



Dynamic Energy Burst Scaling for Transiently Powered Systems

Conference Paper**Author(s):**

Gomez, Andres; [Sigrist, Lukas](#) ; Magno, Michele; [Benini, Luca](#) ; Thiele, Lothar

Publication date:

2016

Permanent link:

<https://doi.org/10.3929/ethz-b-000116251>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Funding acknowledgement:

157048 - Transient Computing Systems (SNF)

Dynamic Energy Burst Scaling for Transiently Powered Systems

Andres Gomez*, Lukas Sigrist*, Michele Magno*[†], Luca Benini *[†], Lothar Thiele*

*D-ITET, ETH Zurich, 8092 Zurich, Switzerland

Email: {gomeza, sigristl, magnom, lbenini, thiele}@ethz.ch

[†]DEIS, University of Bologna, Bologna, Italy

Email: firstname.lastname@unibo.it

Abstract—Energy harvesting is generally seen to be the key to power cyber-physical systems in a low-cost, long term, efficient manner. However, harvesting has traditionally been coupled with large energy storage devices to mitigate the effects of the source’s variability. The emerging class of transiently powered systems avoids this issue by performing computation only as a function of the harvested energy, minimizing the obtrusive and expensive storage element. In this work, we present an efficient Energy Management Unit (EMU) to supply generic loads when the average harvested power is much smaller than required for sustained system operation. By building up charge to a pre-defined energy level, the EMU can generate short energy bursts predictably, even under variable harvesting conditions. Furthermore, we propose a dynamic energy burst scaling (DEBS) technique to adjust these bursts to the load’s requirements. Using a simple interface, the load can dynamically configure the EMU to supply small bursts of energy at its optimal power point, independent from the harvester’s operating point. Extensive theoretical and experimental data demonstrate the high energy efficiency of our approach, reaching up to 73.6% even when harvesting only 110 μ W to supply a load of 3.89 mW.

I. INTRODUCTION

Over the past decade, there has been a considerable push to reduce the power consumption of electronic devices. However, the broader problem of how to supply them with the energy they require in an efficient, low-cost, long-term, self-sustainable manner has not yet been adequately solved. Over-provisioning with large energy harvesting and storage elements is not a feasible solution in many application scenarios such as implantable devices or "smart dust" systems.

Transiently powered systems are systems that are supplied by volatile energy sources which can, at most, directly power the system for only a limited amount of time. During this time, the energy harvesting rate is not high enough to complete even one atomic task execution. Therefore, such a system needs to be able to buffer at least the amount of energy needed to bridge this power deficit in order to guarantee the completion of a task. In this application domain, attempting to use a buffer that stores large amounts of energy inevitably leads to high losses due to power harvesting costs, self-discharge, and converter inefficiencies. Devices such as batteries and supercapacitors are particularly unsuited because they are expensive in terms of cost and area, have limited charge cycles and high self-discharge rates, and are not easily integrated on board [1].

Typical low-power cyber-physical systems have components such as microcontrollers, memories, and peripherals, for example, sensors and transceivers. Microcontrollers usually have a wide operating voltage range, but on-chip converters

operate most efficiently at lower supply voltages [2]. External peripherals such as sensors and radios can have substantially different voltage requirements, but system designers usually avoid multiple voltage domains to reduce converter losses and simply choose the highest minimum voltage required to supply the entire system. In order to design a flexible platform that is able to efficiently harvest energy from different sources, it is necessary to decouple the source and load voltages, allowing each to operate at their respective optimal power point. In this work, we argue that transiently powered systems should take all of these issues into consideration. More precisely, we believe these systems should have the following properties in order to be considered useful and efficient:

- 1) Source and load power points are decoupled.
- 2) System buffers minimal energy.
- 3) Load receives the energy required for task completion.

The first property ensures functionality and maximum power point tracking [3] for a wide range of power and voltage inputs. The second limits the energy that the system can buffer to the absolute minimum, since anything more leads to unnecessary losses. We define this minimum to be the energy known to be required for the execution of one task. If an application consists of several tasks, the maximum energy level allowed corresponds to the task with the highest energy requirement. The third property implies that when the load is activated, its minimum energy requirement can be guaranteed.

This paper presents an Energy Management Unit (EMU) which allows a system with limited energy buffering to operate predictably and efficiently, even under very lower power harvesting conditions. Existing works [4]–[7] have looked at low power systems with energy harvesting and storage capabilities. However, these systems are extremely expensive in terms of harvesting and storage requirements for long-term, efficient functionality under transient power conditions. Our proposed EMU has an optimally sized capacitor which minimizes the required start-up time and energy from zero, while maintaining a low cost, small form factor, high efficiency and virtually unlimited charge cycles. Furthermore, we propose the novel concept of Dynamic Energy Burst Scaling (DEBS) to track the load’s optimal power point and minimize its energy. We summarize the main contributions of this work as follows:

- Energy Management Unit that efficiently converts low power levels to short, high power energy bursts.
- Feedback-based Dynamic Energy Burst Scaling technique to track the load’s optimal power point.
- Accurate model to optimize system’s application-specific parameters for low input power scenarios.

II. RELATED WORK

Cyber-physical systems have traditionally been used in conjunction with energy harvesting and energy storing. More recently, the research community has focused on systems with a limited energy storage capacity. Broadly speaking, there are three types of architectures for transient systems:

Directly Coupled: When the energy source has a V-I curve compatible with the load, they can be directly connected. The authors of [8], [9] have proposed a combined hw/sw approach to perform computation when the source can directly sustain the load during short periods of time. These works use volatile logic that requires state-retention mechanisms. In [10]–[13], the authors present storage-less and converter-less harvesting systems in which the load uses frequency scaling to track the maximum power point of the source. While frequency scaling can maximize the energy input, it does not minimize the load’s energy consumption and is limited to a narrow power range. Even though directly-coupled systems typically enjoy a high energy efficiency, if the power input is below this narrow range the load cannot be powered and the system’s efficiency immediately drops to 0%. Unfortunately, this is often the case in typical transiently powered systems. When the energy source and load have incompatible operating points, decoupling them with converters becomes a necessity. As opposed to traditional, battery-based systems, decoupled transient systems have a limited energy buffer between the source and load.

Boost Converter Only: In [14], [15], the authors propose a low-power management system that requires very low input voltage and current. Using a large buffer capacitor at the converter input, they are able to start the energy conversion at very low input power level. However, both approaches suffer from excessively long cold-start times due to charging a large input capacitance, 140 mF, at a constant low input power of 2.5 μ W. As will be explained in Sec. IV-B, our capacitance is chosen to minimize the cold-start energy and time.

Boost Buck Converter Combination: The authors of [16] also use a boost converter for optimal power point tracking. However, their proposed system utilizes RF harvesting to accumulate charge in a supercapacitor and then power a camera application with a buck converter. The boost/buck converter topology with an energy buffer also serves as basis for the approach presented in this work. While a charge-state model is used to characterize the capacitor’s self-discharge rate, energy losses such as impedance matching and converter inefficiencies are neglected. More importantly, the system has a large startup cost and can only supply the load with bursts of a constant size and voltage. In Sec. VI, it will be shown that this approach can lead to a substantially higher energy consumption.

To summarize this work, we propose an Energy Management Unit (EMU) to decouple the load from the source, and efficiently build up charge with minimum start-up costs. In addition, we propose a feedback-loop technique called Dynamic Energy Burst Scaling (DEBS) that follows the load’s optimal power point and minimizes its energy.

III. PRELIMINARIES

In a new cyber-physical system paradigm, transiently powered systems are designed to operate in limited energy harvesting scenarios. In order to execute an atomic task, such as reading a sensor value or transmitting a data packet, these




Fig. 1: Feedback loop for Dynamic Energy Burst Scaling.

systems need to be able to buffer the required energy, otherwise its completion cannot be guaranteed. Consequently, we argue that a novel Energy Management Unit (EMU) is needed to provide energy guarantees in such disadvantageous scenarios. Due to the limited energy intake in transiently powered systems, the unit should self-start requiring as little time and energy as possible. During those short periods of limited energy intake, it maximizes the energy build-up by harvesting at the source’s optimal power point. When powering the load with short energy bursts, it should provide a control interface to the load so its optimal power point can be tracked. In this work, we present an EMU that satisfies these requirements, shown in Fig. 1.

The proposed Dynamic Energy Burst Scaling (DEBS) technique aims to exploit the EMU’s control interface by closely following the application’s minimum required power envelope. To illustrate with an example, imagine a simple low-power camera application with two tasks: 1) acquisition and 2) processing. The first requires a camera supplied with 3 V, while the second requires only 2 V. One approach, used in [16], uses bursts of constant size and supply voltage. Using DEBS, our proposed EMU is able to produce one burst at 3 V to acquire an image, and another burst at 2 V to process it, thus minimizing the total energy.

IV. SYSTEM MODEL

In this section, we describe our model of the proposed Energy Management Unit (EMU), shown in Fig. 3. One of the main goals is to derive equations which can apply to a wide variety of energy sources and loads. The proposed model will then be used to optimize important system parameters, namely the EMU’s start-up costs and the load’s energy. The accuracy of the proposed model will be experimentally validated in Sec. VI.

A. Energy Buffering and Losses

The amount of energy buffered in the EMU depends on several parameters including the input power and load powers, and the system’s non-idealities. The equation governing the time-dependent energy level in a capacitor is as follows:

$$E'_{cap}(t) = \frac{d}{dt} E_{cap}(t) = \eta_{boost} (V_{in}(t), I_{in}(t)) \times P_{in}(t) - P_{load}(S_i)/\eta_{buck} - P_{leak}(t) \quad (1)$$

In this equation, the positive term represents the energy intake, while the negative ones represent the energy consumption.

Input Power: The system has only one power input, $P_{in}(t)$, supplied by the harvester’s transducer. This work focuses on the scenario where $P_{in} < P_{load}$ and $V_{in} < V_{load}$. In order to maximize the transducer’s efficiency, the maximum power point must be tracked to account for variable harvesting conditions.

Load Power: In the proposed model, the load can have two states (S_i): active or inactive. When active, the load is characterized by three quantities: $E_{burst,i}$, $V_{load,i}$, $P_{load,i}$; where $E_{burst,i}$ defines the energy burst size required for one




Fig. 2: Start-up time and cold-start energy overhead.

execution of task i , $V_{load,i}$ its supply voltage and $P_{load,i}$ the power consumption during the execution of task i . These parameters were characterized experimentally. In the inactive state, the load is in deep sleep, consumes very little power, and awaits the trigger from the energy management unit.

Converter Efficiencies: Since decoupled systems have the source and load operating at different power points, voltage converters are used. This step, while necessary, introduces non-negligible losses, which are represented by boost and buck converter efficiencies $\eta_{boost}(V, I)$ and η_{buck} . The boost converter's efficiency is particularly sensitive to the operating voltage and current, meaning it must be parameterized. These efficiencies were also characterized experimentally, and a simple look-up table is used for simulations.

Other Energy Losses: Unfortunately, converter inefficiencies are not the only sources of energy losses. The maximum power point tracking unit and the control circuit also consume energy. The consumption of the control circuit I_{ctrl} and buck converter I_{buck} consists of a constant current and resistive component and hence depends on V_{cap} . For the energy buffer, a capacitor of size C_{cap} , a resistive leakage R_{cap} is assumed. Considering these components, the system leakage is summarized as:

$$P_{leak}(t) = V_{cap}(t) \times (I_{ctrl}(V_{cap}(t)) + I_{buck}(V_{cap}(t))) + V_{cap}(t)^2 / R_{cap}. \quad (2)$$

Equations (1) and (2) can accurately describe the time evolution of the system's energy levels, as will be shown in Sec. VI-D. They will be used in the remainder of this section to estimate how different parameters impact the system's losses, to then calculate the optimal parameters that minimize the losses.

B. Minimizing Cold-Start Energy and Start-up Time

Given the system model presented above, we can start optimizing the cold-start energy and start-up time. By definition this is the fixed start-up cost to turn a transient system on. Fig. 2 shows that after a period of energy unavailability, the capacitance first needs to be recharged to the level of V_{load} . In order to minimize these fixed costs for a given input power, we need to minimize the start-up time defined as:

$$t_{start-up} = \left\{ t \mid V_{cap}(t) = \sqrt{\frac{2 \int_0^t E'_{cap}(\tau) d\tau}{C_{cap}}} = V_{load} \right\} \quad (3)$$

However, the minimum capacitance is limited by the EMU's maximum supported voltage swing, as shown in the following equation:

$$C_{min,i} = \frac{2E_{load,i}}{\eta_{buck}(V_{max}^2 - V_{load,i}^2)}, \quad (4)$$




Fig. 3: Transient system architecture with proposed EMU.

where $E_{load,i}$ and $V_{load,i}$ are the energy and voltage required to execute task i , and V_{max} is the EMU's maximum supported voltage. The optimal capacitor value is then selected as the highest $C_{min,i}$ among all tasks i .

C. Minimizing Load Energy

To show the advantages of our EMU's boost-buck architecture compared to the boost-only architecture, let us consider the case of supplying a constant current load. Assuming the load has a maximum supply voltage tolerance from V_{max} down to V_{min} , we have the following power consumption: for boost-only architecture the average power of a task is $P_A = (V_{min} + V_{max})/2 \cdot I_{load}$, while the buck has a constant power of $P_B = (V_{min} \cdot I_{load})/\eta_{buck}$. By comparing these two power consumptions, it directly follows that buck converter reduces the load's power consumption, if the following condition for the buck converter efficiency holds:

$$\eta_{buck} > \frac{2V_{min}}{V_{min} + V_{max}} \quad (5)$$

To illustrate with a numerical example, suppose a load has a voltage tolerance of 3 to 5 V. This means that a buck converter has a lower power consumption if $\eta_{buck} > 75\%$. Furthermore, the use of a buck adds the possibility of tracking the load's optimal power point. When an application consists of multiple tasks with different voltage requirements, we can use Dynamic Energy Burst Scaling (DEBS) to minimize the load's energy.

V. SYSTEM ARCHITECTURE

In this section we present the architecture of a transient system with the proposed Energy Management Unit (EMU). Fig. 3 shows the system's main components. Since the EMU can work with a wide variety of sources, this section focuses on the rest of the system. In Sec. VI we will discuss one specific source used to evaluate the proposed system.

A. Energy Management Unit (EMU)

This component is tasked with building up energy and producing short bursts to power the load. The EMU provides a control interface to dynamically adjust the bursts' size and voltage. Our proposed Dynamic Energy Burst Scaling (DEBS) technique exploits this by using a feedback loop to track the load's optimal power point and minimize its energy consumption.

Converters: The harvesting part of the system is based on the commercial bq25505 energy harvesting chip. This chip uses a boost converter to convert the input voltage to a level where the energy can be stored in a storage device. Using its integrated maximum power point tracking (MPPT), the boost converter adjusts the input impedance such that the power source always operates at its optimal power point to maximize




Fig. 4: Load's software execution flow.

the harvested energy. To provide the required output voltage to the load, the TPS62740 buck converter is directly connected to the energy buffer. This buck converter was chosen for two reasons: 1) its high efficiency, and 2) its digitally controlled adjustable output voltage. Since the measured η_{buck} is $\geq 85\%$ and the EMU supports a maximum voltage swing from 5-2 V, condition (5) confirms that using a buck converter is more efficient. Furthermore, it allows tracking the load's power point, which will be discussed in Sec. V-C.

Energy Buffer: The energy buffer between the input voltage boosting and output voltage regulation guarantees complete separation of the harvesting and load supply unit and therefore allows independent optimization of these parts. The buffer used in our implementation is a small SMD capacitor of 80 μF . It should be noted that this value is not the optimal theoretical value of 62.5 μF calculated using (4), but it is the minimum value supported by our harvester chip. The chosen capacitance minimizes the EMU's start-up time and energy cost from zero.

Control Circuit: The control circuit manages the burst size as well as the output voltage and oversees the energy accumulation in the buffer. For the first, the battery OK signal of the bq25505 is used to trigger the activation of the load, once the capacitor voltage reached the threshold level V_{th} at which the requested energy burst was accumulated. The threshold voltage V_{th} is configured using a resistor network, which is controlled by a digital switch to dynamically adjust V_{th} and therefore the burst size. The load supply voltage V_{load} can be controlled using the TPS62740 buck converter's digital input.

B. Application Circuit (Load)

As an example for a typical sensing load, we use a low power image acquisition application. The hardware is composed of an MSP430FR5969 microcontroller and a Centeye Stonyman image sensor, which both feature low power consumption and ultra-low power deep sleep modes. The IO state lock mechanism and non-volatile FRAM features of the microcontroller are important to keep the interface state of the energy manager during deep sleep and maintain the task configuration across periods of energy unavailability. The application consists of two tasks: 1) an image acquisition task to read the sensor, and 2) a basic image processing task. The former has an average power consumption of 3.77 mW and minimum voltage requirement of 3 V due to the external camera, while the latter can operate at 2 V and consumes 2.74 mW.

The execution flow of the load's software is shown in Fig. 4. When the system exits cold-start after a long period of energy unavailability, known as a Power-On-Reset (POR), the microcontroller performs some basic initialization and immediately enters deep sleep. With a measured power consumption of $<50\text{ nW}$, it minimizes losses during the buildup of energy for

the next burst. When the next burst is generated, the EMU triggers a control signal to wakeup the load. The system then reads the next task configuration and starts its execution after initializing the peripherals needed for that task. At the end of the task, the configuration is updated and the next required burst is configured. Afterwards, the load enters deep sleep again and waits for the next energy burst to build up.

C. Feedback Control for DEBS

As was discussed in Sec. IV-C, there are many application scenarios where the load has a varying optimal power point. This occurs when tasks use peripherals with substantially different voltage requirements. For this scenario, our proposed EMU provides a control interface to dynamically adjust the burst size and voltage. Our proposed DEBS technique is based on a feedback loop (Fig. 1) that allows the load to configure the EMU to supply the energy burst at the optimal operating point. This configuration takes so little instructions that it is negligible with respect to simple tasks. Following our image acquisition example, when DEBS is used, the EMU generates two bursts. During the first burst, 215.0 μJ at 3 V were requested. Once enough charge has been built up, the EMU's control circuit configures the buck converter's digital input to set the output to 3 V and triggers the load to acquire the image. Afterwards, the load uses the EMU's interface again and requests the second burst by setting the energy and voltage to 149.7 μJ and 2 V, respectively. So long as the EMU's buffer has energy, the buck converter will maintain this output voltage until the next burst is generated, the next task executed, and the load requests the next energy burst size and voltage.

Without DEBS, the EMU would generate one burst at 3 V to acquire and process the image, similar to the approach proposed in [16], which leads to significantly larger bursts sizes due to the non-optimal operating point. These two approaches, dynamic and constant bursts, will be evaluated experimentally in the following section.

VI. EVALUATION

This section evaluates the costs, performance and efficiency of the Energy Management Unit (EMU). To this end, we begin the EMU evaluation by testing its cold-start energy and start-up time, an important result of our optimized buffer size. We will then test the performance the EMU by powering an image acquisition system using two execution profiles: 1) Dynamic Bursts (using DEBS), and 2) Constant Bursts (no DEBS). Both execution profiles will be tested under a) constant power input, and b) variable power input. These last experimental results will also be compared to a discrete-time simulation of the model presented in Sec. IV using Matlab.

A. Experimental Setup

As sample energy source, the MP3-25 flexible solar panel from PowerFilm was used and it was exposed only to low light levels (125-600 lux), since we focus on low power harvesting scenarios. For the analysis of the system performance, the following metrics are used in all experiments:

- $E_{in} = \int_0^{T_{exp}} P_{in}(t) dt$, for the total input energy,
- $E_{app,j} = \sum_{i=1}^{N_{tasks}} \int_{t_{active,i,j}} P_{load}(t) dt$, for active energy consumed by the j -th application execution,




Fig. 5: System evaluation under constant input power conditions: execution rate R_{exec} and system energy efficiency η_{sys} .

- $E_{load} = \sum_j E_{app,j}$, the total energy consumed by the load for all application executions,
- $\eta_{sys} = E_{load}/E_{in}$, the total system efficiency, and
- $R_{exec} = N_{bursts}/(N_{tasks} \cdot T_{exp})$, the application execution rate.

In the formulas above, $t_{active,i,j}$ denotes time interval of task i in j -th application execution, N_{tasks} the number of tasks in the application, and N_{bursts} the number of bursts during the experiment of duration T_{exp} . By measuring the currents and voltages at the relevant points, we can accurately measure P_{in} and P_{load} , and experimentally calculate these metrics.

B. Start-Up Time and Cold-Start Energy Costs

As was discussed in Sec. IV-B, the energy buffer was optimized to minimize the cold-start's required energy and start-up time and still guarantee the completion of atomic tasks. To characterize these costs, the capacitor was completely discharged, and the flexible solar panel was exposed to constant illumination level until the cold-start phase ended. The measured start-up time and cold-start energy as a function of the input power can be seen in Fig. 6. The maximum cost, which occurs at the minimum input power of $20 \mu\text{W}$ was 118 s and 2.49 mJ. This was expected since the harvester, by definition, cannot operate efficiently in this region. It should be noted that with an input power of $400 \mu\text{W}$, the start-up costs go down to 3.9 s and 1.54 mJ. This translates to reduced reaction times and energy costs in transient systems, and highlights the importance of a minimized capacitance.

C. Constant Input Power

In this experiment, the flexible solar panel was exposed to a constant illumination level for 5 min, supplying the EMU with a constant power. This experiment was repeated for different power levels using Dynamic and Constant Bursts. The system was then analyzed using the previously discussed performance metrics dependent on the input power level.

The resulting analysis of the task execution rate and system efficiency is shown for both task execution profiles in Fig. 5, together with model-based simulation results for the same

scenarios. The results show up to 50 and 39 task executions per minute when using Dynamic and Constant Bursts, respectively. For both profiles, the system efficiency η_{sys} reaches more than 70% for a wide range of input power, with a peak system efficiency of up to 75.1% for Dynamic Bursts. It should be noted that the system model presented in Sec. IV allows accurate simulation of the number of task executions as well as system efficiency. However, some additional non-linear leakage effects of the boost converter at very low input power of $< 50 \mu\text{W}$ in combination with high buffer capacitor voltages cannot be captured by the model and result an optimistic efficiency for Constant Bursts.

The experimental results show that execution rate R_{exec} when using Dynamic Bursts is on average 27% higher than Constant Bursts. Further, Dynamic Bursts lowers the minimal system operating input power down to $19 \mu\text{W}$ compared to $36 \mu\text{W}$ for Constant Bursts. Lastly, the system efficiency η_{sys} is increased across the whole input power range for Dynamic Bursts, with significant improvements for input powers of $200 \mu\text{W}$ and below.

D. Variable Input Power

In this experiment the system performance was evaluated in an indoor real-world scenario, again for both task execution profiles. The two execution profiles were each evaluated with a 15 min experiment that included walking around with the setup in the office hallway partly illuminated by natural and artificial light, walking in a dimly lit basement and sitting at a well illuminated office desk.

The experimental metrics for Dynamic and Constant Bursts under variable input power conditions are shown in Table I. The first thing to note is that Dynamic Bursts reduces the average energy per application execution by 19.7% when compared to Constant Bursts. Even though the Dynamic Bursts experiment had on average a lower input power P_{in} , both the execution rate and energy efficiency η_{sys} are still higher. This can be explained by the lower energy consumption per task execution due to DEBS' minimization of load energy. Normalized to the average input power, this results in 22% more task executions with Dynamic Bursts compared to Constant Bursts, which shows the considerable advantage of using DEBS.




Fig. 6: Measured start-up time and cold-start energy.

TABLE I: Execution rate and efficiency for variable P_{in} .

Burst Size	Avg. P_{in}	Metric	Simulation	Experiment	Error
Dynamic	92.3 μW	R_{exec}	9.93 min^{-1}	10.33 min^{-1}	-3.9%
		avg. $E_{app,j}$	368.4 μJ	369.0 μJ	-0.2%
		E_{load}	54.9 mJ	57.2 mJ	-4.0%
		η_{sys}	66.11%	68.82%	-3.9%
Constant	111.9 μW	R_{exec}	9.87 min^{-1}	9.93 min^{-1}	-0.7%
		avg. $E_{app,j}$	460.8 μJ	459.7 μJ	+0.2%
		E_{load}	68.2 mJ	68.5 mJ	-0.4%
		η_{sys}	67.76%	68.01%	-0.4%




Fig. 7: Time domain comparison between model and experimental evaluation of Dynamic Bursts under variable input power.

The table also compares the experimental results to the model simulation that uses the experimental P_{in} data as input. Here, the comparison to experimental values shows that even in a real world scenario with variable input power, the model is able to predict the system behavior with maximum error of 4% for both task execution profiles. This fact is also reflected in Fig. 7: it shows the input power, simulated and measured energy level of the buffer capacitor during a 40 second sample time window of the Dynamic Bursts experiment. Beside a small time drift in the energy accumulation during very low input power, where not all effects can be represented by our model, it tracks the buffer's energy level and bursts with high accuracy. This high accuracy results only in small deviation in the time diagram, despite the accumulation of simulation errors in the time domain.

E. Result Discussion

The results from the constant power characterization and variable input experiment highlight the four main advantages of our proposed approach. First, thanks to our minimized energy buffer, the cold-start energy and start-up time are minimized: at $400 \mu W$, they were only 1.54 mJ and 3.9 s, respectively. Second, in the very common low power harvesting scenario for transient systems, the EMU completely decouples the source's and the load's power points. Even though the harvested power never surpassed $400 \mu W$, the EMU still provided 3.83 mW load with a 75.1% energy efficiency using DEBS. With direct coupling, it is simply impossible to power the same load. Third, the proposed DEBS technique uses the EMU in a feedback loop to track the load's optimal power point and significantly reduce its energy consumption. Fourth, the proposed model is able to accurately predict the experimental results. This validates the minimization of our model parameters, namely the minimized energy buffer and start-up costs.

VII. CONCLUSIONS

In this work, we have presented an energy management unit (EMU) that minimizes the cold-start energy and start-up time for transiently powered systems. By accumulating only the minimum amount of energy in an optimally-sized buffer, the EMU is able to supply generic loads predictably and efficiently, even when harvesting only a small fraction of the load's power. Furthermore, we proposed a Dynamic Energy Burst Scaling (DEBS) technique to track the load's optimal power point. Using a simple interface consisting of only a few digital inputs, the EMU is able to dynamically adjust the burst size and voltage according to an application's needs, thus minimizing the load's energy consumption. Our model is able to predict the system's performance and energy efficiency within 4.0% of the experimental values, even under variable power input conditions. Lastly, the proposed principles

push the limits of energy proportionality by lowering the input power requirements and maintaining a high energy efficiency.

Acknowledgments This research was funded by the Swiss National Science Foundation under grant 157048: Transient Computing Systems. The authors would like to thank Lukas Cavigelli, Pascal Hager and David Bellasi for their insightful discussions.

REFERENCES

- [1] P. Zhang, D. Ganesan, and B. Lu, "QuarkOS: Pushing the operating limits of micro-powered sensors," *Proc. 14th USENIX Conf. Hot Top. Oper. Syst.*, 2013.
- [2] A. Gomez, C. Pinto, A. Bartolini, D. Rossi, L. Benini, H. Fatemi, and J. P. de Gyvez, "Reducing energy consumption in microcontroller-based platforms with low design margin co-processors," in *Proc. DATE*, 2015.
- [3] D. Brunelli, L. Benini, C. Moser, and L. Thiele, "An Efficient Solar Energy Harvester for Wireless Sensor Nodes," in *Proc. DATE Conf.*, 2008.
- [4] M. Magno, S. Marinkovic, D. Brunelli, E. Popovici, B. O'Flynn, and L. Benini, "Smart power unit with ultra low power radio trigger capabilities for wireless sensor networks," in *Proc. DATE Conf.*, 2012.
- [5] A. Yakovlev, "Energy-modulated computing," in *Proc. DATE Conf. IEEE*, 2011.
- [6] K. Ahmed and S. Mukhopadhyay, "A 190nm bias current 10mv input multi-stage boost regulator with intermediate node control to supply rf blocks in self-powered wireless sensors," *IEEE Trans. Power Electronics*, vol. PP, no. 99, 2015.
- [7] C. Lu, S. P. Park, V. Raghunathan, and K. Roy, "Efficient power conversion for ultra low voltage micro scale energy transducers," in *Des. Autom. Test Eur. Conf. Exhib. (DATE)*, 2010, 2010.
- [8] D. Balsamo, A. S. Weddell, G. V. Merrett, B. M. Al-hashimi, D. Brunelli, and L. Benini, "Hibernus : Sustaining Computation during Intermittent Supply for Energy-Harvesting Systems," *Embed. Syst. Lett. IEEE*, vol. 7, no. 1, 2015.
- [9] H. Jayakumar, A. Raha, and V. Raghunathan, "QUICKRECALL: A Low Overhead HW/SW Approach for Enabling Computations across Power Cycles in Transiently Powered Computers," *Proc. Int. Conf. VLSI Design*, 2014.
- [10] H. G. Lee and N. Chang, "Powering the iot: Storage-less and converter-less energy harvesting," in *Proc. ASP-DAC Conf.*, 2015.
- [11] Y. Wang, Y. Liu, C. Wang, Z. Li, X. Sheng, H. Lee, N. Chang, and H. Yang, "Storage-less and converter-less photovoltaic energy harvesting with maximum power point tracking for internet of things," *IEEE Trans. CAD*, vol. PP, no. 99, 2015.
- [12] C. Wang, N. Chang, Y. Kim, S. Park, Y. Liu, H. G. Lee, R. Luo, and H. Yang, "Storage-less and converter-less maximum power point tracking of photovoltaic cells for a nonvolatile microprocessor," in *Proc. ASP-DAC Conf.*, 2014.
- [13] Y. Kim, N. Chang, Y. Wang, and M. Pedram, "Maximum power transfer tracking for a photovoltaic-supercapacitor energy system," in *Proc. ISLPED*, 2010, pp. 307-312.
- [14] E. Dallago, A. Barnabei, A. Liberale, P. Malcovati, and G. Venchi, "An interface circuit for low-voltage low-current energy harvesting systems," *IEEE Trans. Power Electronics*, vol. 30, no. 3, 2015.
- [15] E. Dallago, A. Lazzarini Barnabei, A. Liberale, G. Torelli, and G. Venchi, "A 300 mv low-power management system for energy harvesting applications," *IEEE Trans. Power Electronics*, vol. PP, no. 99, 2015.
- [16] S. Naderiparizi, A. N. Parks, Z. Kapetanovic, B. Ransford, and J. R. Smith, "WISPCam : A Battery-Free RFID Camera," in *Proc. IEEE RFID*, 2015.