

SolarCore: Solar Energy Driven Multi-core Architecture Power Management

Chao Li, Wangyuan Zhang, Chang-Burm Cho and Tao Li

Intelligent Design of Efficient Architectures Laboratory (IDEAL)
Department of Electrical and Computer Engineering, University of Florida
{chaol, zhangwy, chorenol}@ufl.edu, taoli@ece.ufl.edu

Abstract — *The global energy crisis and environmental concerns (e.g. global warming) have driven the IT community into the green computing era. Of clean, renewable energy sources, solar power is the most promising. While efforts have been made to improve the performance-per-watt, conventional architecture power management schemes incur significant solar energy loss since they are largely workload-driven and unaware of the supply-side attributes. Existing solar power harvesting techniques improve the energy utilization but increase the environmental burden and capital investment due to the inclusion of large-scale batteries. Moreover, solar power harvesting itself cannot guarantee high performance without appropriate load adaptation. To this end, we propose SolarCore, a solar energy driven, multi-core architecture power management scheme that combines maximal power provisioning control and workload run-time optimization. Using real-world meteorological data across different geographic sites and seasons, we show that SolarCore is capable of achieving the optimal operation condition (e.g. maximal power point) of solar panels autonomously under various environmental conditions with a high green energy utilization of 82% on average. We propose efficient heuristics for allocating the time varying solar power across multiple cores and our algorithm can further improve the workload performance by 10.8% compared with that of round-robin adaptation, and at least 43% compared with that of conventional fixed-power budget control. This paper makes the first step on maximally reducing the carbon footprint of computing systems through the usage of renewable energy sources. We expect that the novel joint optimization techniques proposed in this paper will contribute to building a truly sustainable, high-performance computing environment.*

1. Introduction

The continuing decline of conventional fossil fuel has resulted in increasing energy costs all around the world. Meanwhile, fossil fuel and its induced greenhouse gas emissions have had profound impact on the environment and the climate of our planet. These trends have led the IT community into the renewable energy driven, green computing era. Among various renewable energy options, photovoltaic (PV) generation is gaining increased importance due to its advantages such as absence of fuel cost, low maintenance, and no noise and wear due to the absence of moving parts. It is expected that PV system installations will provide 15%-20% of the global electricity in 2040 [1]. Today, solar energy is not only being used for powering various low-power electronic devices and household appliances, but it is also drawing increased attentions in performance-focused computing environments such as data centers [2]. Internet giants, such as Google, Microsoft and Yahoo all power part of their data centers using renewable resources such as solar farms. In this paper, we extend the benefits of solar

energy to the high performance computer architecture design area. Specifically, we explore the joint optimization of green energy utilization and workload performance for multi-core processors, which are the mainstream hardware design choice for today's IT industries and demand increasing amounts of power to unleash their full computation potential.

Although solar energy is an abundant energy source, its efficient utilization presents a new challenge. In PV systems, sunlight is converted into DC electricity using PV cells. Implemented as a p-n junction fabricated in a thin wafer, PV cells have a non-linear relationship between their output photocurrent and terminal voltage. Under uniform irradiance (solar energy per unit area of the solar panel's surface) and temperature, photovoltaic arrays exhibit a current-voltage characteristic with a unique point, called the maximum power point (MPP) [3, 4], where the module produces maximum output power. Due to the high fabrication cost (e.g. material, ingot, and wafering) and low conversion efficiency (typically 13% ~19%) of present solar cells, it is crucial to operate the PV generator at the maximum power point in order to achieve the maximum efficiency. Nevertheless, the power generated by PV systems changes over time due to the unpredictable nature of weather pattern, as shown in Figure 1. The unstable working environment along with the existence of optimum operating points necessitate the rethinking of multi-core power management policy to achieve higher efficiency in utilizing solar energy. Moreover, the processor load should be regulated and tuned in such a way that the overall workload performance is also optimized during the successive MPP tracking.

There are various ways to connect PV array to the load, as shown in Figure 2. For instance, solar-powered systems can employ energy storage elements (i.e. batteries or super-capacitors) [5] and charge controller (Figure 2-C) [5] to backup energy which reduces voltage fluctuations and performs maximum power tracking. However, energy storage elements introduce many limitations in PV system utilization. For example, the maximal energy that can be delivered is limited by the battery capacity and this will unavoidably affect the performance of multi-core systems. The large current drawn by multi-core chips requires battery with large capacity, which is both bulky and expensive. In addition, the turn-around efficiency of battery is low due to internal resistance and self-discharge. These de-rating factors can contribute to significant loss in energy efficiency (up to 30% energy loss, detailed in Section 5). Last but not least, existing rechargeable batteries all have limited lifetime. Frequent charge/discharge (in case of using low-capacity batteries) and self-discharge further aggregate the aging effect. Without proper maintenance (which requires human intervention and causes loss of availability), the aging problem will directly lead to the capacity reduction and output voltage change. As a result, over the lifetime of the




Figure 1: Solar energy utilization efficiency varies significantly for fixed load under different solar irradiance intensity

solar-powered systems, batteries cost (e.g. initial and maintenance cost) can be the most expensive component of the renewable energy systems [6, 7].

In this study, we consider cost-effective direct-coupled PV system design (Figure 2-B), which has been used in applications such as heating and water pumping. For direct-coupled PV system, the solar-panel load is DC load and no power inverter will be used. In our design, when the solar power supply drops below a certain threshold, a secondary power supply (e.g. grid utilities) will be switched in (via power switch) and used as a power supply until sufficient solar power is available. Different from the conventional direct-coupled PV systems, we propose techniques that allow multi-core processors to autonomously harvest the renewable energy supply to the maximal degree. As a result, our design eliminates the limitations (i.e. low performance, short-lifetime and high-cost) due to inclusion of batteries, while delivering satisfactory green energy utilization and high workload performance. Note that the purpose of our proposal is not to completely replace existing power supply. The key benefit of using *SolarCore* is that it maximally reduces the power dependence on the fossil-fuel-based supply while simultaneously optimizing the workload performance.

This paper makes the following contributions:

- We perform detailed characterization of PV cell/module under varying atmospheric conditions. We show that unlike conventional power sources, matching the power consumption level with the supply level can make a significant difference in solar energy utilization.
- We propose *SolarCore*, the multi-core architecture power management schemes that address a two-fold challenge. The first is to maximize multi-core processor’s total solar energy utilization by performing load matching under variable PV power output. The second is to intelligently allocate the dynamically varied power budget across multiple cores so that the maximum workload performance can be achieved.
- We perform extensive simulations to evaluate the effectiveness of the proposed techniques. We characterize the MPP tracking efficiency and workload behavior using real-world solar irradiance data across different geographic locations and seasons.
- We characterize the green energy utilization and multi-core performance of various solar energy based power management schemes. We demonstrate that *SolarCore* is able to extract 82% solar energy on average without relying on storage elements. Compared with the most efficient battery-based MPPT systems, *SolarCore* yields less than 1% performance degradation and overcomes the




Figure 2: Three typical PV systems: (A) grid-connected, (B) direct-coupled, and (C) battery-equipped

major drawbacks such as short life time, high capital cost and negative environment impact.

The rest of this paper is organized as follows: Section 2 describes circuit model, I-V, and power characteristics of PV cell. Section 3 presents a characterization of solar array I-V and power under varying atmospheric conditions. Section 4 proposes solar energy driven multi-core power management techniques. Section 5 describes the experimental methodologies. Section 6 evaluates the power and performance impact of the proposed schemes. Section 7 discusses related work and Section 8 concludes the paper.

2. The Power Behavior of PV Systems

A photovoltaic (PV) cell (a.k.a. solar cell) is a p-n junction fabricated in a thin wafer or layer of semiconductor. It is the basic building block of PV systems, where the electromagnetic radiation of sun can be directly converted to electricity through the photovoltaic effect. In this section, we describe solar power characteristics and important factors in designing solar energy driven computing systems.

2.1 PV Power Model

During darkness, the PV cell does nothing; it behaves like a diode, generating a current I_d (a.k.a. dark current) under forward bias. When it is switched on by sufficient sunlight irradiance, the solar cell can generate a voltage. As shown in Figure 3, the simplest equivalent circuit of an ideal solar cell can be modeled by a current source in parallel with a diode. When switched on, the ideal cell produces a photocurrent (I_{ph}) proportional to the light intensity. To obtain a more accurate current-voltage (I-V) characteristic of a solar cell, the series resistance R_s , representing the internal losses due to the current flow, and parallel resistance R_p , representing the leakage current to the ground, can be added to the ideal model. In addition, a second non-ideal diode can be added in parallel to the current source [8]. In this paper, we choose a model of moderate complexity, which captures the solar radiation and temperature impact on the photocurrent (I_{ph}) and saturation reverse current (I_0) of the diode. Our model only considers the series resistance since the impact of shunt resistance (parallel resistance) is negligible.

The electrical behavior of a solar cell is given by its I-V characteristic equations, which relates solar cell parameters and environment conditions to the output current (I_{out}) and voltage (V_{out}). In Figure 3, the output current is a complex function of the sun irradiance (G), the operating temperature (T), and the terminal voltage V_{out} . A detailed derivation of I-V characteristics can be found in [8, 9].




Figure 3: The equivalent circuit of a solar cell

2.2 I-V Characteristics and MPP

Figure 4 shows a conceptual plot of solar cell I-V and P-V curves, assuming constant sun irradiance and temperature. When the cell is operated at open circuit, $I_{out} = 0$ and the voltage across the output terminals is defined as the open circuit voltage (V_{oc}). Similarly, when the cell is operated at short circuit, $V_{out} = 0$ and the current through the terminals is defined as the short circuit current (I_{sc}). For a high-quality solar cell, R_s and I_0 are low and R_p is high, therefore $I_{sc} \approx I_{ph}$. For a given load R_L , the cell develops a terminal voltage between $0 \sim V_{oc}$. Assuming there is no energy loss between the PV output power and the load input power, the operating point is the intersection between the I-V curves of the load and the PV generator.

For a given I-V curve under the constant ambient irradiation and temperature, there is a single operating point (MPP) where the values of the voltage (V_{mpp}) and current (I_{mpp}) of the cell result in the maximum power output (P_{max}), which is the optimal operation point for the efficient use of the solar energy. As shown in Figure 4, a load connected PV generator may operate (as determined by the intersection of the I-V curve and the load line) far from its optimal condition (i.e. MPP or P_{max}), resulting in inefficient solar power utilization. Moreover, since the maximum power point varies with the solar irradiance and temperature, the optimal operating point of the photovoltaic cell changes dynamically.

2.3 PV Systems and Load Matching

When connected to electrical loads, solar panels generally employ power converters to achieve appropriate output voltage. Such power converters can be implemented by using controllable transformers [10], such as a Pulse Width Modulation (PWM) based transformer. Assuming $P_{in} = P_{out}$, the transformers can be described as $V_{out} = V_{in} / k$ and $I_{out} = k \cdot I_{in}$, where k is the transfer ratio that can be adjusted by changing the duty cycle of the PWM. The actual operating point of the PV system occurs at the intersection of the electrical characteristics of the solar panel and the load. By load matching, maximum power point tracker (MPPT) is able to extract the optimal power from the panel under varying atmospheric conditions. Such load matching can be achieved by either adjusting the DC-DC converter transfer ratio, k , or tuning multi-core load, w (e.g. via DVFS). For a multi-core processor, its electrical characteristics are largely dependent on parameters such as clock frequency and supply voltage. For




Figure 4: The I-V (a) and P-V (b) characteristics of a PV cell, operating at specific resistive load and MPP


instance, as the clock frequency increases, the multi-core processor will exhibit lower impedance and draw more current and power.

Figure 5 shows how the operating point changes if we tune the two system parameters, namely, the multi-core load w or transfer ratio k . The output voltage will decrease when the load line moves counterclockwise and increase when the load line moves clockwise. However, due to the existence of maximum power point, the actual output power may either increase or decrease as the load line moves, depending on whether the operating point is approaching the MPP or not. Note that simply increasing the multi-core processor utilization rate and voltage level will not automatically draw maximal power because the load adaptation scheme (i.e. increasing or decreasing the load) varies with different operating point positions, as shown in Figure 5 (a) and (b). On the other hand, tuning the transfer ratio alone cannot increase the multi-core performance either due to the lack of effective multi-core load adaptation schemes. Moreover, without appropriate coordination between the power converter and load adaptation, the system will suffer from severe voltage fluctuations.

3. Characterizing the I-V and Power Behavior of PV Modules

Solar cells are the basic building blocks of solar photovoltaic systems. When charged by the sun, each cell can generate approximately 0.5 to 1V photo-voltage and tens of milliamps short current. PV modules consist of multiple interconnected PV cells which are arranged in series-parallel structure to achieve desired power output. PV modules can be further combined to form PV array. In this study, we model the BP3180N polycrystalline 180W solar module [11].

We built the PV cell/array equivalent circuit models described in Section 2 and simulated them using SPICE. We modeled both current versus voltage and power versus voltage characteristics of the studied solar panel at various insolation and temperature conditions. As shown in Figures 6 and 7, higher sun irradiance generates more photocurrent and the MPPs consequently move upward. When the environment temperature rises, the open circuit voltage is reduced and the short circuit current increases. Since the voltage decrease is faster than the current increase, MPP shifts left and the total power that a PV module can produce is reduced at a higher temperature. As can be seen, unlike conventional energy sources, the solar power produced by the PV generator varies significantly under different atmospheric conditions. To harvest the maximum energy, it is crucial that the multi-core power management schemes are capable of identifying and tracking the MPP.



(a) Operating point is on the right side of MPP

(b) Operating point is on the left side of MPP

Figure 5: The characteristics of load matching (P , V and I are all for DC/DC output values). In scenario (a), increasing load or decreasing k will approach MPP. In scenario (b), decreasing load or increasing k will approach MPP.




Figure 6: I-V and P-V curves for various irradiances ($T = 25^\circ\text{C}$)




Figure 7: I-V and P-V curves for various temperatures ($G = 1000\text{W}/\text{m}^2$)

4. Solar Energy Driven Multi-core Architecture Power Management

In this section, we propose *SolarCore*, solar energy driven multi-core architecture power management that can achieve the maximal solar energy utilization and the optimal workload performance simultaneously. Our design makes the first step on reducing the dependence on the utility power of a high performance computing system. A lower dependence saves more on utility bills and produces a lower carbon footprint, which is more sustainable.

4.1 SolarCore: An Overview

Figure 8 illustrates an overview of the *SolarCore* power control architecture. The system is powered by solar energy with grid utility as backup. An automatic transfer switch (ATS) is employed to seamlessly select between the primary (i.e. the solar panel) and backup power sources and an uninterruptable power supply (UPS) ensures continuous power delivery to the load. An AC/DC converter is used only if the ATS switches to the utility. A tunable power-conservative matching network (i.e. DC/DC converter) is used to convert the PV output voltage to the level (e.g. 12V) that is compatible with existing multi-core processors and systems.

Today's power supply unit (PSU) has multiple output rails [12] which can be leveraged to power different system components with different power supplies (e.g. solar/utilities). In this study, we assume that the processor, which contributes to the significant power dissipation in typical multi-core systems, is the only component powered by the renewable energy while the rest of the system is powered by the energy from the utilities. The output power of the PSU is allocated across all running cores with the objective of maximizing the overall workload performance. At the front end, both load current and voltage are measured via I/V sensors and the

results are fed-back to the *SolarCore* controller, which is responsible for identifying and tracking the maximal power point. The controller adjusts the DC/DC converter through MPPT control signal and communicates with the processor through an adapter, which is responsible for per-core load tuning and workload performance optimization.

We decompose solar energy management control into multi-core aware MPP tracking (detailed in Section 4.2) and dynamic load adaptation (detailed in Section 4.3). In our study, we use per-core dynamic voltage/frequency scaling (DVFS) [13] and per-core power gating (PCPG) for load adaptation on multi-core chips. The purpose of performing load adaptation is to move the solar array operating point close to the MPP under changing atmospheric conditions. To achieve the per-core DVFS, we use an on-chip voltage-regulator module (VRM) [14] for each core to provide the appropriate supply voltage. The corresponding supply voltage is communicated between the *SolarCore* controller, the VRMs, and cores via a number of bits called the Voltage Identification Digital (VID), which has been implemented in commercial microprocessors. For example, the Intel Xeon processor employs a 6-bit VID to specify the input voltage from 0.8375 to 1.6V DC with 32 different voltage steps. Our techniques use optimization methods to maximize workload performance under the variable maximal power budget.

4.2 Multi-core Aware MPP Tracking

Our multi-core aware MPP tracking technique relies on successive tuning of both the DC/DC converter transfer ratio k and the multi-core load w . The *SolarCore* controller aims at coordinating the power supply converter and the multi-core load adaptation to achieve the maximal power drawn. In addition, our technique ensures 1) a correct tracking direction and 2) a stable load operating voltage V_l .

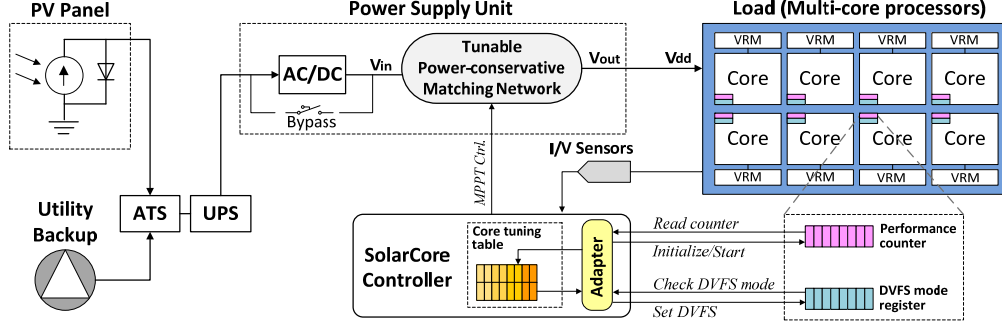


Figure 8: An overview of SolarCore power management architecture

Table 1 summarizes the electrical behavior of load tuning. Tuning k or w can increase output power P_{out} , load voltage V_l and load current I_l , depending on the location of the operating point. By simultaneously adjusting k and w , one can increase the output current while maintaining constant output voltage. As a result, the operating point moves closer to MPP and the multi-core processor can utilize more power. Based on the above observations, we propose a three-step control strategy to perform MPP tracking, as shown in Figure 9.

Step 1. The algorithm starts with a normal operating voltage (i.e. $V_l = V_{dd}$). Due to supply variation, V_l may not equal to V_{dd} at each beginning phase of periodically triggered MPP tracking. In this step, the *SolarCore* controller will restore the output voltage to V_{dd} by decreasing or increasing the load to an appropriate level. This step avoids system overloading and serves as preparation for the following stepwise tuning.

Step 2. As Table 1 shows, the transfer ratio tuning direction that approaches MPP depends on the system's current operating position. To determine the tuning direction, our techniques set the transfer ratio of the DC/DC converter from k to $k + \Delta k$ and observe the current. An increase in output current suggests that the panel generates more power and the system is approaching to the MPP. In this case, the actual operating point is on the left side of the MPP (Figure 5-b) and our algorithm proceeds to step (3) for load matching.

On the other hand, a decrease in output current suggests a wrong tuning direction. In this case, we further decrease the transfer ratio by $2\Delta k$, which results in a net change of $-\Delta k$ in transfer ratio. Consequently, our techniques resume the correct tuning direction and proceed to step (3) to perform load matching.

Step 3. In this step, we tune the load until V_l equals V_{dd} . Due to the adjustment of k in step (2), the output voltage changes as well. By increasing the load, we can change the PV operating point and decrease V_l until it reaches V_{dd} .

During each control period, our algorithm increases the load successively with the aid of tuning transfer ratio k , as discussed in steps (2) and (3). The load adaptation scheme used in steps (1) and (3) will be discussed in detail in Section 4.3. Through the stepwise tuning of transfer ratio and load matching, our methods progressively adapt multi-core power consumption and eventually reach to the new MPP. The goal of *SolarCore* power management is appropriate coordination of the variable power supply and demand. Our MPPT-aware power control is orthogonal to the underlying processor microarchitecture (e.g. out-of-order vs. in-order) and core-count (single-core/SMT vs. multi-core). For example, in the light of single-core processor, the chip-level dynamic voltage/frequency scaling (DVFS) or microarchitecture resource adaptation can be used for load tuning.

Table 1: The characteristics of load tuning (P_{out} , V_l and I_l are DC/DC output)

Position	Operation	Power	Load voltage	Load current
Left of MPP	$k \uparrow$ (const. load)	$P_{out} \uparrow$	$V_l \uparrow$	$I_l \uparrow$
	$w \uparrow$ (const. transfer ratio)	$P_{out} \downarrow$	$V_l \downarrow$	$I_l \uparrow$
	$k \uparrow + w \uparrow$	$P_{out} \uparrow$	$\Delta V_l \approx 0$	$I_l \uparrow$
Right of MPP	$k \downarrow$ (const. load)	$P_{out} \uparrow$	$V_l \uparrow$	$I_l \uparrow$
	$w \uparrow$ (const. transfer ratio)	$P_{out} \uparrow$	$V_l \downarrow$	$I_l \uparrow$
	$k \downarrow + w \uparrow$	$P_{out} \uparrow$	$\Delta V_l \approx 0$	$I_l \uparrow$

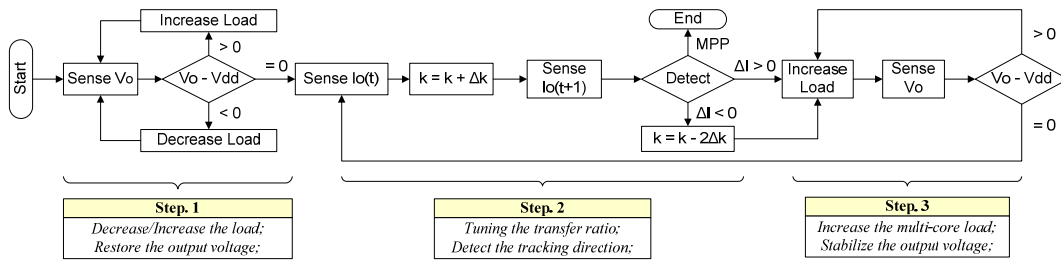


Figure 9: The flowchart of SolarCore MPP tracking

4.3 Dynamic Per-Core Load Tuning

In this work, we explore using DVFS to perform successive load tuning. To be more specific, we apply per-core DVFS to manage multi-core power at a fine granularity and achieve more effective maximum power tracking. There has been prior work [15] that formulates DVFS power control as linear programming optimization problem and using online linear programming to optimize power and performance trade-offs has been proven to be efficient.

The processor dynamic power model we assume in this study is $P = \alpha CV^2 f$ [16]. We further assume that: (1) the voltage scales approximate linearly with frequencies within typical operating ranges, i.e. $f_i = \mu V_i + \lambda$, $V_i \in \{V_1, V_2, \dots, V_n\}$, $i \in \{1, 2, \dots, N\}$, where n is the number of voltage levels and N is the total number of cores; (2) the total power drawn by a core is approximately a cubic function of its voltage level, i.e. $P_i = \alpha CV_i^2(\mu V_i + \lambda) \approx a_i V_i^3 + c_i$; (3) the voltage scaling has little impact on IPC and the throughput $T_i = IPC_i \times f_i$ of a core is proportional to its frequency; and (4) the available power for an individual core does not exceed its maximum power and all cores are homogeneous. Given the above assumptions, the per-core throughput can be expressed as $T_i = IPC_i \times (\mu V_i + \lambda) = b_i V_i + d_i$. The average throughput across all cores is $\bar{T}_i = \frac{1}{N} \sum (b_i V_i + d_i)$. Since the system working duration, which is subject to the solar power generation level, plays a key role in determining the overall performance, our goal is to maximize $\bar{T}_i \times \text{Runtime} / \text{day}$. This performance-time product (PTP) can be measured as the total instructions committed per day. Intuitively, to maximize the PTP, one needs to increase both the throughput and the system operation duration.

In contrast to conventional power management problem, the solar power budget is not fixed, i.e. $P_{total} = \sum P_i = \sum (a_i V_i^3 + c_i) \leq P_{MPP} = f(I, T)$, where I and T is the sun irradiance and environment temperature, respectively. The unpredictable nature of PV output makes performance optimization much more challenging. In this paper, we propose *throughput-power ratio (TPR)* optimization, a cost-

effective PTP optimization method that can be performed along with load adaptation. The throughput-power ratio specifies the throughput speedup of a processor when additional power is available. When adjusting the core voltage by ΔV , a power consumption change of $\Delta P = 3a_i V_i^2 \Delta V$ will occur. The relationship between throughput change and voltage scaling is given by $\Delta T = b_i \Delta V$. Hence the throughput power ratio can be defined as $\frac{\Delta T}{\Delta P} = \frac{b_i}{3a_i V_i^2}$. Both b_i and a_i can be derived using profiling information obtained from the performance counters and the current/voltage sensors at each beginning of MPP tracking. Given the throughput power ratio, our scheme identifies the appropriate core voltage setting through a heuristic process. The cores that exhibit large *TPR* will have higher priority to receive the available power.

We store the core ID and per-core voltage level in a table sorted by the throughput-power ratio, as shown in Figure 10. As described in Section 4.2, the MPPT algorithm either increases or decreases the load, depending on the output characteristics of PV panel. Each load tuning is achieved by increasing or decreasing the voltage level and frequency of the selected core. When we increase the load, we choose the core with higher throughput-power ratio, which will maximize the performance by using additional power. When the load needs to be decreased, our scheme chooses the core with lower throughput-power ratio. Doing so decreases the multi-core power consumption to meet the power budget while minimizing the impact on the overall system performance. The above step is performed iteratively until the aggregated multi-core power approximates the new budget.

Figure 11 shows the per-core load adaptation scenario. The optimum load is achieved when the inflection point (the point where the slope of power trace changes sign) is reached. Since the inflection point can be stable operation states or transition states (i.e. depending on whether the inflection point is under the power budgets), the MPPT controller should not stop tuning immediately once an inflection point is met. For example, if the inflection point is above the power budget, the MPPT controller needs to decrease the load further. As a result, the system returns back to the stable load level and the optimum power consumption is met. Note that the load




Figure 10: Per-core load adaptation policies




Figure 11: Per-core load adaptation scenario

```

// Eight cores are indexed in Core_ID[]
1. Update a_i, b_i for each core;
2. Quick-Sort Core_ID[] based on the throughput-power ratio of each core;
3. Restore front pointer to Core_ID[0]; restore rear pointer to Core_ID[7];
4. if (additional power available == TRUE) { // We have additional power-headroom
5.     if (previous operation is to decrease the workload)
6.         break; // An optimum operating point is obtained
7.     else { // We still have to increase load successively
8.         while (Core_ID[front] operates in its highest voltage level)
9.             front++;
10.            Increases the voltage level of Core_ID[front];
11.        }
12.    }
13. if (power budget decrease == TRUE) { // We have to decrease the workload
14.     if (previous operation is to increase the workload) {
15.         // We are now above the power budget
16.         decrease the voltage level of Core_ID[rear];
17.         break;
18.     }
19.     else {
20.         while (Core_ID[rear] operates in its lowest voltage level)
21.             rear--;
22.         Decrease the voltage level of Core_ID[rear];
23.     }

```

Figure 12: Pseudo code for per-core load tuning

adaptation cannot guarantee perfect utilization of the maximum power. On the other hand, the existence of a power margin is necessary since it improves the robustness of the system. Figure 12 shows the pseudo code of our per-core load-tuning algorithm. The processor starts tuning its load when the controller detects a change in PV power supply. Successive load adaptation is performed at the beginning of each triggered tracking period.

Note that our run-time load adaptation optimization applies to other components although the tuning knobs may vary. For example, the power tracking technique can be combined with dynamic disk speed control [17] to adjust the storage current drawn so that the optimal power consumption can be achieved. Exploring these alternatives and opportunities is beyond the scope of this paper.

5. Experimental Methodologies

We evaluated our design through heavily extended simulators and real-world solar traces. Our PV power model is built using the SPICE-based equivalent circuit simulations discussed in Section 2. The model takes irradiance and temperature profiles as input and generates the I-V and power profiles. We validated the correctness of maximal power point tracking algorithm using MATLAB and Simulink tool box before incorporating it into our architecture simulator.

We used meteorological data from the Measurement and Instrumentation Data Center (MIDC) [18] of National Renewable Energy Laboratory. The MIDC provides real-time records from measurement stations located throughout the U.S. Those records cover a variety of meteorological data (e.g. irradiances, temperature, and wind speed) and have satisfactory accuracy. We chose locations that have different solar energy resource potentials, as shown in Table 2. For each location, we evaluated our design with daytime irradiance data (i.e. 7:30am~5:30pm) across different seasons (e.g. the middle of Jan., Apr., Jul. and Oct. in the year of 2009). We trigger the MPP tracking at each 10-minute interval. We observed that the tracking duration within each interval is less than 5ms. We applied sampling techniques in our simulations to capture the irradiance changes at a large time scale.

We also compared the total solar energy drawn by *SolarCore* with that of battery-equipped standalone PV system. We assume that the battery is optimally charged using MPPT circuit and the multi-core processor runs with full

speed using stable power supply. We assume that dynamic power monitor is used to ensure that all the solar energy stored in the battery is consumed in our simulation. We use de-rating factors [20] to estimate the upper bound of the utilization on conventional battery-equipped solar power systems. Typically, battery loss ranges from 5% to 25% [21] and a modern maximum power tracking controller has conversion efficiency between 93% ~ 97% [22]. We consider three levels of performance ranges, as shown in Table 3.

We simulated a multi-core system comprised of 8 Alpha 21264 processors, where each core has private L1 and L2 caches. We assume a frequency of 2.5GHz and a 90nm technology with a maximum supply voltage of 1.45V. Table 4 summarizes our core configuration. We used a cycle-accurate multi-core simulator integrated with the modified power models from Wattch [23] and CACTI [24]. Both dynamic and leakage power are modeled. We assume that the dynamic voltage and frequency scaling schemes employed in each core are similar to Intel’s SpeedStep techniques [25]: each core supports 6 frequency and voltage operating points and the frequency can be scaled from 2.5GHz down to 1GHz with a 300MHz stepping. The voltage scales approximate linearly with frequency within typical operating ranges; six voltages ranging from 1.45V to 0.95V with a 0.1V stepping are used. To support DVFS, the Wattch and CACTI power models are extended to use the voltage and frequency as input and then adjust the power estimation for a core when the operating status changes. For the performance analysis, we used the performance-time products rather than IPC since the frequency varies across the entire program execution.

We use workloads from the SPEC2000 benchmark suite shown in Table 5. All workloads were simulated using reference inputs. To form multi-programming workloads, we use a program’s average energy-per-instruction (EPI) and categorize the benchmarks as being high-EPI ($EPI \geq 15nJ$), moderate-EPI ($15nJ \geq EPI \geq 8nJ$), or low-EPI ($EPI \leq 8nJ$). We run each benchmark in their representative execution intervals and the EPI is obtained by calculating the average-energy consumed per-instruction. In Table 5, the High-, Moderate- and Low- EPI workloads (H1, H2; M1, M2; L1, L2) consist of 8 benchmarks from each category. The High-Moderate- and Moderate-Low-EPI workloads are formed using four benchmarks from each category. We generated both homogeneous and heterogeneous workloads.

Table 2: The evaluated geographic locations [18, 19]

Station	Location	KWh/m ² /day on average	Solar energy potential
PFCI	Phoenix, AZ	> 6.0	Excellent
BMS	Golden, CO	5.0 ~ 6.0	Good
ECSU	Elizabeth City, NC	4.0 ~ 5.0	Moderate
ORNL	Oak Ridge, TN	< 4.0	Low

Table 3: Performance levels of battery-based PV systems

Levels	High	Moderate (typical)	Low
MPP Tracking efficiency	97%	95%	93%
Battery round-trip efficiency	95%	85%	75%
Overall de-rating factors	92%	81%	70%
Efficiency range	81%~92%	70%~81%	< 70%

Table 4: Simulated machine configuration

Parameter	Configuration	Parameter	Configuration
Frequency	2.5/2.2/1.9/1.6/1.3/1.0GHz	ROB Size	98 entries
Voltage	1.45/1.35/1.25/1.15/1.05/0.95V	LSQ	48 entries
Width	4-wide fetch/issue/commit	Integer ALU	4 I-ALU, 2 I-MUL/DIV, 2 Load/Store
Issue Queue	64	FP ALU	2 FP-ALU, 2 FP-MUL/DIV/SQRT
ITLB	128 entries, 4-way, 200 cycle miss	DTLB	256 entries, 4-way, 200 cycle miss
Branch Predictor	2K entries Gshare 10-bit global history	L1 I-D- Cache	64KB, 4-way, 64 B/line, 2 ports, 3 cycle access
BTB	2K entries, 4-way	L2 Cache	2MB, 8-way, 128 B/line, 12 cycle access
Return Address Stack	32 entries RAS	Memory Access	64 bit wide, 400 cycles access latency

Table 5: The simulated multi-programmed workloads

EPI	High	H1(<i>art</i> × 8)	H2 (<i>art</i> × 2, <i>apsi</i> × 2, <i>bzip</i> × 2, <i>gzip</i> × 2)
	Moderate	M1(<i>gcc</i> × 8)	M2 (<i>gcc</i> × 2, <i>mcfl</i> × 2, <i>gap</i> × 2, <i>vpr</i> × 2)
	Low	L1(<i>mesa</i> × 8)	L2 (<i>mesa</i> × 2, <i>equake</i> × 2, <i>lucas</i> × 2, <i>swim</i> × 2)
	High-Moderate	HM1(<i>bzip</i> × 4, <i>gcc</i> × 4)	HM2 (<i>bzip</i> , <i>gzip</i> , <i>art</i> , <i>apsi</i> , <i>gcc</i> , <i>mcfl</i> , <i>gap</i> , <i>vpr</i>)
	Moderate-Low	ML1(<i>gcc</i> × 4, <i>mesa</i> × 4)	ML2 (<i>gcc</i> , <i>mcfl</i> , <i>gap</i> , <i>vpr</i> , <i>mesa</i> , <i>equake</i> , <i>lucas</i> , <i>swim</i>)

Table 6: The evaluated power management schemes

Algorithm	MPPT	Load adaptation	Scheduling method
<i>Fixed-Power</i>	No	DVFS	Linear programming optimization with a fixed power budget
<i>MPPT & IC</i>	Yes	DVFS	Tuning individual core until reaching its highest/lowest V/F level
<i>MPPT & RR</i>	Yes	DVFS	Round-robin scheduling
<i>MPPT & Opt</i>	Yes	DVFS	Optimized scheduling based on throughput-power ratio

6. Evaluation Results

We compared the efficiency of *SolarCore* with various power management policies, as summarized in Table 6. The *Fixed-Power* is a non-tracking power management scheme which assumes a constant power budget during the entire workload execution. In contrast to *SolarCore*, *MPPT&IC* and *MPPT&RR* both apply tracking control techniques but employ different scheduling policies to perform load adaptation. To be more specific, *MPPT&IC* keeps tuning individual core until reaching its highest or lowest V/F level while *MPPT&RR* tends to distribute the power budget variation evenly across all the cores in a round-robin fashion. The *MPPT&Opt* (i.e. our default configuration for *SolarCore*) selects cores using the proposed *throughput-power ratio* optimization method.

Since we are primarily interested in green energy powered duration, the operation duration we refer to in this paper will always be the daytime duration. In addition, the effective operation duration means the duration that electrical load successfully draws power from the solar panel. During effective operation duration, *SolarCore* performs maximal power point tracking and workload run-time optimization. Due to the intermittency of renewable resource, effective operation duration will be shorter than the operation duration.



6.1 MPP Tracking Accuracy

We first examined the MPP tracking accuracy across different workloads. Figures 13 and 14 provide a graphic view of the tracking results. In each Figure, we plotted the maximum power trace and the actual power extracted using MPP tracking. We present results for regular (e.g. Jan @AZ)

and irregular (e.g. July @AZ) weather patterns. As can be seen, MPP tracking technique ensures that actual power consumption closely follows the ideal maximal power.

For high EPI workloads (e.g. H1), large ripples in power tracking are generated due to the high variation in load power. Power ripples, together with the unpredictable nature of the environment, affect the use of renewable energy and the system reliability. Low EPI workloads and heterogeneous workloads are observed to have small ripples, as can be seen in Figures 13(b)-(c) and 14(b)-(c). The former manifests small power ripples since they have lower average energy per instruction, which results in relatively small power fluctuation amplitude; the latter exhibits small ripples because a mix of diverse workloads can smooth the power fluctuation and relatively steady power consumption is obtained.

Our dynamic load adaptation scheme reduces the impact of load power ripples by leaving a small power margin between the actual load power consumption and the solar power budget. This power margin slightly degrades the tracking accuracy but improves the reliability. To quantify the impact of power margin on the tracking accuracy, we calculated the relative tracking error. In each tracking period t , the relative tracking error is defined as $\bar{\eta} = |P - P'|/P'$, where P is the actual power consumption and P' is the power budget. Table 7 shows the geometric mean of the errors on each geographic location across different weather patterns. Due to the reduced power margin, the low EPI workloads (e.g. L1) show relatively small tracking errors. Similarly, compared with H1, heterogeneous workloads (e.g. HM2) also have small tracking error as well.

**Figure 13: MPP tracking accuracy across different workloads (Jan. @ AZ)****Figure 14: MPP tracking accuracy across different workloads (Jul. @ AZ)****Table 7: The average relative error**

		Homogeneous			Less homogeneous		Less heterogeneous			Heterogeneous	
		H1	M1	L1	HM1	ML1	H2	M2	L2	HM2	ML2
AZ	Jan	10.0%	9.2%	7.7%	8.9%	7.1%	8.5%	7.4%	8.4%	6.9%	7.6%
	Apr	12.6%	9.2%	8.3%	10.6%	9.3%	10.7%	9.3%	9.0%	10.2%	8.3%
	Jul	10.0%	7.3%	7.7%	9.0%	7.6%	9.6%	7.1%	7.0%	7.9%	6.8%
	Oct	9.8%	7.8%	7.7%	8.6%	8.0%	7.0%	7.6%	7.4%	7.4%	6.6%
CO	Jan	14%	11.3%	10.8%	12.0%	10.6%	11.2%	10.0%	9.2%	9.0%	11.3%
	Apr	13.2%	8.9%	7.8%	10.8%	7.5%	9.7%	8.0%	6.8%	9.8%	8.9%
	Jul	10.1%	8.2%	7.6%	8.9%	7.9%	8.5%	7.6%	6.6%	8.1%	6.5%
	Oct	12.1%	9.2%	7.9%	9.2%	8.1%	9.5%	8.6%	7.8%	8.1%	7.0%
NC	Jan	13.0%	10%	11.4%	12.4%	10.4%	11%	8.9%	10%	13.3%	9.3%
	Apr	22.0%	17%	19%	18.5%	16.4%	21%	14.3%	16.3%	16.3%	13.8%
	Jul	11.6%	8.3%	5.7%	8.6%	5.4%	7%	4.4%	4.0%	4.9%	4.5%
	Oct	15.4%	14.1%	12.0%	14.3%	12.1%	14.2%	12.5%	11.9%	13.3%	11.1%
TN	Jan	10%	8.2%	6.6%	8.0%	6.00%	13.2%	6.1%	7.4%	9.0%	8.7%
	Apr	15.1%	14.4%	15.3%	16%	14%	13.1%	12.9%	13%	14.4%	12.9%
	Jul	13.4%	9.7%	8.5%	9.4%	8.5%	13%	9.0%	8.2%	7.7%	7.6%
	Oct	18.0%	13.8%	14.3%	14.7%	13.3%	12.9%	11.9%	13.4%	12.7%	11.6%

EPI: High \longrightarrow Low

6.2 Fixed Power Budget

For the direct-coupled PV systems, the load starts to operate using solar power when the amount of renewable energy exceeds a power-transfer threshold. We evaluated the solar energy utilization on multi-core systems that use power-transfer threshold as a fixed power budget. This fixed power budget ensures that the multi-core systems operate reliably with sufficient power. If the power supply from renewable power source falls below the power-transfer threshold, the multi-core system will switch to the secondary power source (e.g., grid utility).

We simulated both homogeneous and heterogeneous workloads with high and low EPI. For each weather pattern, we calculated the effective operation duration under different power budget thresholds, as shown in Figure 15. The threshold affects processor performance-time product (PTP) in both throughput and effective operation duration. A higher power-transfer threshold will make the multi-core processor run at higher voltages and frequencies but only for short durations. A conservative, low power-transfer threshold operating mode will have longer duration but lower clock frequencies. Intuitively, the processor should run under higher (lower) power budget to maximize its performance when the effective operation duration declines slowly (rapidly). Extensive simulations shown in Figures 16 and 17 further proof this.

We calculated the accumulated energy drawn by the multi-core processor and the performance in terms of PTP. All the results are normalized to those obtained on *SolarCore*. As

shown in Figures 16 and 17 (average results across all benchmarks), the maximal solar energy drawn does not guarantee maximal performance (e.g. July @CO and Apr @TN). The maximal workload performance may occur under a high power budget (e.g. Apr @AZ), a moderate power budget (e.g. Jan @TN) or even a low power budget (e.g. Apr @NC). Therefore, a single, optimal fixed power budget for the multi-core system does not exist. Even under the optimal fixed power budget, the best energy utilization and PTP that the *Fixed-Power* schemes can achieve is less than 70% of that yielded on maximum power tracking. In other words, *SolarCore* outperforms *Fixed-Power* control scheme by at least 43% (i.e. $(100\% - 70\%)/70\%$) in terms of both energy utilization and workload performance.

6.3 Energy Utilization

The energy utilization of *SolarCore* depends on the solar resource potential of the geographic location. When the PV generation is too low to power the load, *SolarCore* will draw power from the utility via transfer switch and the multi-core processor will act as a traditional CMP. We calculated the energy utilization (i.e. actual total solar energy consumed / theoretical maximum solar energy supply) with various load adaptation scenarios. As shown in Figure 18, the average solar energy utilization drops when the renewable resource potential is low (e.g. TN). For locations with abundant solar resource (e.g. AZ), *SolarCore* draws 5% more power compared to a typical battery equipped PV system which has an energy utilization upper bound of 81%.

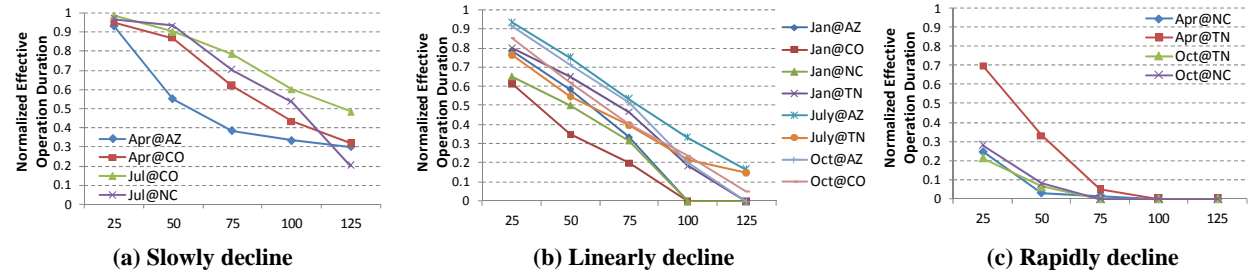


Figure 15: Operation durations on different power budget thresholds (x-axis: watts). The operation duration may decline slowly (a), linearly (b) or rapidly (c).




Figure 16: Normalized solar energy usage under different fixed power budget (x-axis: watts)




Figure 17: Normalized PTP under different fixed power budget (x-axis: watts)




Figure 18: Average energy utilization across different geographical locations (The upper bound of utilization on battery-based systems is estimated using de-rating factors)




Figure 19: Effective operation duration of SolarCore under different weather patterns

The primary reason of low energy utilization in NC and TN is that long effective operation duration cannot be guaranteed. As shown in Figure 19, the effective operation duration ranges from 60% to 90% during the daytime, depending on the weather patterns. The effective operation duration in AZ is always longer and therefore, the average energy utilization is higher. Figure 20 shows the average energy utilization with different effective *SolarCore* operation duration. The solar energy utilization decreases significantly if *SolarCore* has to be powered by the backup supply for a long duration (e.g. Oct @NC and Oct @TN). Given that 80% of the daytime is operated using power tracking, *SolarCore* can guarantee 82% or more solar energy utilization on average.

In addition to the weather patterns, different load adaptation methods and workload characteristics affect the utilization as well. In Figure 18, we observe that *MPPT&Opt* load adaptation method has 2.6% lower energy utilization compared with that of *MPPT&RR*. This is because *MPPT&Opt* relies on successively tuning low-power, high-throughput cores to improve the performance. The optimization method that *MPPT&Opt* used implicitly reduces




Figure 20: Average solar energy utilization under different tracking durations

the overall load energy consumption while improving the overall performance. Compared with the high EPI workloads, low EPI workloads exhibit higher energy utilization due to the fine-grained load power tuning and the reduced power margin. By increasing the granularity of DVFS level, one can increase the control accuracy of MPPT and the power margin can be further decreased without significantly affecting the MPPT robustness.

6.4 Performance Improvement

We compared *SolarCore* with the high efficiency battery-equipped systems. *Battery-L* denotes the lower bound of a high-efficiency battery-equipped system, which has a total energy conversion efficiency of 0.81. *Battery-U* denotes the upper bound of a high-efficiency battery-equipped system, which has a total energy conversion efficiency of 0.92. Figure 21 shows the performance of *SolarCore* with different load scheduling methods across different weather patterns. All the results are normalized to *Battery-L*.

In Figure 21, *MPPT&IC* shows the lowest performance because it concentrates the solar power into fewer cores. The

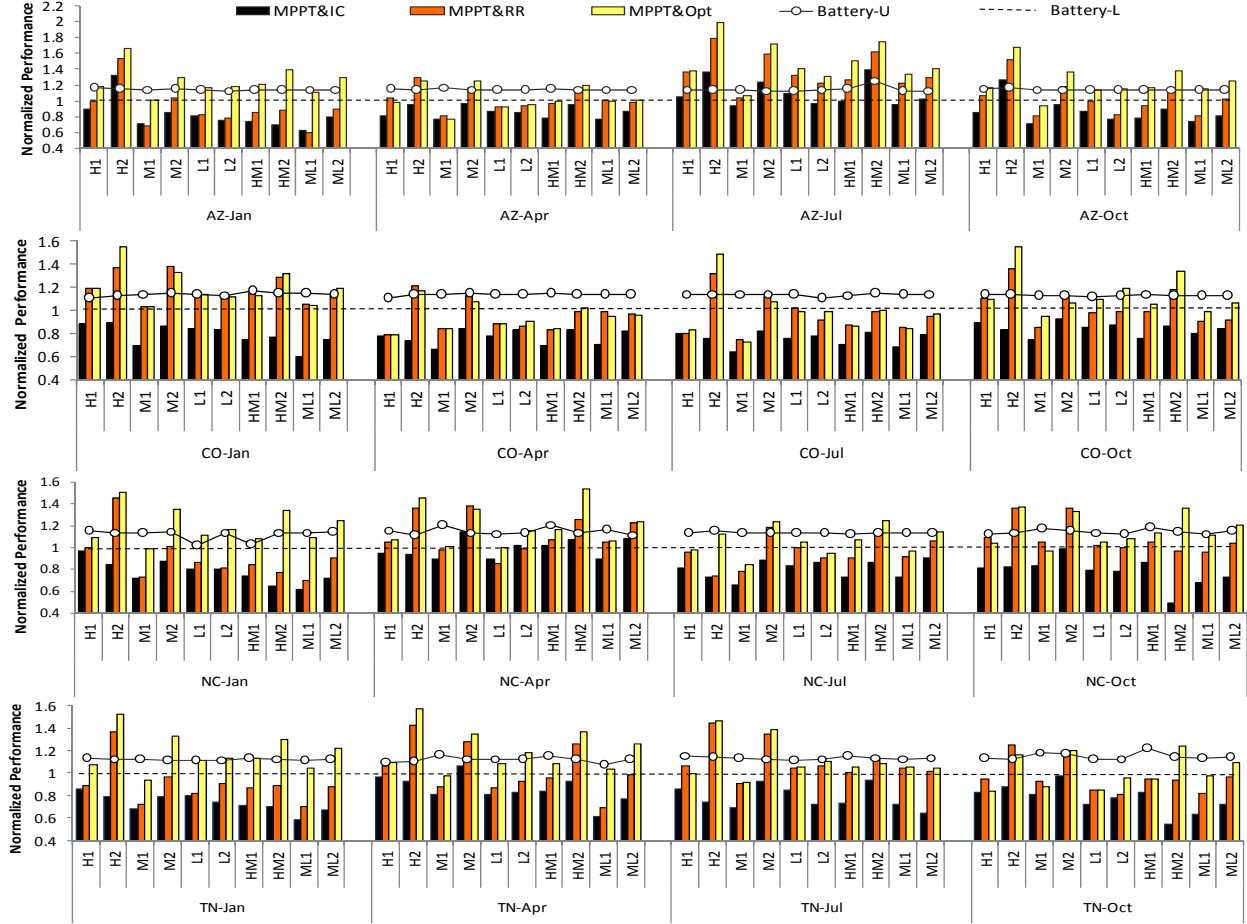


Figure 21: Normalized performance time product (PTP) with different load scheduling methods

throughput-power ratio $\frac{\Delta T}{\Delta P} = \frac{b_i}{3a_i V_i^2}$ shows that the performance return decreases when we allocate the limited available power to those high V/F level cores. *MPPT&RR* increases the performance by distributing the solar power to each individual core in a round-robin fashion. In this case, each core runs in a moderate V/F level with a relatively high performance per watt. The average normalized performance of *MPPT&IC*, *MPPT&RR* and *MPPT&Opt* is 0.82, 1.02 and 1.13, respectively. The normalized performance of *Battery-U* is 1.14. As can be seen, with runtime *throughput-power ratio* optimization, *MPPT&Opt* improves the performance by 37.8% compared with *MPPT&IC* and 10.8% compared with *MPPT&RR*. Compared with the most efficient battery-based MPPT systems, *SolarCore* yields less than 1% performance degradation. Moreover, *SolarCore* is environmental friendly and does not have the life-time problem as the expensive battery-equipped PV systems.

7. Related Work

To fully exploit the energy potential, systems powered by renewable source require power source aware control techniques [26, 27]. Pioneering work is most seen in low-power embedded system area [28, 29, 30]. Li et al. [27] proposed using load matching [10, 31] to exploit the CPU

slack time of a solar-powered wireless handheld system. Clark et al. [30] demonstrated a heuristic control scheme to manage a battery-less solar-powered RFID system. Stewart et al. [26] discussed the intermittency nature of renewable energy and proposed the power transfer system. Different from prior work, we take the first step to design a renewable energy aware power management technique to coordinate power provisioning control and load matching. Our technique enables joint optimization of solar energy utilization and workload performance.

Many algorithms have been proposed for MPP tracking [3, 29, 32]. A detailed comparison between different MPP tracking methods can be found in [33]. Different approaches of modeling a PV panel is discussed in [8, 9]. In this paper, we show that the multi-core system has the potential to operate autonomously without relying on complex control or energy storage elements.

Conventional multi-core power management techniques assume constant power supply and therefore they are largely workload-driven. The goals of these techniques are either reducing power consumption by improving power efficiency or maximizing the performance under limited power supply. For example, [15] proposed a variability-aware DVFS control method using linear programming to optimize average throughput based on given power budget. [13] showed that on

chip regulators have the potential of achieving fast fine-grained DVFS power control. There has been prior work that utilized DVFS to achieve multi-core power management [34, 35, 36, 37, 38]. [34] investigated global power management and evaluated different control policies in light of high performance and power efficiency. [36] proposed a fine-grained power management scheme working on the nanosecond scale. In this study, we exploit DVFS for dynamic load tuning. Different from conventional energy source, solar energy is not a constant power source and the solar power generated can vary significantly with different loads.

8. Conclusions

Computer systems operated using energy generated from conventional fossil fuel increases the total cost of ownership and causes environment impact. Clean and renewable energy sources such as photovoltaic power generation are one of the most promising technologies for the future of green computing in IT infrastructures. Unlike conventional energy sources, PV modules exhibit non-linear I-V characteristics and their output power changes with the variation of light intensity and ambient temperature. Therefore, the best design practice of solar energy driven high performance computing system requires 1) *appropriate power provisioning control to harvest the renewable energy and reduce the impact of supply variation* and 2) *dedicated load adaptation schemes that optimize the workload performance*.

In this paper, we propose *SolarCore*, novel power management schemes for solar energy driven multi-core processors. While existing techniques seek to minimize multi-core power dissipation under performance constraints or to maximize throughput for a given power budget, our techniques harvest the maximum amount of solar energy to maximize the power budget for optimal throughput without using short lifetime and expensive storage elements. Furthermore, *SolarCore* applies load optimization based on the workload throughput-power ratio to ensure that the dynamic load tuning across multiple cores achieves the optimal performance. Because of its ability to extract additional solar energy and its ability for load optimization, *SolarCore* boosts multi-core processor performance by 43% compared with conventional fixed power budget scheme and 10.8% compared with round-robin load adaptation. We expect that the enabling techniques we proposed will open new research opportunities on multi-core power management in light of the clean, renewable energy. As our future work, we will explore the full-system based solar power management techniques for server architecture that includes other hardware components such as memory, disk and network interface.

Acknowledgements

This work is supported in part by NSF grant 0845721(CAREER) and by Microsoft Research Safe and Scalable Multi-core Computing Award. Chao Li is also supported by a University of Florida Graduate Fellowship.

References

- [1] G. Pistoia, Battery Operated Devices and Systems: From Portable Electronics to Industrial Products, *Elsevier*, 2009
- [2] Solar Power at Datacenter Scale: <http://www.datacenterknowledge.com/archives/2009/06/16/solar-power-at-data-center-scale/>
- [3] Maximum Power Point Tracking, National Instrument, <http://zone.ni.com>

- [/devzone/cda/tut/p/id/8106](http://devzone/cda/tut/p/id/8106)
- [4] A. Luque and S. Hegedus, Handbook of Photovoltaic Science and Engineering, *John Wiley & Sons Inc*, 2003
- [5] D. Linden and T. Reddy, Handbook of Batteries, 3rd Ed, *McGraw Hill*, 2002
- [6] A. Jossen, J. Garche and D. Sauer, Operation Condition of Batteries in PV Applications, *Solar Energy*, Vol. 76, 2004
- [7] Battery Power for Your Residential Solar Electric System, White Paper of National Renewable Energy Lab, <http://www.nrel.gov/docs/fy02osti/31689.pdf>
- [8] D. King, J. Dudley and W. Boyson, PVSIM: A Simulation Program for Photovoltaic Cells, Modules, and Arrays, *IEEE Photovoltaic Specialists Conference*, 1996
- [9] D. Sera, R. Teodorescu and P. Rodriguez, PV Panel Model based on Datasheet Values, *IEEE International Symposium on Industrial Electronics*, 2007
- [10] D. Shmilovitz, On The Control of Photovoltaic Maximum Power Point Tracker via Output Parameters, *IEEE Proceedings on Electric Power Applications*, Vol. 152, 2005
- [11] BP3180N datasheet: <http://www.solardepot.com/pdf/BP3180N.pdf>
- [12] ATX12V Power Supply Design Guide v2.01
- [13] W. Kim, M. Gupta, G. Wei and D. Brooks, System Level Analysis of Fast, Per-Core DVFS using On-Chip Switching Regulators, *HPCA 2008*
- [14] VRM Design: <http://www.intel.com/design/pentium4/guides/249205.htm>
- [15] R. Teodorescu and J. Torrellas, Variation-Aware Application Scheduling and Power Management for Chip Multiprocessors, *ISCA 2008*
- [16] S. Kaxiras and M. Martonosi, Computer Architecture Techniques for Power-Efficiency, *Morgan & Claypool*, 2008
- [17] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir and H. Franke, DRPM: Dynamic Speed Control for Power Management in Server Class Disks, *ISCA 2003*
- [18] Measurement and Instrumentation Data Center, <http://www.nrel.gov/mid/>
- [19] Geographic Information System (GIS) Maps, Power Technologies Energy Data Book, 4th Ed, Technical Report of National Renewable Energy Lab
- [20] Calculating Actual PV System Output: <http://www.xantrex.com/>
- [21] IEEE Guide for Array and Battery Sizing in Stand-Alone Photovoltaic (PV) Systems, IEEE Std 1562™-2007
- [22] MPPT Tracking: <http://www.windsun.com/ChargeControls/MPPT.htm>
- [23] D. Brooks, V. Tiwari and M. Martonosi, Wattech: A Framework for Architectural-level Power Analysis and Optimizations, *ISCA 2000*
- [24] N. Muralimanohar, R. Balasubramanian and N. Jouppi, CACTI 6.0: A Tool to Model Large Caches, Technical Report, <http://www.hpl.hp.com/techreports/2009/HPL-2009-85.html>
- [25] Enhanced Intel SpeedStep Technology for the Intel Pentium M Processor, <ftp://download.intel.com/design/network/papers/30117401.pdf>
- [26] C. Stewart and K. Shen, Some Joules Are More Precious Than Others: Managing Renewable Energy in the Datacenter, *Workshop on Power Aware Computing and Systems, HotPower 2009*
- [27] D. Li and P. Chou, Application/Architecture Power Co-Optimization for Embedded Systems Powered by Renewable Sources, *DAC 2005*
- [28] X. Jiang, J. Polastre and D. Culler, Perpetual Environmentally Powered Sensor Networks, *IPSN 2005*
- [29] J. Taneja, J. Jeong and D. Culler, Design, Modeling, and Capacity Planning for Micro-solar Power Sensor Networks, *IPSN 2008*
- [30] S. Clark, J. Gummeson, K. Fu, D. Ganesan, Towards Autonomously-Powered CRFIDs, *Workshop on Power Aware Computing and Systems, HotPower 2009*
- [31] D. Li and P. Chou, Maximizing Efficiency of Solar-Powered Systems by Load Matching, *ISPLED 2004*
- [32] N. Femia, G. Petrone, G. Spagnuolo and M. Vitelli, Optimization of Perturb and Observe Maximum Power Point Tracking Method, *IEEE Transactions on Power Electronics*, Vol.20, 2005
- [33] T. ESRAM and P. Chapman, Comparison of Photovoltaic Array Maximum Power Point Tracking Techniques, *IEEE Transactions on Energy Conversion*, Vol.22, 2007
- [34] C. Isci, A. Buyuktosunoglu, C. Cher, P. Bose and M. Martonosi, An Analysis of Efficient Multi-Core Global Power Management Policies: Maximizing Performance for a Given Power Budget, *MICRO 2006*
- [35] J. Lee and N. Kim, Optimizing Throughput of Power- and Thermal-Constrained Multicore Processors Using DVFS and Per-Core Power-Gating, *DAC 2009*
- [36] K. Rangan, G. Wei and D. Brooks, Thread Motion: Fine-Grained Power Management for Multi-Core Systems, *ISCA 2009*
- [37] T. Horvath and K. Skadron, Multi-mode Energy Management for Multi-tier Server Clusters, *PACT 2008*
- [38] S. Herbert and D. Marculescu, Variation-Aware Dynamic Voltage /Frequency Scaling, *HPCA 2009*