

Efficient, Long-Term Logging of Rich Data Sensors using Transient Sensor Nodes

ANDRES GOMEZ, ETH Zürich
LUKAS SIGRIST, ETH Zürich
THOMAS SCHALCH, ETH Zürich
LUCA BENINI, ETH Zürich and Univeristy of Bologna
LOTHAR THIELE, ETH Zürich

While energy harvesting is generally seen to be the key to power cyber-physical systems in a low-cost, long term, efficient manner, it has generally required large energy storage devices to mitigate the effects of the source's variability. The emerging class of transiently powered systems embrace this variability by performing computation in proportion to the energy harvested, thereby minimizing the obtrusive and expensive storage element. By using an efficient Energy Management Unit (EMU), small bursts of energy can be buffered in an optimally-sized capacitor and used to supply generic loads, even when the average harvested power is only a fraction of that required for sustained system operation. Dynamic Energy Burst Scaling (DEBS) can be used by the load to dynamically configure the EMU to supply small bursts of energy at its optimal power point, independent from the harvester's operating point. Parameters like the maximum burst size, the solar panel's area as well as the use of energy-efficient Non-Volatile Memory Hierarchy (NVMH) can have a significant impact on the transient system's characteristics such as the wake-up time and the amount of work that can be done per unit of energy. Experimental data from a solar-powered, long-term autonomous image acquisition application show that, regardless of its configuration, the EMU can supply energy bursts to a 43.4 mW load with efficiencies of up to 79.7% and can work with input power levels as low as 140 μ W. When the EMU is configured to use DEBS and NVMH, the total energy cost of acquiring, processing and storing an image can be reduced by 77.8%, at the price of increasing the energy buffer size by 65%.

CCS Concepts: • **Computer systems organization** → **Embedded hardware**;

General Terms: Design, Algorithms, Energy

Additional Key Words and Phrases: Energy Harvesting, Energy Efficiency, Low Power Design

ACM Reference Format:

Efficient, Long-Term Logging of Rich Data Sensors using Transient Sensor Nodes *ACM Trans. Embedd. Comput. Syst.* X, X, Article X (January 2017), 24 pages.

DOI: <http://dx.doi.org/10.1145/3047499>

1. INTRODUCTION

Over the past decade, there has been a considerable research effort to reduce the energy consumption of electronic devices. While there has been considerable progress, the lifetime of battery-based devices remains the bottleneck in their development. Broadly speaking, the problem of how to supply low power embedded systems with the energy they require in an efficient, low-cost, long-term, scalable, and self-sustainable manner

This work is supported by the Swiss National Science Foundation, under grant 157048: Transient Computing Systems.

Author's addresses: A. Gomez and L. Sigrist and T. Schalch and L. Benini and L. Thiele, Department of Electrical Engineering and Information Technology, ETH Zürich. L. Benini, Department of Electrical Engineering, University of Bologna.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 1539-9087/2017/01-ARTX \$15.00

DOI: <http://dx.doi.org/10.1145/3047499>

has not yet been adequately solved. Over-provisioning with large energy harvesting and storage elements is either infeasible or unnecessary in many application scenarios such as wearable, autonomous, miniaturized or "smart dust" systems. Fortunately, a purely harvesting driven system can still meet application requirements in many of these scenarios.

Transiently powered systems 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 can be highly variable but not necessarily high enough to complete even one atomic task execution, such as performing a sensor reading or transmitting a radio packet. Consequently, such systems need to be able to buffer at least the amount of energy needed to bridge this power deficit and thereby to guarantee the completion of any single task to be executed.

1.1. Transient System Challenges

The design and optimization of transient systems must take into account some general properties to be able to execute applications in an efficient and reliable manner. The main properties we have identified are the following:

- (1) *Energy sources are variable and environmental conditions unknown*
Micro-level energy harvesting, which is capable of powering low power Wireless Sensor Network (WSN) nodes, has received considerable attention in recent years. Interested readers can read [Bhatti et al. 2016] for a detailed survey on existing energy harvesting and wireless energy transfer solutions. One important characteristic to all of them, is the intermittent availability of harvested energy. In many cases, harvesting the maximum power from a particular source (e.g. solar, thermal, etc) requires an impedance matching circuit which dynamically adjusts to changing environmental conditions. Since the system designer has no control over the environmental conditions, no assumptions can be made about the evolution of source's maximum power point. Unless the complete sensor node and its peripherals are fully scalable in terms of voltage and current, a system which maximizes the input energy must decouple the source's power point from that of the load.
- (2) *Loads have highly variable I/V characteristics and energy requirements*
Typical low-power cyber-physical systems have components such as microcontrollers, memories and peripherals (e.g. sensors and wireless transceivers). Microcontrollers usually have a wide operating voltage range, but on-chip converters operate most efficiently at lower supply voltages [Gomez et al. 2015]. External peripherals such as sensors and radios can have substantially different voltage requirements, but to minimize cost system designers avoid having multiple voltage domains and simply choose the highest minimum voltage required to supply the entire system. Different tasks, however, can have highly varying energy/voltage requirements and current consumption since these are highly dependent on the application, peripherals used, etc. In many cases, like ultra low power RFID-scale devices or wearable TEG harvesters, the operating voltage is so low that high voltage peripherals cannot be used [Salajegheh et al. 2013; Thielen et al. 2017]. 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.
- (3) *Storage element must be minimized*
In many application domains, such as wearable systems, there are stringent form factor restrictions, making storage devices such as batteries and supercapacitors

particularly unsuited for long-term use. They are also expensive in terms of cost and area, can have limited charge cycles and high self-discharge rates, impose current peak limits and might not be easily integrated on board [Zhang et al. 2013]. Consequently, designing a system with an over-dimensioned buffer, if at all possible, invariably leads to higher losses due to power harvesting costs, additional leakage current, and converter inefficiencies [Hester et al. 2015]. It is thus indispensable to optimize the energy buffer's size according to application-specific parameters, in order to limit these negative effects.

The aforementioned properties, when combined, form a veritable challenge from the system design perspective. Conventional low power design dictates that any harvesting-based system should, at all times, maximize the energy input and minimize the energy output. To maximize the harvested energy in dynamic scenarios, the sources' maximum power point needs to be tracked. To minimize the load's energy, its power needs to be minimized by dynamically adjusting its operating voltage. These two criteria can differ significantly, especially given the fact that environmental conditions can exhibit great variability. This variability in turn demands that the design use energy storage to provide some minimal energy guarantees, otherwise no program progress can be guaranteed. Designing an architecture that addresses these issues requires innovative methods that combine both hardware and software aspects. In particular, the main challenge is to design a system that can still operate efficiently, have minimized storage and wake-up times given these conditions.

1.2. Transient System Configurations

As has been previously discussed, the different properties of transiently powered systems require novel approaches to operate efficiently in such disadvantageous scenarios. In this work, we argue that an additional Energy Management Unit (EMU), shown in Fig. 1, is needed to maximize the harvested energy, minimize the load's power, and provide the load with the energy guarantees necessary for program progress. 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 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 such that the application circuit's optimal power point can be tracked.




Fig. 1: Superset of transient node configurations. Dashed lines indicate the components or connections proposed in this work.

While existing works [Magno et al. 2012; Yakovlev 2011; Ahmed and Mukhopadhyay 2015; Lu et al. 2010] have looked at low power systems with energy harvesting and storage capabilities, these are expensive in terms of harvesting and storage requirements for long-term, efficient functionality under transient power conditions. State of

the art transient system design [Balsamo et al. 2015; Jayakumar et al. 2014], connects the energy source directly to the load, without any other intermediaries. However, these works only focus on non-atomic execution of processing tasks, and only work when specific harvesting conditions generate a safe operating voltage on the solar panel. In our previous work [Gomez et al. 2016], an Energy Management Unit (EMU) was first proposed to decouple the operating point of the source from the load. Additionally, the EMU can apply Dynamic Energy Burst Scaling (DEBS) to track the load's optimal power point. The EMU allows a system to operate predictably and efficiently with limited energy buffering, even under very low power harvesting conditions where the harvested power is much lower than the load's minimum power requirement. In this work, we focus on specific design aspects of transient nodes which do long term logging of rich data sensors. In particular, we consider the example of a solar-powered vision sensors targeting life-logging applications. These transient sensor nodes have the property of guaranteed information and energy availability, since darkness does not provide neither energy nor information and light provides both.

1.3. Our Contributions

Rich data sensors such as cameras bring their own challenges to transient system design. Logging applications are particularly costly, due to the large volume of data that rich data sensors produce. We thus propose a novel Non-Volatile Memory Hierarchy (NVMH), which increases the energy efficiency of rich data sensor logging applications. We will show how the addition of the NVMH introduces a trade-off between the energy cost per stored byte and the minimum energy buffer size. Our proposed EMU-based design uses 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.

The main contributions of this work are summarized as follows:

- Energy Management Unit (EMU) that efficiently converts variable low power levels to short, high power energy bursts.
- Feedback-based Dynamic Energy Burst Scaling (DEBS) technique to track the load's optimal power point.
- Non-Volatile Memory Hierarchy (NVMH) that reduces the energy cost of long-term data logging.
- Accurate model to optimize system's application-specific parameters for low input power scenarios, including energy and data buffer sizes as well as harvester's dimension.
- Experimental validation of the high energy efficiency and proportionality of the proposed transfer scheme in long-term image acquisition application, as well as the trade-off between energy cost per image stored and minimum buffer size.

The remainder of this work is organized as follows. In Sec. 2, we will give a detailed overview of the state of the art in transient system design and commercially available non-volatile memory technologies. The model and architecture of the EMU are presented and discussed in detail in Sec. 3. The main application scenario, long-term logging of a rich data sensor, as well as the architecture of our baseline application circuit is explained in Sec. 4. Our novel energy-efficient Non-Volatile Memory Hierarchy (NVMH), composed of both Flash and FRAM memories is presented in Sec. 5. The optimized system design, which includes the NVMH for long-term logging, is described in detail, along with its design trade-offs, in Sec. 6. The experimental evaluation of our different load configurations and analysis of the energy efficiency and proportionality are shown in Sec. 7. Lastly, we conclude our work in Sec. 8.

2. RELATED WORK

Cyber-physical systems have traditionally been used in conjunction with energy harvesting and energy storing. When coupled with aggressive duty-cycling techniques, they are able to significantly reduce their average power consumption, possibly to the point of self-sustainability. Due to the prohibitive costs of storing energy, there is a new trend to design systems with minimized storage capacity. As a consequence of this limited capacity and the variability of energy harvesting, Non-Volatile Memory (NVM) is required to ensure data consistency. For long-term logging applications, this poses a challenge to store large amounts of data. While there are many different memory technologies, each has its own characteristics in terms of storage density, read/write cycles, and power consumption. In this section, we will summarize the state of the art of transient systems and non-volatile memory technologies.

2.1. Transient System Architectures

Broadly speaking, we can identify three types of architectures for transient systems:

Directly-Coupled. When the energy source has an I-V curve compatible with the load, they can be directly connected. The authors of [Balsamo et al. 2015; Jayakumar et al. 2014] 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. An approach to federating energy proposed in [Hester et al. 2015] increases the computational ability by using multiple independent capacitors, each dedicated to a specific peripheral. In [Lee and Chang 2015; Wang et al. 2015; Wang et al. 2014; Kim et al. 2010], 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 in CPU bounded applications, it does not minimize the load's energy consumption and is limited to a narrow active power range. Even though directly-coupled systems avoid converter losses, if the power input is below this narrow active 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. In contrast to traditional battery-based systems, these decoupled transient systems have a limited energy buffer between the source and load.

Boost-Only. In [Dallago et al. 2015a; Dallago et al. 2015b], 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 $2.5 \mu\text{W}$. As will be explained in Sec. 3.2, our capacitance is chosen to minimize the cold-start energy and time.

Boost-Buck. The authors of [Naderiparizi et al. 2015] 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. 7, it will be shown that this approach can lead to a substantially higher energy consumption, larger storage elements and longer start-up times.

2.2. Non-Volatile Memory Technologies

In the design of transient sensor nodes, one of the key considerations is the choice of Non-Volatile Memory (NVM). Due to the inherent power cycling of transient nodes, any data saved in volatile memory will be lost. Consequently, the choice of NVM technology is closely related to the application's power envelope, reliability and storage requirements.

One of the most mature non-volatile technologies used in cyber-physical systems today is Flash Memory. While it offers one of the highest storage densities available, it suffers from two main drawbacks: power consumption and reliability. SD Cards alone, for example, can have a capacity in the order of 100's of GB and a power consumption in the range of 50-100 mW. Each memory block also has a restrictive read/write cycle limit of $10^4 - 10^5$ [Torres et al. 2013].

For many years, researchers have been actively searching for new NVM technologies that can (ideally) offer ultra-low power consumption, unlimited read/write cycles, ultra high densities, and compatibility to standard fabrication processes. Unfortunately, no silver bullet has been found yet, and commercially available technologies offer different trade-offs between these important parameters. We will not provide a summary of recent advances in NVM technologies, though readers can find surveys in [Mittal and Vetter 2016; Moreau 2013]. Instead, we will focus on one specific technology, Ferroelectric Random Access Memory (FRAM), which is a promising candidate for unified (instruction and data) memory due to its high endurance (10^{15} read/write cycles), and its low power consumption [Texas Instruments 2015]. The main limitation of FRAM for long-term logging of rich data sensors is its capacity, since the largest commercially available capacity is in the order of 100's of KB.

3. ENERGY MANAGEMENT UNIT (EMU)

In this section, we describe the model and architecture of the Energy Management Unit (EMU), shown in Fig. 3. One of the main goals is to derive an analytical model which can be applied to a wide variety of energy sources and loads. The 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 model will be experimentally validated in Sec. 7.

3.1. Modeling 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:

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

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 $P_{in} < P_{load}$. In order to maximize the transducer's efficiency, the maximum power point must be tracked to account for variable harvesting conditions.

Load Power. In our model, the load can have two states (S_j): 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 execution of task i , $V_{load,i}$ its sup-




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

ply voltage and $P_{load,i}$ the power consumption during the execution of task i . These parameters will be 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 $\eta_{buck}(V, I)$. The boost converter's efficiency is particularly sensitive to the operating voltage and current, meaning it must be parameterized. For these efficiencies a look-up table is used for simulations. The overall system efficiency of the EMU will be bounded by the product of the boost and buck converter efficiencies. While this depends on both the input and output voltage/current, broadly speaking for our application domain, it goes up to $\sim 75\%$, see our experimental results in Sec. 7.

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 a resistive component and hence depends on V_{cap} . For the energy buffer, a capacitor of size C_{cap} , a resistive leakage R_{cap} in parallel 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. 7.4. 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 these losses.

3.2. Optimizing 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 $V_{load,min} = \min_i\{V_{load,i}\}$. In order to minimize these fixed costs for a given input power, we need to minimize the start-up time defined as:

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

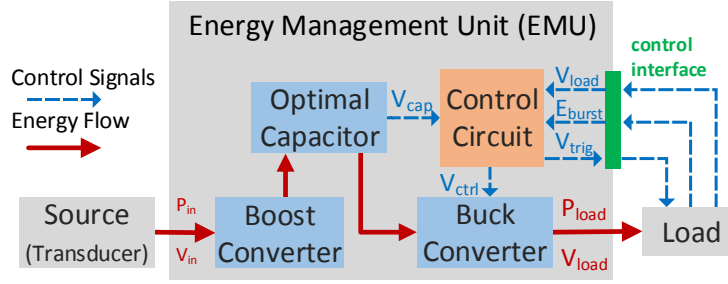


Fig. 3: Architecture of the Energy Management Unit (EMU).

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

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

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. These values must be known at design time, such that the optimal capacitor value can be selected as the highest $C_{min,i}$ among all tasks i , i.e. $C_{optimal} = \max_i\{C_{min,i}\}$. For the implementation we selected the next higher available capacitor size C_{buffer} to guarantee task completion.

3.3. 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, consuming I_{load} . The harvesting power of a transiently powered system which is typically much smaller than the load's power consumption, therefore has a negligible impact on the linear voltage decrease during the time in which the load is supplied with an energy burst. Assuming the load has a maximum supply voltage tolerance from V_{max} down to V_{min} , this results in the following power consumption: for the boost-only architecture the average power of a task is $P_A = (V_{min} + V_{max})/2 \times I_{load}$, while the buck converter provides a constant power of $P_B = (V_{min} \times I_{load})/\eta_{buck}$ to the load. By comparing these two power consumptions, it directly follows that a buck converter reduces the load's power consumption, if the following lower bound for its conversion 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 system using a buck converter has a lower power consumption if $\eta_{buck} > 75\%$. Furthermore, the use of a buck converter adds the possibility of tracking the load's optimal power point for all tasks by dynamically switching the voltage level. 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.

3.4. EMU Architecture

The Energy Management Unit (EMU) controls the buildup of energy from the source, and controls the energy transfer to the load, or application circuit. Its main components can be seen in Fig. 3, which are discussed individually in the following subsections.

Harvesting and Buffering. 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 that 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 the harvested energy. To provide the required output voltage to the load, the TPS62740 buck converter is directly connected to the 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. As was shown in Sec. 3.2, the storage element minimization is application-specific. In this work, we will show that two different implementations of the same application can exhibit a trade-off between storage size, and the amount of work completed with the same energy budget.

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 a threshold level V_{th} . At this voltage the energy level is reached at which enough energy has been accumulated to provide the requested energy burst to the load. The variable burst size dependent threshold voltage V_{th} is configured using a resistor network. This comparator threshold can be switched digitally from the control circuit by selecting between different resistor networks. Beside very large resistor values, the bq25505 control circuit uses duty cycling to reduce the energy consumption of the comparator and resistor network. The load supply voltage V_{load} can be controlled directly using the TPS62740 buck converter's digital input.

Requirements for EMU Operation. Thanks to its inherent decoupling of source and load power points, the minimum requirements for EMU operation are conceptually independent from the load and are only tied to the EMU's circuit implementation. In our case, the first requirement is a minimum input voltage of 330 mV, which is required to turn on a diode in the BQ harvester. If this requirement is met, charge is transferred from the solar panel to a small capacitance. After a certain voltage on this capacitance is reached, the main boost converter is turned on. This transition requires a minimum input current of $\sim 60 \mu\text{A}$. This means that so long as the input power is greater than $\sim 20 \mu\text{W}$, the EMU is guaranteed to exit the cold-start phase and enter the energy build-up phase. During this phase, the charge on the capacitor will increase as long as $\overline{P_{in}} > P_{leak,max}$, where

$$P_{leak,max} = \{P_{leak}(t') \mid V_{cap}(t') = V_{max}\}. \quad (6)$$

After some time, which depends on the input power, enough charge will be accumulated in the capacitor to guarantee a task completion. This makes only the frequency of task activations dependent on the harvesting conditions, but not the task itself. Once triggered, the actual task execution is guaranteed regardless of the harvesting conditions.

4. LONG-TERM LOGGING OF RICH DATA SENSORS

Up to now, we have discussed how energy can be efficiently buffered, even in low power harvesting scenarios. How this energy is consumed, however, is a function of the load itself. To highlight the flexibility and efficiency of our approach, we will focus on nodes which do long-term logging of rich data sensors. More specifically, we focus on transiently powered vision sensors, which acquire and store images in Non-Volatile Memory (NVM).

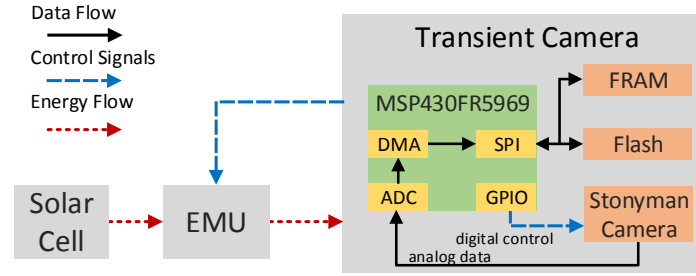


Fig. 4: Transient Camera Architecture for Long-Term Logging.

4.1. Transient Node Architecture

The architecture of our proposed system can be seen in Fig. 4. The transiently powered vision sensor is composed of an MSP430FR5969 microcontroller and a Centeye Stonyman image sensor, both of which feature low power consumption and ultra-low power deep sleep modes. The IO state lock mechanism and the microcontroller’s non-volatile features are important to keep the interface state of the energy manager during deep sleep and maintain the task configuration across periods of energy unavailability. In order to store long-term data, external memories with high capacities are needed. One of the most commonly used technologies for this purpose is an SD Card, which is based on Flash. As will be highlighted in Sec. 5, SD Cards generally suffer from very high power and energy consumption but offer the highest densities.

4.2. Energy Burst Configuration

In order to configure the EMU for correct operation, it is important to characterize the application’s voltage and energy needs. For our rich data sensor logger, the baseline application consists of three tasks: 1) image acquisition to read the sensor, 2) basic image processing, and 3) image storage to copy data from volatile to non-volatile memory. Table I shows the energy burst configuration (V_{load}, E_{burst}) for each task. As was previously mentioned in Sec. 3.2, the EMU’s capacitor has to be dimensioned according to the largest task: transferring an image to the SD Card. The energy cost for this transfer has two components, a constant initialization cost and a transfer cost which depends on the amount of data transferred. For the baseline application using DEBS, which acquires, processes and transfers a single image, the costliest task is the image transfer which requires 11.67 mJ. The minimum available capacitance that can store this energy between 5.1V and 2.7V is 1470 μ F. If DEBS is not used then all tasks would need to execute in a single burst at constant voltage, requiring a minimum capacitance of 2200 μ F.

Table I: Baseline voltage requirements and energy costs for individual task execution.

Task	Voltage (V_{load})	Energy (E_{burst})
Acquire One Image	3.0 V	156 μ J
Process One Image	2.0 V	527 μ J
SD Card Initialization	2.7 V	10 536 μ J
SD Card Data Transfer (per Image)	2.7 V	1 137 μ J

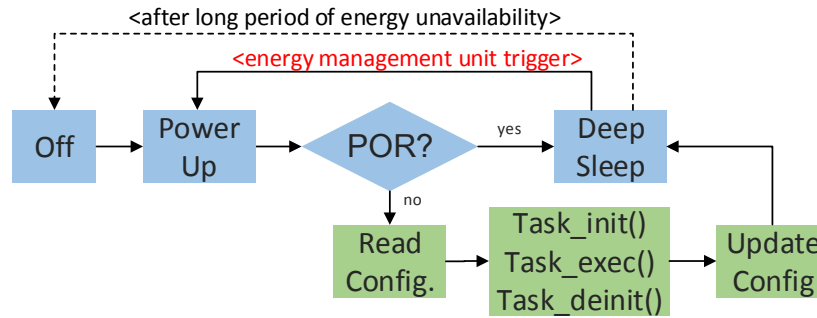


Fig. 5: Load's software execution flow.

4.3. Software Execution Flow

The execution flow of the load's software is shown in Fig. 5. 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 < 600 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. For all of our evaluated applications, all of the tasks will be repeatedly executed in a static schedule. For example, the baseline application will repeatedly execute the following tasks: 1) acquire one image, 2) process one image, 3) store one image on the SD Card.

4.4. Feedback Control for Dynamic Energy Burst Scaling (DEBS)

As was discussed in Sec. 3.3, 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, the EMU provides a control interface to dynamically adjust the burst size and voltage. Our 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 baseline image acquisition example, when using DEBS the EMU generates three bursts, one for each task. During the first burst, $156 \mu\text{J}$ 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 (image processing) by setting the energy and voltage to $527 \mu\text{J}$ and 2 V, respectively. Lastly, at the end of the second burst, the load requests the third burst for storing an image with energy and voltage set to 11.67 mJ and 2.7 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 only be able to generate bursts at a constant voltage of 3 V. This results in an approach similar to the one proposed in [Naderiparizi et al. 2015], where one large burst would be used to acquire, process, and store one image.

This approach leads to significantly larger burst sizes due to the grouping of tasks with a non-optimal operating point. These two approaches, single burst and DEBS based bursts, will be evaluated experimentally in Sec. 7.

5. ENERGY-EFFICIENT NON-VOLATILE MEMORY HIERARCHY

In order to make long-term logging of rich data sensors viable, the energy cost per unit of data stored needs to be as small as possible. To this end, we propose the use of a novel Non-Volatile Memory Hierarchy (NVMH) that combines the best characteristics of different technologies. The use of a NVMH is independent from Dynamic Energy Burst Scaling (DEBS). The former is an application-side optimization to reduce the cost of storing data to non-volatile memory, while the latter is EMU-based feedback technique to split a single transient application into task-based burst executions. Over the following sections, we will discuss how to combine both techniques to optimize the system design of a transient image sensor for long-term logging applications.

5.1. Non-Volatile Memories (NVM)

As was discussed in Sec. 2.2, Flash is the most mature NVM technology available and boast very large storage capacities, but suffers from high energy consumption. Contrarily, Ferro-electric Random Access Memory (FRAM) offers ultra-low power/energy consumption, but is limited to very low storage capacities.

Flash. While Flash memories can have a very large capacity (up to 100's of GB), they have a high power consumption during the initialization and write phases. The time-consuming initialization procedure results in a very high constant energy overhead as it is needed before the actual memory can be accessed. In burst-powered transient systems, this initialization has to be paid for every energy burst that requires SD Card access and can lead to prohibitively large overheads. Furthermore, Flash memories alone have low durability (around 100 000 read/write cycles) and typically use a controller to spread the wear of individual memory cells evenly, also called wear-leveling [Hu et al. 2013]. Unfortunately, wear-leveling leads to additional delays and energy consumption and introduces additional variability in the total overhead.

FRAM. In recent years, FRAM has emerged as a viable alternative with very low power consumption. FRAM provides ultra-low power read and write actions at high read and write speeds. Beside high energy efficiency, FRAM also provides very high durability. Unfortunately, due to the incompatibility with standard manufacturing processes FRAM components have relatively low storage capacities (~ 100 's KB).

In the following section, we show how these two memory technologies can be combined in long-term logging applications to exploit the advantages of each technologies: Flash's high density and FRAM's energy efficiency.

5.2. Non-Volatile Memory Hierarchy (NVMH)

Novel NVM technologies have not reached densities that allow them to replace Flash in data-intensive applications. However, since they consume significantly lower energy than traditional Flash, they open the door to reducing the energy requirements of long-term logging.

We propose to use energy efficient FRAM to increase the efficiency of transiently powered logging applications that rely on an SD Card as a large storage device for long-term logging. More specifically, we introduce a small FRAM data cache before the SD Card to distribute its high initialization cost among several image transfers. As shown in Fig. 4, this can be added to the sensor node as an additional component on the SPI bus. Thanks to the ultra-low energy read and write operations in FRAM, multiple images can be cheaply buffered before transferring them in one single batch to the

SD Card. Compared to writing the images one by one on the SD Card, as described in Sec. 4, the high SD Card initialization cost is required only once for writing *all* images buffered in the FRAM to the SD Card. This reduces the SD Card initialization cost by a factor determined by the FRAM buffer size. Because of the high energy overhead for initialization, this non-volatile memory hierarchy allows significant reductions in the energy needed to acquire, process and store one image. The experimental results described in Sec. 7 will show that on average 294.8% more images can be stored in Flash when the baseline application discussed in Sec. 4.2 uses a 10 image FRAM buffer, compared to the same application without it.

It is important to note that the non-volatility property of the FRAM is a key requirement in this memory hierarchy, because this buffer is only supplied with energy when an acquired image is buffered or the full buffer is flushed to the SD Card. In between bursts with memory access, it is powered off to reduce leakage losses. Deploying any volatile component in the memory hierarchy, like SRAM, is not suitable for a transiently powered system because the content would be lost at the point where the memory is turned off, or when the system enters cold start due to low input power.

The proposed memory hierarchy combines the advantages of low energy consumption available in novel NVM technologies with the high density of traditional Flash. This allows building long-term rich data sensor logging applications with large storage requirements and high energy efficiency.

5.3. Energy Cost vs Minimum Capacitor Size Trade-Off

The key design parameter of the proposed memory hierarchy is the size of its intermediate FRAM buffer. The larger the buffer, the more images can be buffered before flushing them to the SD Card. By distributing the initialization cost of the SD Card, the average energy requirement per image decreases. The energy overhead for reading and writing the image to the FRAM buffer is negligible compared to the SD Card: even with a buffer of only 2 images, the energy saved for one initialization is larger than the additional energy cost for the FRAM. However, because of the larger transfer size from FRAM to the SD Card during the flush operation, the amount of energy that needs to be guaranteed by the EMU also increases. This results in a trade-off between required energy buffer size and energy cost of the storing data in NVM. In this work, which focuses on long-term logging of rich data sensors, our primary concern is minimizing the energy cost of non-volatile data. How one can select the optimal parameter values at design time will be discussed in Sec. 6.1.

6. OPTIMIZED SYSTEM DESIGN

In this section we demonstrate how the Energy Management Unit (EMU) model with Dynamic Energy Voltage Scaling (DEBS) presented in Sec. 3.1 is used during the design of a transient system for determining individual system parameters. For that, we use the following two use-cases: 1) selecting the FRAM size in the Non-Volatile Memory Hierarchy (NVMH), and 2) dimensioning the solar panel for desired application performance.

6.1. FRAM Buffer Size and Cost/Capacitance Trade-Off

As was previously mentioned, the main focus of this work is to minimize the total energy cost E_{img} required to acquire, process and store image data in non-volatile memory for long-term logging applications. It has already been shown that for generic applications, using DEBS reduces the load's energy requirements through task-level optimizations. Now we will focus on optimizing the non-volatile storage component of E_{img} through the use of both DEBS and NVMH. To design the NVMH, we will first select the FRAM buffer size, since this has a direct impact on the energy cost reduction.

As a second step, we will select the value for the capacitance such that all other costs like form factor and start-up time/energy are minimized.

As long as Flash memory remains the cheapest and densest technology available for embedded systems, it will always be required in long-term image logging applications. In the transient application scenario, which involves duty-cycling the SD Card, it is inevitable to pay the SD card's initialization cost before transferring any data. Since this initialization cost is constant, the main purpose of the NVMH is to divide this cost, on average, over as many images as possible. Because of this, the first relevant parameter is how large the FRAM buffer can be. Currently, the main limitation is technology scaling, which has so far prevented FRAM to have capacities greater than 100's KB. At the time of writing, the largest commercially available FRAM can only buffer 10 images from our vision sensor.

Once the FRAM buffer size has been determined, the only remaining question is how small the buffer capacitance can be such that the energy savings are as high as possible for a given FRAM buffer size. To minimize the capacitance, it is important to understand how much energy is needed to flush the entire FRAM buffer to Flash. To simplify the discussion, we assume that once the FRAM buffer is full, the application schedule ensures the buffer is fully flushed before acquiring new pictures. Depending on the capacitor size, which limits the maximum burst size, flushing the entire buffer might take one or more bursts. The data-dependent energy cost to transfer a block of one or more images from the FRAM to the SD card is given by

$$E_{trans}(N_{img}) = N_{img} \times (E_{FRAM,img} + E_{Flash,img}), \quad (7)$$

where N_{img} is the total number of images that can be buffered in FRAM before being flushed to the SD card and $E_{x,img}$ is the data-dependent cost of storing one image in memory X . As mentioned earlier, if the EMU buffer capacitor is not big enough to guarantee that the FRAM can be flushed in one single burst, then several bursts will be needed for the complete data transfer. The number of bursts N_{burst} required to complete the FRAM flush can be calculated as

$$N_{burst}(N_{img}, E_{buff}) = \left\lceil \frac{E_{trans}(N_{img})}{E_{buff} - E_{mem,init}} \right\rceil, \quad (8)$$

where $E_{mem,init}$ is the initialization cost of both FRAM and Flash memories, and $E_{buff} = \frac{1}{2}C_{buff}(V_{max}^2 - V_{min}^2)$ is the energy stored in the EMU buffer capacitor C_{buff} between the flush task's lowest operating voltage of $V_{min}=2.7$ V and the EMU's highest buffer voltage of $V_{max}=5.1$ V. Using (7) and (8), it is possible to calculate the total energy E_{flush} required to flush the FRAM buffer of N_{img} images as

$$E_{flush}(N_{img}, E_{buff}) = N_{burst}(N_{img}, E_{buff}) \times E_{mem,init} + E_{trans}(N_{img}). \quad (9)$$

Fig. 6 plots the normalized energy per image stored transfer cost $\bar{E}_{flush} = E_{flush}/N_{img}$. Each line in the plot represents a specific EMU configuration with a different C_{buff} . As discussed earlier, this has a direct influence on the number of bursts required to flush the FRAM buffer. For a given capacitance value, the image cost decreases monotonically up until the point where the data flush needs to be split into more than one burst. This step increase in $E_{flush,img}$ is introduced by the need of an additional Flash initialization for the subsequent bursts and decreases as the transfer size becomes multiple of the specific buffer size. From (9) it follows that all E_{trans} minima of a given C_{buff} are equal, since the NVMH mechanism amortizes the SD Card initialization cost by the same ratio. The optimized C_{buff} is the minimum capacitance that has not reached its first minimum for a given FRAM buffer size, since any larger capacitance will not introduce higher savings but would require larger start-up costs.

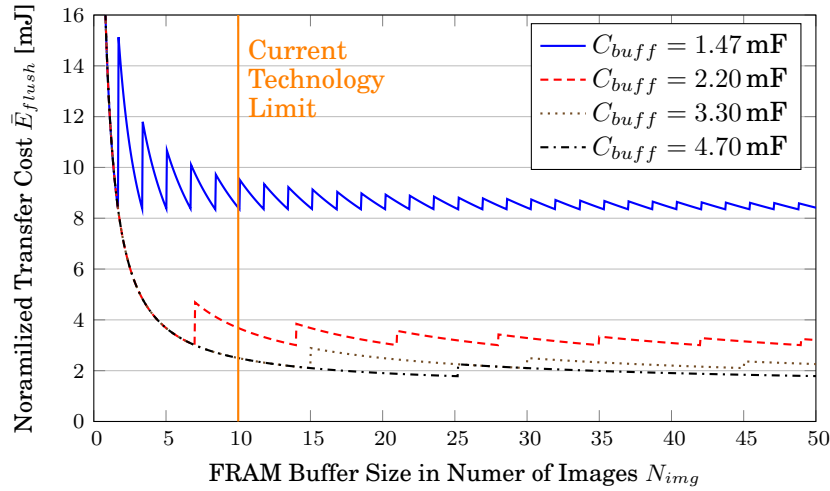


Fig. 6: DEBS+NVMH simulation results show how a given FRAM buffer size can have very different energy cost per image stored if the capacitor size is not chosen appropriately.

Optimized Burst Configuration. As mentioned earlier in Sec. 4.2 for the baseline DEBS application, it is important to characterize the application’s voltage and energy needs in order to guarantee correct operation of the EMU. With the introduction of the Non-Volatile Memory Hierarchy (NVMH), the application scheduling needs to be modified to 1) image acquisition to read the sensor, 2) basic image processing, 3) buffer processed image in the FRAM, and 4) when the buffer is full, flush the entire buffer to the SD Card. With respect to the energy costs presented in Table I, there is only the additional task of buffering one image in FRAM. This task was characterized to require $32 \mu\text{J}$ at 2.0 V. Since the energy cost of transferring data to the SD card has a data proportional term, the new cost of writing to Flash is much larger than without NVMH. This is because a new transfer of 10 images is needed to flush the entire FRAM buffer, instead of transferring a single picture to Flash. Compared to the DEBS Only configuration, the minimum capacitance increases from $1470 \mu\text{F}$ to $3300 \mu\text{F}$, while the energy cost per stored image was lowered from 11.11 mJ to 2.73 mJ when using DEBS+NVMH. This is the fundamental trade-off presented by NVMH: significantly reduced energy costs of stored images at the expense of a larger capacitance.

Design Space. So far, three critical parameters for a transient vision sensor with NVMH have been discussed: FRAM buffer size, EMU buffer capacitance, and energy cost per image stored. Fig. 6 shows that when using DEBS+NVMH, the minimum energy cost and start-up time/energy is achieved when the FRAM buffer is maximized and the capacitance minimized for that buffer size. Fig. 7 shows the energy cost per image stored vs the minimum capacitor size of all four configurations in a Pareto plot. It should be noticed that whenever DEBS is applied to a baseline configuration, the result is a reduced capacitor *and* a reduced energy cost per image stored. In fact, the Pareto front consists of all configurations that employ DEBS. By contrast, whenever NVMH is added to a baseline application, the energy cost per image stored always decreases, but at the expense of a larger buffer capacitor. This corresponds to the expected behavior from previously presented models and highlights the fact that there is a trade-off between energy per image stored and buffer capacitance.

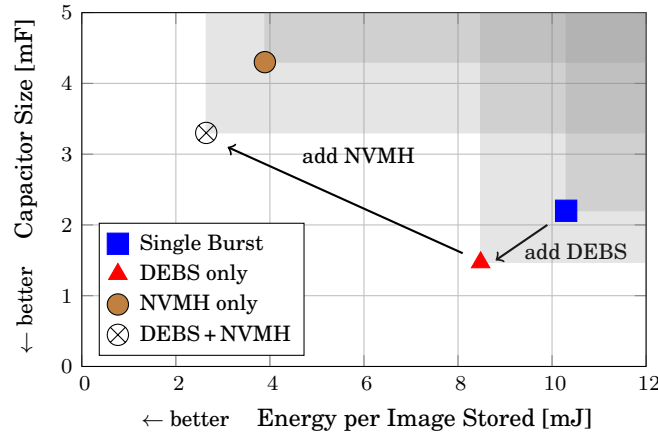


Fig. 7: Pareto plot of the trade-off between minimal buffer size and energy used to acquire, process and store an image. Regions dominated by the shown design points are grayed out. The design goal is to minimize both the energy per image stored and the capacitor size.

6.2. Minimum Harvester Size

In order to calculate the necessary size of the solar panel for an application, it is important to determine the desired performance. Depending on the lighting conditions the sensor node will be exposed to, the size of the solar panel might need to be bigger or smaller. Simulation can be used to optimize the solar panel area.

Four different PowerFilm flexible solar panels with sizes ranging from $12.7\text{ mm} \times 64\text{ mm}$ to $37\text{ mm} \times 114\text{ mm}$, with corresponding areas from 8 cm^2 to 42 cm^2 , were considered as harvesting sources. We first evaluated the output power of the individual solar panels for illuminance levels in the range from 100 lx up to 2000 lx . To do so, the solar panels were connected to the bq25505 harvester with maximum power point tracking (MPPT). The buffer capacitor of the harvesting circuit was replaced by a source meter (Keithley SMU 2450), which was configured to keep the buffer voltage constant at the expected average capacitor voltage of 4 V . For measuring the actual power extracted from the solar panel a small shunt resistor of $10\ \Omega$ was inserted between the solar panel and the bq25505 to measure the current. The power was then calculated using the measured current and the solar panel voltage.

The input power derived from these measurements was then used to simulate the application's behavior under the given constant illuminance using the model presented in Sec. 3.1. For this use-case the application configuration that includes both Dynamic Energy Burst Scaling (DEBS) and the Non-Volatile Memory Hierarchy (NVMH) with a FRAM buffer size of 10 images was used in the model. The application was simulated for a time window of one hour and the output is then analyzed to determine the average number of images that can be acquired, processed and stored per minute.

In Fig. 8, the results of these simulations are shown for solar panels with different areas A_{panel} depending on the illuminance level. It can be seen that for a typical indoor illuminances ranging from 250 lux to 1000 lux , a solar panel with an area of 42 cm^2 can acquire, process and store from 1 up to 11 images per minute. The dependency of the application performance on the solar panel area A_{panel} is also clearly visible, with lower areas the number of images acquired, processed and stored decreases. Considering an example where the expected room illuminance is 750 lux and the application should

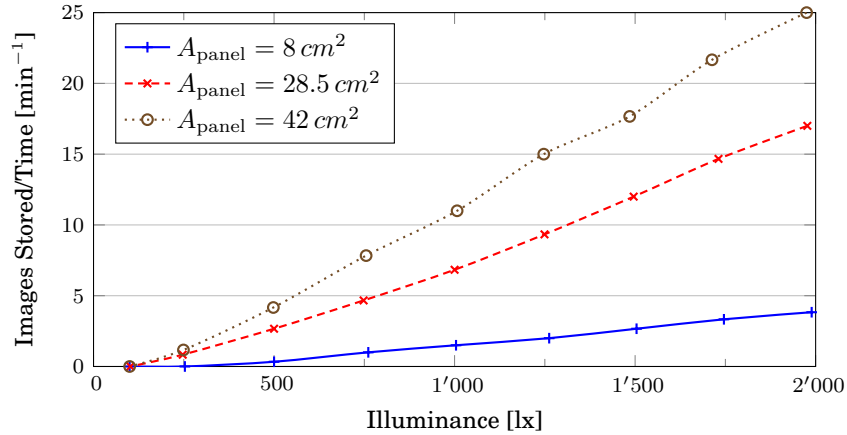


Fig. 8: The application’s performance is shown as the average number stored images per minute vs the illuminance level for different solar panel areas A_{panel} .

acquire 6 images per minute (one every 10 seconds on average), the results show that the solar panel with $A_{\text{panel}} = 42 \text{ cm}^2$ should be used to get the desired performance. This is also the solar panel that will be used for experimental evaluation in Sec. 7.

7. EVALUATION

This section evaluates the costs, performance and efficiency of an EMU-based transient vision sensor in four different configurations. These configurations represent all the possible combinations of our two main contributions: Dynamic Energy Burst Scaling (DEBS) and Non-Volatile Memory Hierarchy (NVMH). We will compare the performance vs efficiency trade-offs that these different combinations introduce. For the purposes of our sample life-logging application, our main goal is to minimize the energy costs/requirements since it will allow the transient camera to acquire, process and store the greatest number of pictures during a day.

7.1. Experimental Setup

The performance of our wearable prototype will be tested for each configuration listed in Table II using long-term measurements. Experiments will be done in both controlled (constant) and real-world (variable) harvesting conditions. All evaluation results shown in this section were done with the same MP3-37 flexible solar panel from PowerFilm, which has an area of 42 cm^2 . For the analysis of the system and application performance, all the relevant voltages and currents in the source, EMU and load were measured. To compare the experimental and simulation results, the input power traces were also recorded and used as input to our Matlab model.

Table II: Description of the evaluated transient configurations.

Configuration Name	Burst Size	Execution Profile	DEBS	NVMH
Single Burst	$\sum_i E_{task,i}$	Entire Application	X	X
DEBS Only	$\max_i \{E_{task,i}\}$	Single Tasks	✓	X
NVMH Only	$\sum_i \{E_{task,i}\}$	Entire Application	X	✓
DEBS + NVMH	$\max_i \{E_{task,i}\}$	Single Tasks	✓	✓

Transient Load Configurations. In order to identify the effect of both DEBS and NVMH, different configurations will be individually tested for prolonged periods of time. The performance of each configuration, under different harvesting conditions, is measured, compared and contrasted.

Single Burst. This configuration is the baseline for all comparisons. It buffers the energy for one entire application execution and does it in a single burst with constant voltage. This means that within a single burst, one picture is taken, processed and stored in Flash directly.

DEBS Only. This configuration, described in Sec. 4.2, uses Dynamic Energy Burst Scaling for each task. This means that the first burst does acquisition, the second does processing, and the third saves to Flash. Each of these bursts is configured to its optimal voltage.

NVMH Only. This configuration introduces the Non-Volatile Memory Hierarchy (NVMH) with a FRAM buffer size set to ten, as was described in Sec. 6.1. This means that, conceptually, each burst executes an application cycle. The first 9 bursts write only to the FRAM buffer: each burst doing acquisition+processing+buffering. The 10th burst performs an additional SD card flush, which transfers the ten buffered images from FRAM to Flash.

DEBS + NVMH. This configuration combines the previous two, but with one task executed per burst. This means that it takes 30 bursts (acquisition, processing and buffering 10 times each) to fill up the FRAM buffer. One additional burst transfers the buffered data to Flash. Again, each burst is executed at its optimal voltage.

Performance Metrics. In order to compare the performance of different transient configurations, the following metrics are calculated for 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,
- $E_{load} = \sum_j E_{app,j}$, the total energy consumed by the load for all application executions,
- $\eta_{sys} = E_{load}/E_{in}$, the total system efficiency,
- $E_{image} = E_{load}/N_{image}$, the average energy cost for acquiring, processing and storing one image in NVM, and
- $\Theta_{app} = N_{image}/T_{exp}$, the average number of images acquired, processed and stored per time.

In the formulas above, $t_{active,i,j}$ denotes the execution time of task i in j -th application execution, N_{tasks} the number of tasks in the application, and N_{image} the number of images acquired, processed and stored in Flash memory during the experiment of duration T_{exp} .

7.2. Start-Up Time and Cold-Start Trade-Offs

As was discussed in Sec. 3.2, each configuration requires a different minimized capacitor that guarantees the completion of its largest atomic task. This minimized capacitor in turn minimizes the required energy and start-up time for each configuration's cold-start. To characterize these costs that occur after an input power loss and depletion of the buffer capacitor, the buffer capacitor was completely discharged, and the flexible solar panel was exposed to constant illuminance level until the cold-start phase ended. The time measured to go through this cold-start phase as function of the input power is shown in Fig. 9 for all four configurations. As expected, the start-up time for all configurations decreases with higher input power. More specifically, a maximum start-up




Fig. 9: The cold-start time vs the input power for the EMU with different storage capacitor sizes.

time of 1055 s for the NVMH configuration ($C_{buff} = 4300 \mu\text{F}$) was reached at the minimum measured input power of $110 \mu\text{W}$. For an input power higher than $1000 \mu\text{W}$, all start-up times decrease to values below 44 s.

The start-up time overhead analysis shows the need for a minimized buffer capacitor, since this also minimizes the time/energy overhead of a given configuration. So long as the harvesting scenario provides enough energy, the transient node can exit cold-start and begin executing its application. In our specific life-logging scenario, we assume that typical human activities satisfy the requirements for EMU operation and last long enough for the transient vision node to start storing images.

7.3. Constant Input Power

In this part of the evaluation, the solar panel was exposed to a constant illumination level, resulting in an energy management unit with constantly supplied power. The experiments lasted for 10 min. For low input power levels of $200 \mu\text{W}$ and lower this time was extended to 15 min to observe a sufficient number of application triggers. For each application configuration, *Single Burst*, *DEBS only*, *NVMH Only* and *DEBS + NVMH*, the experiment was repeated for constant power levels ranging from $145 \mu\text{W}$ up to $1875 \mu\text{W}$. Measuring the currents and voltages at the EMU's input, output and buffer capacitor, as well as the load supply, the system's state and energy flow can be tracked to later calculate the performance metrics of the experiment. With these measurements, the previously introduced metrics system efficiency η_{sys} and the number of images stored per time Θ_{app} were calculated. The results of these metrics are analyzed depending on the different input power levels and discussed in detail in the following sections.

System Efficiency. The analysis of the system efficiency η_{sys} is shown in Fig. 10 for the four configurations mentioned earlier. The results show a consistent behavior: despite small variations between individual configurations, they all show a system efficiency η_{sys} that reaches at least 70% when the input power is greater than $1000 \mu\text{W}$. At the higher end of the evaluated input power the efficiencies asymptotically approach the maximum efficiency dictated by the product of the boost and buck converter efficiencies. During the experiments the highest observed efficiency reached a value of 78.6% at $1875 \mu\text{W}$, the highest input power for which the system was evaluated. Also common for all configurations is the fact that the efficiency drops sharply for input power levels

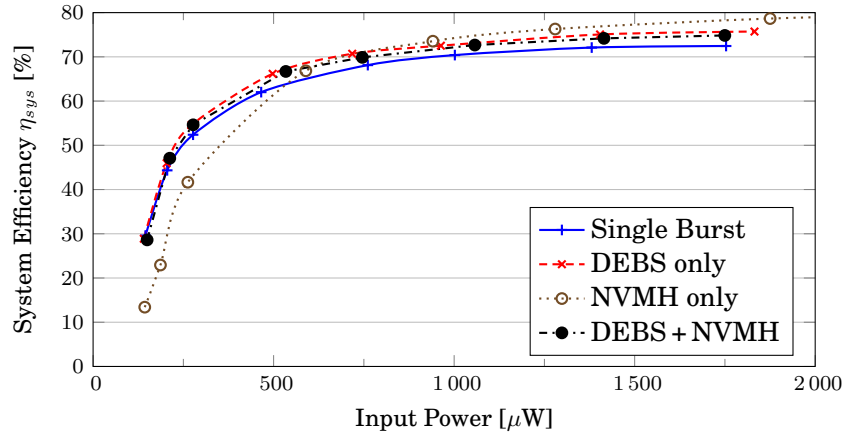


Fig. 10: Evaluation of the EMU's system efficiency at different input power levels.

close to the minimum required input power of $140 \mu\text{W}$. During experiments below that minimum input power level, no task executions were observed, resulting in a system efficiency η_{sys} of 0%.

Number of Stored Images. While the results for the system efficiency are very consistent, a large difference can be observed in the application performance that is characterized by the average number of stored images per time Θ_{app} . The results for this metric as a function of the input power level are shown in Fig. 11. The most noticeable difference is the increase in the number of images stored per minute when deploying *DEBS + NVMH* instead of *NVMH Only* or *DEBS only* instead of *Single Burst*: in either case, the inclusion of DEBS significantly increases the number of images stored per minute. In the case of DEBS Only, the increase was 30.1% at an input power of $1750 \mu\text{W}$, compared to Single Burst. Using the NVMH Only configuration offers a significant performance boost of 294.8% on average compared to Single Burst. It is the DEBS+NVMH combination, however, that clearly offers the highest number of stored images per time, thanks both of the proposed enhancements. The properties that all configurations have in common are the minimum input power of $140 \mu\text{W}$ and, once that power level is reached, the energy proportional increase of the number of stored images. Comparing the slopes of Θ_{app} for the individual configurations, it is visible that deploying DEBS instead of using the very basic *Single Burst* configuration already results in an improvement of 26% on average. Making use of DEBS+NVMH has the highest impact and increases the application performance by 268% or 365% when compared to the *DEBS only* or *Single Burst* configuration, respectively.

These experiments show the performance gain of deploying not only a dynamic energy burst scheme, but also an efficient memory hierarchy design for transiently powered logging applications.

7.4. Variable Input Power

The experiments discussed in this subsection were performed in an indoor real-world scenario, again for all three configurations. Each configuration was evaluated with a 15 min experiment that included walking around with the setup in the office hallway illuminated by artificial light, walking in a dimly lit basement and sitting at an office desk well illuminated by natural and artificial light.

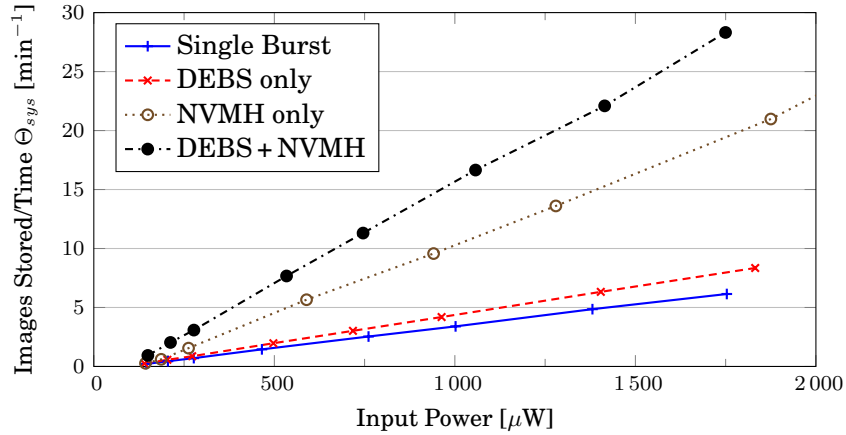


Fig. 11: Evaluation of the average number of images stored on the SD Card per time at different input power levels.

Table III: Results for variable input power experiments: total number of images stored per minute N_{img} , average of energy cost per image E_{img} and EMU efficiency η_{sys} .

Configuration	Avg. P_{in}	C_{buffer}	Metric	Simulation	Experiment
Single Burst	731.94 μW	2 000 μF	Θ_{app}	2.39 min^{-1}	2.13 min^{-1}
			avg. E_{img}	12.06 mJ	12.58 mJ
			η_{sys}	63.51%	60.93%
DEBS Only	706.98 μW	1 470 μF	Θ_{app}	2.62 min^{-1}	2.62 min^{-1}
			avg. E_{img}	10.54 mJ	11.12 mJ
			η_{sys}	64.48%	68.59%
NVMH Only	663.93 μW	4 300 μF	Θ_{app}	5.96 min^{-1}	6.28 min^{-1}
			avg. E_{img}	4.28 mJ	4.35 mJ
			η_{sys}	63.99%	68.65%
DEBS + NVMH	607.71 μW	3 300 μF	Θ_{app}	7.78 min^{-1}	8.29 min^{-1}
			avg. E_{img}	3.27 mJ	2.79 mJ
			η_{sys}	64.48%	63.53%

The experimental metrics for *Single Burst*, *DEBS Only*, *NVMH only*, and *DEBS + NVMH* under variable input power conditions are shown in Table III. The first thing to note is that adding DEBS to a baseline configuration reduces the average energy per image costs (E_{img}). This is the case for DEBS Only, which is an enhancement of the Single Burst configuration, as well as DEBS+NVMH, which is an enhancement of the NVMH Only configuration. This is expected since DEBS optimizes the load's operating point to minimize its energy requirements per task, simultaneously reducing the load's energy requirements as well as the minimum required capacitance. Comparing the buffer capacitances of the configurations without NVMH we see that the minimum required capacitor increases when NVMH is used. This is the fundamental trade-off of NVMH, and it is expected since the improved energy performance requires a larger energy guarantee for the SD Card flush task. Compared to the *Single Burst* configuration, the DEBS+NVMH configuration uses a 65% larger capacitor, but is able to reduce the average energy cost per stored image by 77.8%, down to only 2.79 mJ per image. It should also be noticed that even though the average input power during the *DEBS + NVMH* experiment was only 607.71 μW , the average number of images stored per minute was almost 4 \times that of Single Burst, up to 8.29 images per minute.




Fig. 12: Comparison of the energy buffer state for experiment E_{exp} and simulation E_{sim} based on the same variable input power trace P_{in} .

Table III also compares the experimental results to the model simulation that takes the measured harvested power 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 a maximum error of $\sim 6\%$ for most performance metrics. This fact is also reflected in Fig. 12: it shows the input power, simulated and measured energy level of the buffer capacitor during a 350 second sample time window of the *DEBS + NVMH* experiment. Beside a small time drift in the energy accumulation during very low input power, where not all effects can be represented accurately 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.

8. CONCLUSIONS

In this work, we have shown 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 capacitance, the EMU is able to supply generic loads predictably and efficiently, even when it harvests only a small fraction of the load's active power. Dynamic Energy Burst Scaling (DEBS) can be used with the EMU to track the load's optimal power point and minimize an application's energy. In long-term logging applications, where non-volatile memory can dominate the application's energy requirements, novel memory technologies such as FRAM can be introduced to form a Non-Volatile Memory Hierarchy (NVMH) that reduces the energy cost of storing data. Using a solar-powered camera example, we demonstrate an important design trade-off between the energy cost per image and the minimum required capacitance. By adding only a ten-image FRAM buffer, 77.8% of the energy cost per image can be saved, though the minimum required capacitance grows 65%. Furthermore, the EMU model can be used to dimension the solar panel to achieve a minimum performance for a given lighting condition. Experimental results show that a 42 cm² solar panel under indoor lighting conditions of 870 lux can be used to acquire, process, and store more than 11 images per minute on an SD Card. The EMU powered the 43.4 mW load at 69.90% efficiency requiring only 746 μ W input power. We believe that by minimizing an application's energy cost and storage requirements, both DEBS and NVMH significantly lower the current technology barriers for a new class of transiently-powered rich data sensor nodes capable of reliable and efficient execution of complex tasks.

ACKNOWLEDGMENTS

This research was funded by the Swiss National Science Foundation under grant 157048: Transient Computing Systems. The authors would like to thank Marco Zimmerling for his support in the submission of this manuscript and the anonymous reviewers for their helpful feedback.

REFERENCES

- K.Z. Ahmed and S. Mukhopadhyay. 2015. A 190nA 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* PP, 99 (2015). DOI: <http://dx.doi.org/10.1109/TPEL.2015.2423666>
- Domenico Balsamo, Alex S Weddell, Geoff V Merrett, Bashir M Al-hashimi, Davide Brunelli, and Luca Benini. 2015. Hibernus : Sustaining Computation during Intermittent Supply for Energy-Harvesting Systems. *Embed. Syst. Lett. IEEE* 7, 1 (2015). DOI: <http://dx.doi.org/10.1109/LES.2014.2371494>
- Naveed Anwar Bhatti, Muhammad Hamad Alizai, Affan A Syed, and Luca Mottola. 2016. Energy harvesting and wireless transfer in sensor network applications: Concepts and experiences. *ACM Transactions on Sensor Networks (TOSN)* 12, 3 (2016), 24.
- D Brunelli, L Benini, C Moser, and L Thiele. 2008. An Efficient Solar Energy Harvester for Wireless Sensor Nodes. In *Proc. DATE Conf.* DOI: <http://dx.doi.org/10.1109/DATE.2008.4484670>
- E. Dallago, A.L. Barnabei, A. Liberale, P. Malcovati, and G. Venchi. 2015a. An Interface Circuit for Low-Voltage Low-Current Energy Harvesting Systems. *IEEE Trans. Power Electronics* 30, 3 (2015). DOI: <http://dx.doi.org/10.1109/TPEL.2014.2322521>
- E. Dallago, A. Lazzarini Barnabei, A. Liberale, G. Torelli, and G. Venchi. 2015b. A 300 mV Low-Power Management System For Energy Harvesting Applications. *IEEE Trans. Power Electronics* PP, 99 (2015). DOI: <http://dx.doi.org/10.1109/TPEL.2015.2431439>
- Andres Gomez, Christian Pinto, Andrea Bartolini, Davide Rossi, Luca Benini, Hamed Fatemi, and Jose Pineda de Gyvez. 2015. Reducing energy consumption in microcontroller-based platforms with low design margin co-processors. In *Proc. DATE Conf.*
- Andres Gomez, Lukas Sigrist, Michele Magno, Luca Benini, and Lothar Thiele. 2016. Dynamic Energy Burst Scaling for Transiently Powered Systems. In *Proc. DATE Conf.*
- Josiah Hester, Lanny Sitanayah, and Jacob Sorber. 2015. Tragedy of the Coulombs: Federating Energy Storage for Tiny, Intermittently-Powered Sensors. In *Proc. SenSys Conf.* 5–16.
- Jingtong Hu, Chun Jason Xue, Qingfeng Zhuge, Wei-Che Tseng, and Edwin H.-M. Sha. 2013. Write Activity Reduction on Non-volatile Main Memories for Embedded Chip Multiprocessors. *ACM Trans. Embed. Comput. Syst.* 12, 3, Article 77 (April 2013), 27 pages. DOI: <http://dx.doi.org/10.1145/2442116.2442127>
- Hrishikesh Jayakumar, Arnab Raha, and Vijay Raghunathan. 2014. QUICKRECALL: A Low Overhead HW/SW Approach for Enabling Computations across Power Cycles in Transiently Powered Computers. *Proc. Int. Conf. VLSI Design* (2014). DOI: <http://dx.doi.org/10.1109/VLSID.2014.63>
- Younghyun Kim, Naehyuck Chang, Yanzhi Wang, and M. Pedram. 2010. Maximum power transfer tracking for a photovoltaic-supercapacitor energy system. In *Proc. ISLPED Conf.* 307–312.
- Hyung Gyu Lee and Naehyuck Chang. 2015. Powering the IoT: Storage-less and converter-less energy harvesting. In *Proc. ASP-DAC Conf.*
- Chao Lu, Sang Phill Park, V Raghunathan, and K Roy. 2010. Efficient power conversion for ultra low voltage micro scale energy transducers. In *Des. Autom. Test Eur. Conf. Exhib. (DATE), 2010.* DOI: <http://dx.doi.org/10.1109/DATE.2010.5457066>
- M. Magno, S. Marinkovic, D. Brunelli, E. Popovici, B. O'Flynn, and L. Benini. 2012. Smart power unit with ultra low power radio trigger capabilities for wireless sensor networks. In *Proc. DATE Conf.* DOI: <http://dx.doi.org/10.1109/DATE.2012.6176436>
- S. Mittal and J. S. Vetter. 2016. A Survey of Software Techniques for Using Non-Volatile Memories for Storage and Main Memory Systems. *IEEE Trans. Parallel Distr. Syst.* 27, 5 (May 2016), 1537–1550. DOI: <http://dx.doi.org/10.1109/TPDS.2015.2442980>
- Magnus Moreau. 2013. *Estimating the Energy Consumption of Emerging Random Access Memory Technologies*. Master's thesis. Institutt for elektronikk og telekommunikasjon, NTNU.
- Saman Naderiparizi, Aaron N Parks, Zerina Kapetanovic, Benjamin Ransford, and Joshua R Smith. 2015. WISPCam : A Battery-Free RFID Camera. In *Proc. IEEE RFID*.
- Narrative. 2016. Narrative Clip. <http://getnarrative.com/>. (2016). [Online; accessed 10-June-2016].
- Mastooreh Salajegheh, Yue Wang, Anxiao (Andrew) Jiang, Erik Learned-Miller, and Kevin Fu. 2013. Half-Wits: Software Techniques for Low-Voltage Probabilistic Storage on Microcontrollers with

- NOR Flash Memory. *ACM Trans. Embed. Comput. Syst.* 12, 2s, Article 91 (May 2013), 25 pages. DOI: <http://dx.doi.org/10.1145/2465787.2465793>
- Texas Instruments 2015. *MSP430FR59xx Mixed-Signal Microcontrollers*. Texas Instruments. Rev. E.
- Moritz Thielen, Lukas Sigrist, Michele Magno, Christofer Hierold, and Luca Benini. 2017. Human body heat for powering wearable devices: From thermal energy to application. *Energy Conversion and Management* 131 (Jan. 2017), 44–54.
- L. Torres, R. M. Brum, L. V. Cargnini, and G. Sassatelli. 2013. Trends on the application of emerging nonvolatile memory to processors and programmable devices. In *Proc. ISCAS*. 101–104. DOI: <http://dx.doi.org/10.1109/ISCAS.2013.6571792>
- Cong Wang, Naehyuck Chang, Younghyun Kim, Sangyoung Park, Yongpan Liu, Hyung Gyu Lee, Rong Luo, and Huazhong Yang. 2014. Storage-less and converter-less maximum power point tracking of photovoltaic cells for a nonvolatile microprocessor. In *Proc. ASP-DAC Conf.*
- Y. Wang, Y. Liu, C. Wang, Z. Li, X. Sheng, H.G. Lee, N. Chang, and H. Yang. 2015. Storage-less and Converter-less Photovoltaic Energy Harvesting with Maximum Power Point Tracking for Internet of Things. *IEEE Trans. CAD PP*, 99 (2015). DOI: <http://dx.doi.org/10.1109/TCAD.2015.2446937>
- Alexandre Yakovlev. 2011. Energy-Modulated Computing. In *Proc. DATE Conf.* IEEE.
- Pengyu Zhang, Deepak Ganesan, and Boyan Lu. 2013. QuarkOS: Pushing the operating limits of micro-powered sensors. *Proceedings of the 14th USENIX conference on Hot Topics in Operating Systems* (2013).

Received June 2016; revised December 2016; accepted January 2017