alpha \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

To minimise overhead, $dV_C/dt$ is approximated each time $V_C$ crosses a threshold as:

$$\frac{dV_C}{dt} \approx \frac{\Delta V_C}{\Delta \tau} = \frac{V_q}{\tau} \quad (3)$$

Where $\tau$ is the time which has elapsed since the previous approximation, as shown in Fig. 5. Substituting this into (2) and considering the case when the $V_{low}$ threshold is crossed, the core scaling response is:

$$S_b = \begin{cases} -1 & \text{if } \tau < \frac{V_q}{\beta} \\ 0 & \text{otherwise} \end{cases}, S_L = \begin{cases} -1 & \text{if } \tau < \frac{V_q}{\alpha} \\ 0 & \text{otherwise} \end{cases}$$

A 'big' core is removed if the rate of change in $V_C$ is greater than $\beta$, a 'LITTLE' core is removed if the rate of change in $V_C$ is greater than $\alpha$, else, the core configuration remains constant. Again, a similar but opposite control flow occurs in the event of an increase in harvested power.

### III. SYSTEM MODELLING AND SIMULATION

To obtain values for parameters $V_{width}, V_q, \alpha$ and $\beta$, the proposed approach was simulated using Matlab-Simulink. A model describing the power-OPP characteristics of the MP-SoC platform was created using experimentally obtained data.
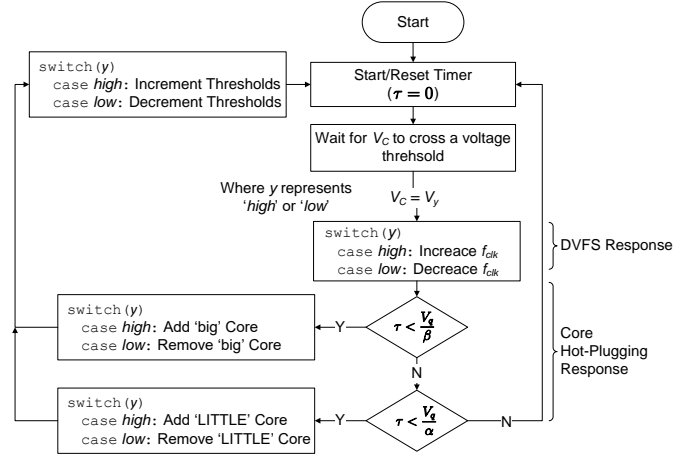


Fig. 5. Flowchart depicting the principles of the power neutral performance scaling approach using DVFS and core hot-plugging.
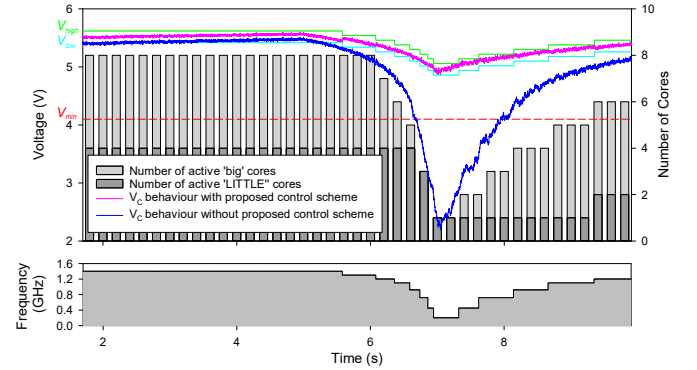


Fig. 6. Simulation showing operation of the control algorithm. Parameters $V_{width}$=0.2V; $V_q$=80mV; $\alpha$=0.1Vs$^{-1}$; and $\beta$=0.12Vs$^{-1}$.

This model was used in conjunction with the standard Simscape blockset to realise the proposed system shown in Fig. 2. The PV EH source was modelled using the solar cell equivalent circuit shown, described mathematically by:

$$I = I_l - I_0\left(\exp\left(\frac{V + R_s I}{N V_T}\right) - 1\right) - \frac{V + R_s I}{R_p} \quad (4)$$

where $I_0$ is the diode saturation current, $R_s$ and $R_p$ are the series and parallel resistances, $V_T$ is the thermal voltage, $N$ is the quality factor, and $I_l$ is the solar generated current. $R_s$, $R_p$ and $N$ were selected to approximate the behaviour of the PV array used for experimental validation in Section IV, and experimentally obtained solar irradiance data was used for $I_l$.

Two congruent systems were developed: one where OPPs are selected by a C program implementing the proposed control approach, and one where the performance is static. Eight DVFS frequencies corresponding to linearly spaced power consumption nodes were chosen, these are 0.2, 0.45, 0.72, 0.92, 1.1, 1.2, 1.3 and 1.4GHz. Simulations were performed using the Matlab ODE23 solver for multiple parameter combinations whilst assessing the control strategy's performance.
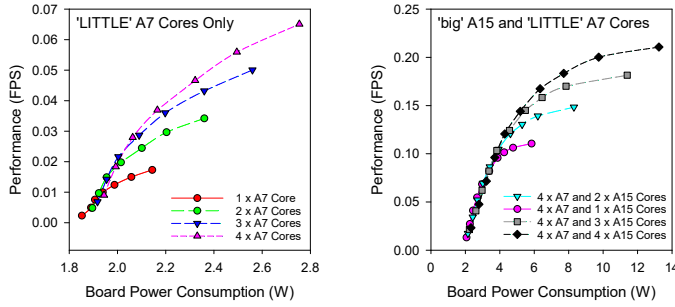
Fig. 7. Raytrace performance vs power consumption for the operating points in Fig. 4, obtained experimentally for the ODROID XU4 platform.
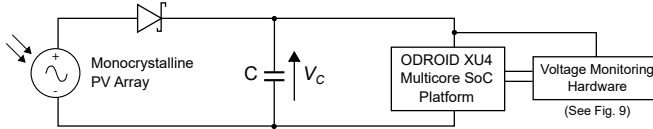


Fig. 8. Schematic of the proposed system architecture.

Fig. 6 shows the simulated behaviour of each system in a period of sudden shadowing. The blue line shows the behaviour of the system with static performance, and the magenta line shows the behaviour using the proposed approach. Here, the control algorithm is effective in scaling performance such that $V_C$ does not fall below the minimum operating voltage, $V_{min}$, as would happen without it. Through analysing the stability in $V_C$, more specifically the proportion of time spent within 5% of the target voltage, best performing values for parameters $V_{width}, V_q, \alpha$ and $\beta$ were determined as 144mV, 47.9mV, 0.120Vs$^{-1}$ and 0.479Vs$^{-1}$ respectively.

## IV. EXPERIMENTAL VALIDATION

The proposed approach was implemented on the ODROID XU4 development board, built around the Samsung Exynos5422 big.LITTLE SoC. The processor features 8 CPU cores (4× 'LITTLE' ARM A7 cores and 4× 'big' ARM A15 cores), and operates between 4.1V and 5.7V. The platform was benchmarked using a CPU intensive ray tracing application (smallpt [12]) to provide a parallelisable and intensive workload. Fig. 7 illustrates how the board's performance varies with power consumption across multiple OPPs. The performance metric used here is the number of frames rendered per second (FPS) at a quality of 5 samples per pixel.

Fig. 8 shows a schematic of the proposed system consisting of three parts: an EH source (in this case a PV array), a capacitor to buffer any latency in the control system (detailed in Section IV-A) and the load, an ODROID XU-4 running the power budgeting software proposed in this paper. In order to minimise software overhead, external low-power circuitry is used to generate hardware interrupts corresponding to $V_{high}$ and $V_{low}$. A schematic of this hardware is shown in Fig. 9. Here, a potential divider is used to coarsely reduce the input voltage to a level appropriate for the analogue comparator. An SPI controlled digital potentiometer is then used to more
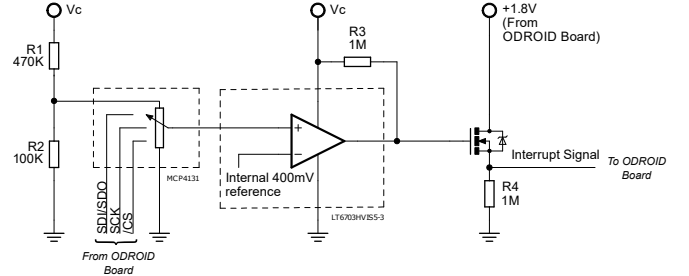


Fig. 9. Schematic of the voltage monitoring hardware.

finely reduce the voltage as directed by the processor, hence allowing the processor to set and adjust the threshold voltages. A comparator generates the interrupt signal which is passed through an n-channel MOSFET stage to convert the signal level for compliance with the development board. Overall, two of these circuits were used in order to facilitate both the *'Low'* and *'High'* dynamic voltage thresholds.

### A. Required Buffer Capacitance

Whilst this work aims to negate the need for the large energy buffers which are employed in energy neutral systems, some additional capacitance is required to support the system through the latency period when performing DVFS or core hot-plugging. Fig. 10 shows the overheads associated with both DVFS and core hot-plugging on the ODROID XU4. Using this data, an estimate of the additional required capacitance ($C$ in Fig. 8) can be obtained by considering the worst case scenario, where it is necessary for the system to modulate its performance from the highest OPP (and hence maximum power consumption) to the lowest OPP (minimum power consumption). If the system capacitance can harbour sufficient energy to tide the processor across this interval, it should respond robustly to any input, provided that it is within the bounds of feasible operation.

There are two ways in which the system may reduce its performance in response to a sudden drop in available power. It may, (a) perform DVFS followed by core hot-plugging, or (b) perform core hot-plugging followed by DVFS. Both of these scenarios were practically evaluated and the results are shown in Table I, in addition to the capacitance which would be required to support the board in each scenario whilst operating at the lowest voltage. It can be observed from Table I that, of the two approaches, approach (b) significantly outperforms (a), and hence the additional required capacitance ($C$ in Fig. 8) is 15.4mF. A 47mF supercapacitor was used for the experiments in this paper to provide a safety margin and to align with available components.

## V. RESULTS AND ANALYSIS

### A. Response to a Controlled Supply

Initial tests were performed using a controlled power supply without the additional capacitor $C$. This permitted verification that the SoC's performance adapts correctly to a changing input voltage. An example of a single data set is illustrated
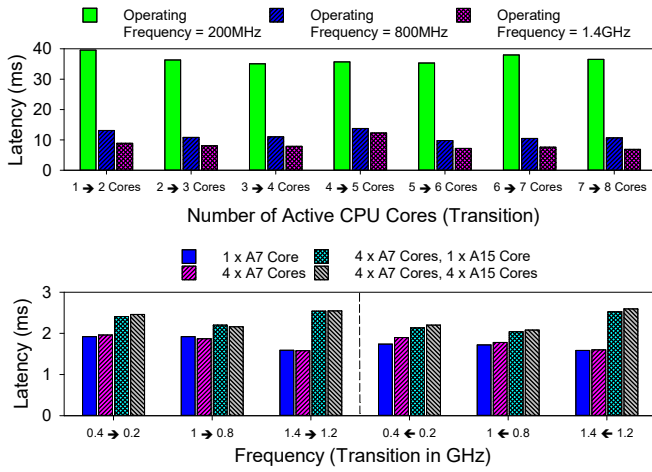
Fig. 10. Latency to switch number of active CPU cores using hot-plugging (top) and to change the operating frequency (bottom), obtained experimentally for the ODROID XU4 SoC platform.

Fig. 12. $V_C$ over time whilst testing the system under full sun conditions.

TABLE I
TIME AND CURRENT EXPENDED WHILST TRANSITIONING FROM THE
HIGHEST TO THE LOWEST OPP.

| Scenario | Transition Time, $\delta$ (ms) | $\int_0^\delta I\,dt = Q$ (C) | Required Capacitance, C (mF) |
|---|---|---|---|
| (a) Frequency, Core | 345.42 | 0.1299 | 84.2 |
| (b) Core, Frequency | 63.21 | 0.0461 | 15.4 |



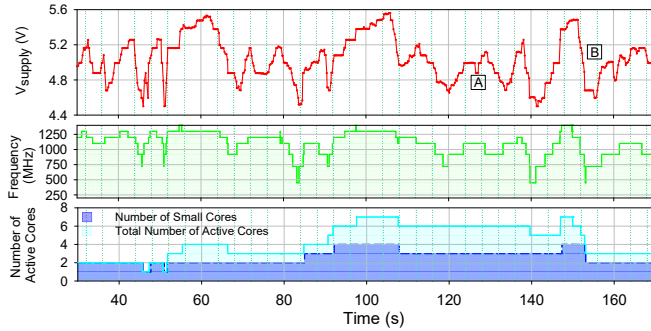Fig. 13. IV characteristics of the PV array and the proportion of time spent at each operating voltage.



Fig. 11. System performance using a controlled variable voltage supply. Parameters $V_{width}$=335mV; $V_q$=190mV; $\alpha$=0.238Vs⁻¹; $\beta$=0.633Vs⁻¹.

in Fig. 11. Here, large values of $V_q$ and $V_{width}$ have been chosen for clarity of illustration. The system can be observed responding as desired, modulating performance in correlation with the supply voltage. It can also be observed that core scaling is applied less often than frequency scaling, implying that the system is selecting long term and transient performance responses well. For example, minor fluctuations at 'A' are dealt with through DVFS only, whereas the response to the sudden reduction at 'B' involves the disabling of some 'big' and 'LITTLE' cores in addition to DVFS.

### B. Response to Energy Harvesting Supply (PV Array)

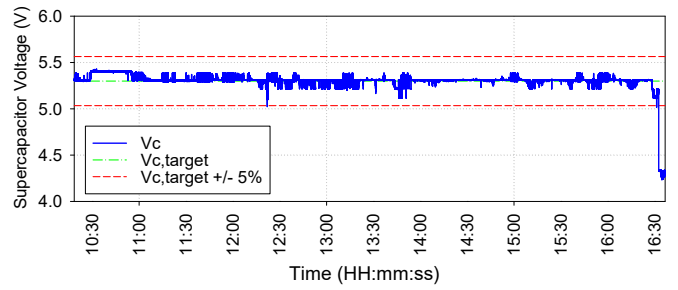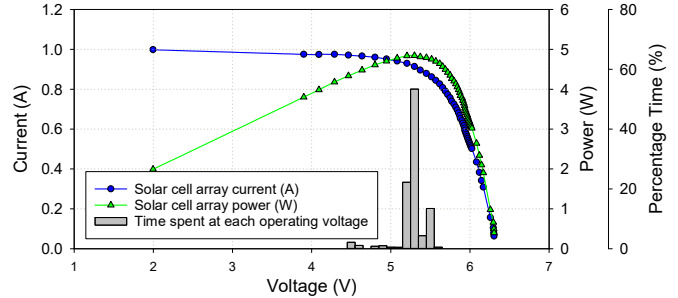The system was then tested using energy harvested from a 1340cm² monocrystalline silicon PV array. Fig. 12 shows the

response of the system over a six hour period, and illustrates the stabilisation effect that the scheme has on $V_C$ in full-sun conditions. The use of the proposed power neutral energy budgeting scheme results in $V_C$ remaining almost entirely (93.3% of the time) within ±5% of the target voltage. The target voltage was set at the solar cell's calibrated maximum power point (MPP), $P_{max}$, of 5.3V. Fig. 13 illustrates the effectiveness of the scheme's voltage stabilisation with reference to the IV characteristics of the PV array, ensuring that the board is always working at, or close to, the MPP. This negates the need for additional sizeable MPPT hardware.

The extent to which power neutrality is achieved by the proposed scheme is shown in Fig. 14 where the estimated power available and power consumed over the course of a day are both plotted for comparison. Available power estimation was performed by logging the open circuit voltage, $V_{oc}(t)$ of an identical, contiguous PV array, and using experimentally obtained IV data to determine corresponding values for $P_{max}(t)$. The board's power consumption can be seen to match closely the available power supply indicating that the system is making good use of the available harvested power supply, without exceeding it. Testing was performed for over 20 hours in a variety of weather conditions (full-sun, partial-sun, cloud, and hail) to examine performance. It was found that, in all cases, the system successfully managed to modulate its own performance with respect to the immediately available harvested power such that (provided the harvested supply was sufficient) the device could perpetually sustain operation.
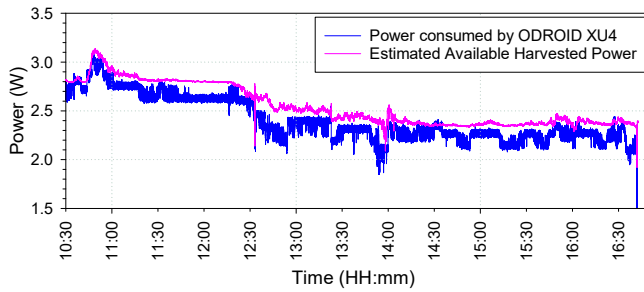
Fig. 14. Available (estimated) and consumed power over the course of a day.



Fig. 15. CPU usage over time, showing overhead of proposed approach.

TABLE II
PERFORMANCE OF POWER MANAGEMENT SCHEMES IN A 60 MINUTE TEST.

| Power Management Scheme | Avg. performance (Render/min) | Lifetime During Test (mm:ss) | Instructions Completed (Billions) |
|---|---|---|---|
| Linux Conservative | 1.0127 | 00:05 | 24.0 |
| Linux Powersave | 0.1456 | 60:00 | 2485.6 |
| Proposed Approach | 0.2460 | 60:00 | 4200.4 |

## C. Comparison with Linux Governors

The proposed performance scaling approach was also compared to each of the default Linux power management governors whilst harvesting energy from the solar PV array. During a one hour test, the results shown in Table II were obtained (Performance, Ondemand, and Interactive governors could not support any operation due to their high current requirements). The proposed performance scaling approach outperforms all of the available static Linux power management governors, allowing the system to operate for the full testing duration in addition to exhibiting enhanced performance when compared to the Powersave governor (which statically reduces performance to a minimum [14]), completing an estimated 69.0% more instructions over the same time period.

## D. Overheads of Proposed Approach

Fig. 15 illustrates the CPU usage of the proposed approach compared to that of the target application. Due to the interrupt based approach, the impact of the control scheme on CPU time is very low, averaging at 0.104% over the full testing duration. The power consumption of the additional voltage monitoring hardware was also measured, and found to be 1.61mW, less than 0.82% of the minimum system power consumption (and 0.01% of the maximum).

## VI. CONCLUSIONS

This paper has proposed a novel approach for achieving power neutral operation in energy harvesting multicore embedded systems. The power neutral performance scaling approach tracks harvested power through the use of two dynamic voltage thresholds, and alters performance accordingly through DVFS and DPM. The approach has been validated through both simulation and practical experimentation using a monocrystalline PV array. Results d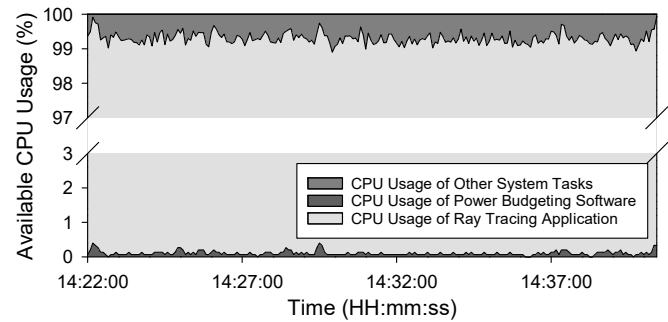emonstrate effective voltage stabilisation under all weather conditions, allowing the system to operate perpetually without the large energy buffers associated with conventional energy harvesting systems.

## REFERENCES

[1] NVIDIA Corporation. (2010) The benefits of multiple CPU cores in mobile devices. [Online]. Available: http://www.nvidia.co.uk/content/PDF/tegra_white_papers/Benefits-of-Multi-core-CPUs-in-Mobile-Devices_Ver1.2.pdf

[2] S. Peng and C. P. Low, "Throughput optimal energy neutral management for energy harvesting wireless sensor networks," in *IEEE Wireless Communications and Networking Conf. (WCNC)*, Apr 2012, pp. 2347–51.

[3] S. Beeby and N. White, *Energy harvesting for autonomous systems*. Norwood, MA: Artech House, 2010.

[4] A. Kansal *et al.*, "Power management in energy harvesting sensor networks," *ACM Trans. Embed. Comput. Syst.*, vol. 6, no. 4, Sept 2007.

[5] A. S. Weddell *et al.*, "Accurate supercapacitor modeling for energy harvesting wireless sensor nodes," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 58, no. 12, pp. 911–915, Dec 2011.

[6] A. Kansal *et al.*, "Harvesting aware power management for sensor networks," in *ACM/IEEE Design Automation Conf.*, 2006, pp. 651–656.

[7] V. Raghunathan *et al.*, "Design considerations for solar energy harvesting wireless embedded systems," in *Int. Symp. on Information Processing in Sensor Networks (IPSN)*, Apr 2005, pp. 457–462.

[8] J. A. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, no. 1, pp. 18–27, Jan 2005.

[9] Y. Wang *et al.*, "Solartune: Real-time scheduling with load tuning for solar energy powered multicore systems," in *Int. Conf. on Embedded and Real-Time Computing Systems and Applications*, Aug 2013, pp. 101–110.

[10] C. Moser *et al.*, "Real-time scheduling for energy harvesting sensor nodes," *Real-Time Systems*, vol. 37, no. 3, pp. 233–260, 2007.

[11] D. Balsamo *et al.*, "Graceful performance modulation for power-neutral transient computing systems," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 35, no. 5, pp. 738–749, May 2016.

[12] K. Beason. (2010) smallpt: Global illumination in 99 lines of C. [Online]. Available: http://www.kevinbeason.com/smallpt/

[13] S. Holmbacka *et al.*, "Accurate energy modeling for many-core static schedules with streaming applications," *Microprocessors and Microsystems*, vol. 43, pp. 14–25, Jun 2016.

[14] E. Ahmad and B. Shihada, "Green smartphone GPUs: Optimizing energy consumption using GPUFreq scaling governors," in *Int. Conf. on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct 2015, pp. 740–747.