

# Selective Policies for Efficient State Retention in Transiently-Powered Embedded Systems: Exploiting Properties of NVM Technologies

Theodoros D. Verykios, Domenico Balsamo, Geoff V. Merrett

*Centre for Internet of Things and Pervasive Systems  
School of Electronics and Computer Science  
University of Southampton, UK*

---

## Abstract

Transiently-powered embedded systems are emerging to enable computation to be sustained during intermittent supply, without the need for large energy buffers such as batteries or supercapacitors. To deal with the intermittent nature of the input source, these systems save the system state (i.e. registers and main memory) to Non-Volatile Memory (NVM) before a power failure, and restore it when the power supply recovers. Existing approaches normally save the entire state of the system upon power failure, but this is both energy and time consuming. In this paper, we analyse existing approaches to identify their inefficiency when used with specific NVM technologies, and propose novel selective policies for efficiently retaining the system state by exploiting properties of different NVM technologies. These policies are based on (1) concatenating multiple images into the available NVM before erasing, and (2) efficiently selecting only the system state that has changed since last saving. The existing and proposed policies are experimentally validated on two embedded platforms featuring different NVM technologies (Flash and FRAM), depending on their characteristics, in order to identify the most energy efficient policy/platform combination. Results show a reduction in energy and time overhead of up to 90.6% for Flash memory using a novel policy, and 86.2% for FRAM, compared to the typical approach of saving the entire system state.

**Keywords:** Selective State Retention, Energy Harvesting, Non-Volatile Memory, Transient Computing

---

## 1. Introduction

Batteries have traditionally been used to power embedded systems. However, requirements such as a long lifetime, low cost, and low weight, pose significant challenges to battery-powered systems. In addition, the nature of some applications such as implantable bio-sensors [1, 2] and underground WSNs [3] implies limited access and, consequently, maintenance for battery replacement or recharging becomes a challenge. Therefore, the need for embedded systems that can operate without batteries has emerged [4].

Energy harvesting (EH) systems scavenge energy from environmental sources such as light, vibration, motion or temperature to power themselves, instead of relying on batteries [5]. However, factors such as the weather condition, availability of light, or the intensity of vibration can have a significant impact on energy availability. Relying solely on these sources can, therefore, result in the system being unable to sustain computation. The traditional solution to tackle this is the use of energy storage (e.g. a supercapacitor or rechargeable battery) to buffer




Figure 1: Typical operation of a transiently-powered system.

harvested energy so that the long-term energy consumed equals the harvested energy [6, 7]. However, these buffers increase the size, weight and cost of the devices, which makes the realisation of some systems infeasible.

Transiently-powered embedded systems are storage-less systems that enable computation to be sustained, despite the variable and unstable energy harvested from the environment [8]. Due to frequent power interruptions caused by the variable source, transient systems achieve forward progress by retaining their state in Non-Volatile Memory (NVM) upon a power failure. This implies that the main memory, core registers and general-purpose registers are saved before a power outage, and restored when the power is available once again.

Several software-based approaches have recently been

---

*Email addresses:* T.Verykios@soton.ac.uk (Theodoros D. Verykios), D.Balsamo@soton.ac.uk (Domenico Balsamo), gvm@ecs.soton.ac.uk (Geoff V. Merrett)

proposed for transient computing [9, 10, 11, 12, 13]; however, these all save the entire volatile state without considering which parts of the memory need to be saved. Furthermore, they consider the NVM to be somewhat ideal, whereas in practice the characteristics of different NVM technologies can have a significant impact on efficiency. Using a universal policy, without regard for the NVM technology, results in spending a considerable amount of energy for the retention process. Therefore, the active time of the system is significantly reduced, resulting in degraded forward execution progress of the application. Figure 1 shows the impact of the saving/restoring process on a system with frequent power intermissions. Here, the ratio of active time against the time required to save/restore ( $\frac{t_{\text{active}}}{t_{\text{save/restore}}}$ ) is low while maximising the time spent on useful computation and, therefore, the forward execution progress, is of vital importance for transiently-powered embedded systems.

In this paper, we propose various novel selective policies for efficient state retention which exploit the characteristics of different NVM technologies and match existing policies with the fundamental properties of each NVM technology, to ensure that state retention is an energy and time efficient operation. Key contributions reported are:

- An exploration and analysis of the inefficiency of existing policies and the effect that the properties of different NVM technologies has on each policy;
- Novel policies for efficiently saving state, which select only memory blocks updated since the last save, and reduce erasing by concatenating multiple images;
- An experimental validation of existing and proposed policies on two platforms from different manufacturers with different NVM technologies (Flash and FRAM) as part of a transiently-powered embedded system, to identify the most energy efficient policy/platform combination.

The remainder of this paper is organized as follows. In Section 2, we discuss the problem and motivate the different policies proposed. Section 3 presents an analysis of current policies for saving and restoring the system state. Novel policies along with their analysis and implementations are then described in Section 4, followed by the experimental design in Section 5. Results are presented in Section 6 and, finally, Section 7 concludes the paper.

## 2. Related Work and Motivation

As highlighted in Section 1, various approaches have been proposed to retain system state and enable computation to be sustained upon power failures. An early software-based approach was Mementos [9], which places static trigger-points in strategic locations (e.g. before a function call or inside each loop) at compile time. Mementos saves the core registers, the stack and the global

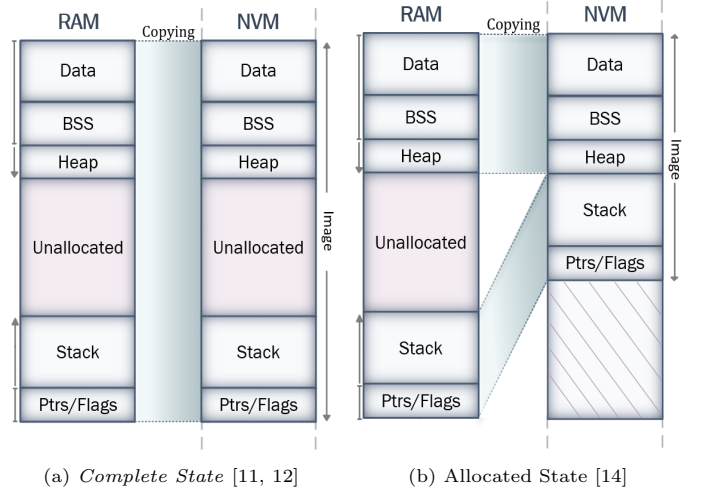


Figure 2: Existing state retention approaches illustrating a) *Complete State* and b) *Allocated State*

variables (part of .bss and .data segments) in Flash memory, captured through analysis at compile time. Furthermore, Mementos often saves the system state even if the power failure is avoided, which results in wasting time and energy. The policy is not generally applicable as it does not address the saving of the heap segment, which is allocated dynamically at run-time.

Hibernus [11] is an interrupt-driven approach that saves the entire system state (main memory, core and general-purpose registers) in Ferroelectric RAM memory (FRAM) and enters a low-power mode when the supply voltage drops below a specific threshold. Hibernus++ [12] is an updated version of the same approach, which dynamically adjusts the saving and restoring thresholds, depending on the on-board decoupling capacitance and the available harvested energy. Figure 2a shows how these approaches can be applied to any system due to their state retention policy which indiscriminately saves and restores the entire RAM memory, including the heap segment and unallocated space. We refer to this policy as *Complete State*, as shown in Figure 2a. However, these approaches do not consider any intelligent policies to identify the unallocated space, introducing a significant amount of time and energy spent on retaining unnecessary data.

To further reduce the time and energy overhead, QuickRecall [10] proposes a unified memory system, replacing the volatile main memory with FRAM. In this case, only the core and general-purpose registers need to be saved in FRAM. However, FRAM is slower and more power-hungry compared to volatile SRAM, making this approach less attractive for low-power embedded systems [15].

The presented software-based approaches save the entire system state every time, without considering what has changed since the last restore. This leads to a sub-optimal state retention process (Hibernus and Mementos) or inefficient use of NVM (QuickRecall). Recently, Bhatti et al. [14] proposed a selective policy for efficient state re-

Prop./Type	Flash <sup>[16]</sup>	FRAM <sup>[17]</sup>	MRAM <sup>[18]</sup>	PCM <sup>[19]</sup>
Read time	70ns	50ns	3-20ns	48ns
Write time	10 $\mu$ s	50ns	3-20ns	150ns
Symmetric	No	Yes	Yes	No
Erase	Yes	No	No	No
Endurance	10 <sup>5</sup>	10 <sup>12</sup>	>10 <sup>15</sup>	10 <sup>12</sup>
Maturity	Established	Production	Production	Testing

Table 1: Comparison of the properties of different NVM technologies

tention which dynamically identifies the unallocated space and only saves to Flash memory the parts of the main memory being used by the application. We refer to this policy as *Allocated State* (see Figure 2b), while the system state being saved in NVM is referred to as an image.

### 2.1. Properties of NVM technologies

To address the challenge of efficiently retaining the system state, we consider the relevant parameters of typical and emerging NVM technologies, which have an impact on the saving and restoring process of a transiently-powered embedded system, as summarised in Table 1.

Flash memory operates by erasing a block of memory cells (page) before writing [16]. The size of this page can range between several bytes to a few kilobytes. A given number of pages form a sector, which is a larger block of memory, and can be erased at the same cost of erasing a single page. Erasing is a highly energy consuming process because it involves generating a voltage pulse using a charge pump. This limits its endurance to approximately  $10^5$  cycles. Moreover, Flash memory is asymmetrical, meaning that reading is faster and more power efficient compared to writing data.

FRAM uses a ferroelectric capacitor as a storage device, offering the benefit of higher power efficiency when compared to Flash memory [17]. However, this technology has a limited lifetime ( $\sim 10^{12}$  cycles), as the ferroelectric material eventually wears out. The read operation with FRAM is destructive because it requires switching the polarisation state to sense the current state. Due to this, the read and write cycles require the same amount of time and energy, thus making FRAM a symmetric memory.

A different symmetric NVM technology is Magnetoresistive RAM (MRAM) that uses a magnetic tunnel junction as a storage device, enabling an unlimited number of read/write cycles [18]. In contrast to FRAM, the read operation is non-destructive, allowing shorter read cycles and improved power efficiency.

Finally, Phase Change Memory (PCM) is an emerging NVM technology which is asymmetric but does not require erasing, offering notably shorter read/write times and significantly higher power efficiency compared to Flash [19]. In addition, erasing is not a prerequisite before writing data which makes PCM an attractive alternative to Flash.

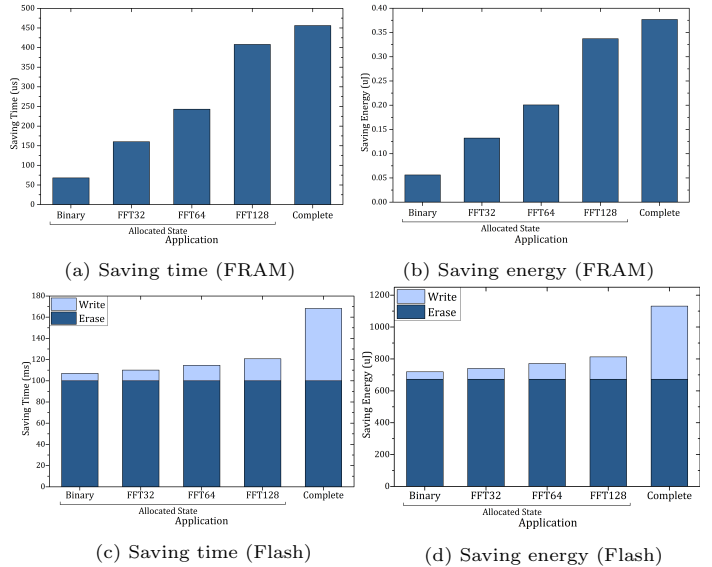


Figure 3: Experimental results showing the time and energy overhead of *Allocated State* with FRAM and Flash memories.

### 2.2. Motivation

To motivate the need for the work presented in this paper, we implemented the *Allocated State* policy (Figure 2b) on two platforms with different NVM technologies (FRAM and Flash). Figure 3 shows results when different software applications (a Binary Counter and FFT) are executed, where the energy required to save the system state is experimentally measured. The *Allocated State* policy allows for substantial energy savings when used with NVM technologies that do not require erasing (such as FRAM). Figures 3a and 3b demonstrate that the cost for saving is proportionally reduced with the size of allocated memory, when compared to saving the entire memory (up to 85% reduction when the Binary Counter application is executed). However, this policy was not validated as part of a transient system and we observe that, when applied on a Flash-based system (as in [14]), it is far less effective as shown in Figures 3c and 3d. This is because the overhead due to the erasing process, that is needed before saving the system state, is neglected in [14]. However, this is a typical property of Flash memory, which accounts for up to 94% of the total cost for saving as highlighted in Figures 3c and 3d. Moreover, we notice that the Figures 3a/3b and 3c/3d are very similar in terms of percentage overhead as the energy required is a function of time. For this reason, we will consider only the energy overhead as the metric of performance for the rest of this paper.

This example motivates the need to apply novel policies depending on the properties of the NVM being used and ensure that each policy is applied on the appropriate NVM technology. In Section 3, we will analyse the inefficiency in existing state retention techniques, while in Section 4, we will consider the properties of NVM technologies (such as symmetry, erasing, endurance and efficiency) to explore different selective policies for system state retention, aiming for more efficient saving/restoring mechanisms and

extended memory lifetime.

### 3. Inefficiencies in Existing Policies

The problem statement and motivation for this work was presented in Section 2, showing that there is a potential for time and energy reductions when saving the system state in transiently-powered embedded systems. In this section, we analyse existing policies to gain insight into the parameters affecting the energy required to save and restore system state. In this way, the factors which play an important role in the saving/restoring process can be determined.

#### 3.1. Complete State

When the *Complete State*, including the entire RAM memory, is saved to a NVM without erase cost, the energy needed to save the state is:

$$E_{\text{Save\_CS}} = M \cdot (P_{\text{RVM}} \cdot t_{\text{RVM}} + P_{\text{WNVM}} \cdot t_{\text{WNVM}}) \quad (1)$$

where  $E_{\text{Save\_CS}}$  is the total energy required for saving the entire system state,  $M$  represents the size of main (volatile) memory in bytes, while  $P_{\text{RVM}}$  and  $t_{\text{RVM}}$  refer to the power and time required to read a byte from volatile memory (VM), and  $P_{\text{WNVM}}$  and  $t_{\text{WNVM}}$  describe the power and time required to save a single byte to NVM. The energy required to restore the system state is given by:

$$E_{\text{Rest\_CS}} = M \cdot (P_{\text{RNVM}} \cdot t_{\text{RNVM}} + P_{\text{WVM}} \cdot t_{\text{WVM}}) \quad (2)$$

In this case, the process is inverted; data is read from NVM and written to Volatile Memory (VM), hence parameters  $P_{\text{RNVM}}$ ,  $t_{\text{RNVM}}$  refer to the power and time required to read a byte from NVM, while  $P_{\text{WVM}}$  and  $t_{\text{WVM}}$  describe the power and time needed to write a byte to the VM. However, when the *Complete State* policy is applied to a system which features a NVM technology with erase cost, Equation (1) becomes:

$$E_{\text{Save\_CS}} = P_{\text{erase}} \cdot t_{\text{erase}} + M \cdot (P_{\text{RVM}} \cdot t_{\text{RVM}} + P_{\text{WNVM}} \cdot t_{\text{WNVM}}) \quad (3)$$

where  $P_{\text{erase}}$  and  $t_{\text{erase}}$  represent the power and time required to erase the NVM.

Table 2 contains typical values of these parameters for three different NVM technologies: FRAM, Flash and PCM. Assuming a platform with a main memory  $M$  equal to 4kB, the energy consumption for saving the system state, when the *Complete State* policy is applied can be estimated as 0.9μJ (FRAM) or 1.1mJ (Flash).

Considering Equations (1)–(3), to reduce the total amount of energy spent for saving or restoring the system state, the state retention policy needs to reduce the amount of data being saved and restored to/from the NVM, and hence reduce the effective  $M$ . For Flash-based systems, the cost for erasing is expected to account for 74% of the total energy overhead (673μJ). This cost is not a function of  $M$  and therefore, to minimise the overhead that the erasing process is contributing, we need to reduce the number of times the NVM needs to be erased.

	FRAM <sup>[17]</sup>	Flash <sup>[17]</sup>	PCM <sup>[19]</sup>	SRAM <sup>[20]</sup>
$P_{\text{erase}}$	–	8mW	–	–
$t_{\text{erase}}$	–	100ms	–	–
$P_{\text{NVM\_R}}$	2mW	6mW	2.5mW	–
$t_{\text{NVM\_R}}$	50ns	70ns	48ns	–
$P_{\text{NVM\_W}}$	2mW	7mW	3mW	–
$t_{\text{NVM\_W}}$	50ns	10μs	150ns	–
$P_{\text{VM\_R}}$	–	–	–	1mW
$t_{\text{VM\_R}}$	–	–	–	10ns
$P_{\text{VM\_W}}$	–	–	–	1mW
$t_{\text{VM\_W}}$	–	–	–	10ns

Table 2: Typical values for reading/writing data from/to different NVM technologies.

#### 3.2. Allocated State Policy

As described in Section 2, the *Allocated State* policy works by distinguishing the memory segments used by the main application, to reduce the amount of data being saved and restored to/from the NVM on every power failure. The amount of allocated memory is defined as:

$$m = \alpha \cdot M$$

where  $\alpha$  is the fraction of the total size  $M$  of the main memory being used. For a system featuring a NVM technology without erase cost, the *Allocated State* policy is described by:

$$E_{\text{Save\_AS}} = P_{\text{track}} \cdot t_{\text{track}} + m \cdot (P_{\text{RVM}} \cdot t_{\text{RVM}} + P_{\text{WNVM}} \cdot t_{\text{WNVM}}) \quad (4)$$

where  $P_{\text{track}}$  and  $t_{\text{track}}$  refer to the power and time required for tracking the location of the end of the heap segment and the top of the stack segment. In contrast, when a NVM technology with erase cost (such as Flash) is used, Equation (4) becomes:

$$E_{\text{Save\_AS}} = P_{\text{erase}} \cdot t_{\text{erase}} + P_{\text{track}} \cdot t_{\text{track}} + m \cdot (P_{\text{RVM}} \cdot t_{\text{RVM}} + P_{\text{WNVM}} \cdot t_{\text{WNVM}}) \quad (5)$$

The energy required to restore the system state when using the *Allocated State* policy depends solely on the size of allocated memory space. Therefore, the energy requirement is given by:

$$E_{\text{Rest\_AS}} = m \cdot (P_{\text{RNVM}} \cdot t_{\text{RNVM}} + P_{\text{WVM}} \cdot t_{\text{WVM}}) \quad (6)$$

Equations (4)–(6) are used with the typical parameter values presented in Table 2 to model the energy consumption of the *Allocated State* policy. However, the cost for tracking the end of the heap segment and the top of the stack ( $P_{\text{track}} \cdot t_{\text{track}}$ ) is neglected as it typically accounts for only a small proportion of the total energy cost. Figures 4




Figure 4: Modelled energy consumption for restoring the system state using the *Allocated State* policy, applied to a system featuring a symmetric memory without erase cost (Equation (4)) and a system featuring an asymmetric memory with erase cost (Equation (6)), using typical parameters of FRAM and Flash from Table 2.

and 5 show the relationship between the energy required to save and restore the system state and the percentage of allocated memory when the *Allocated State* policy is applied. Figure 4 shows the expected effect of this policy both in terms of restoring the system state on two different systems featuring FRAM and Flash memory respectively. For the system with the FRAM memory, the energy overhead is equal for saving and restoring the system state, and it is expected to be reduced by up to 89% when only a small portion of the main memory is used (e.g. 10%), compared to saving the entire system state. The energy overhead for the system featuring Flash memory is more significant as it is a more power-hungry technology. Figure 5 shows the expected energy requirements for saving the system state when the same policy is applied to a system featuring a NVM with erase cost (i.e. Flash). A breakdown of the components contributing to the total cost for saving the system state is shown in Figure 5 which highlights that energy savings of up to 24% can be achieved compared to saving the entire system state (with 10% allocated memory). In this case, the erasing process would account for 74-96% of the total cost for saving the system state. As a consequence, alternative policies need to be devised to tackle the challenge of efficiently saving the system state in transiently-powered embedded systems.

In the following section, a range of policies are proposed which aim to provide a more energy efficient state retention operation, by exploiting the fundamental properties of different NVM technologies.

#### 4. Proposed Policies: Selective State Retention

In this section, we propose a range of policies based on two principles: (a) for NVM technologies with erase cost, we concatenate multiple images and fill NVM before erasing and (b) for asymmetric read/write memory technologies, we save only data that has changed since the last

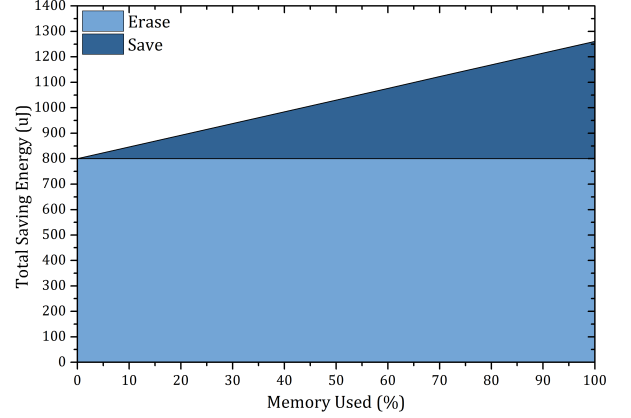


Figure 5: Modelled energy requirement for saving the system state using the *Allocated State* policy, applied to a system featuring NVM with erase cost (Equation (5)), using typical parameters of Flash from Table 2.

restore. The presented policies are developed based on the properties of different NVM technologies and the usage of NVM by the main application.

##### 4.1. Multiple Allocated State Images (MASI)

As presented in Section 2, the *Allocated State* policy works efficiently with NVM technologies that do not need erasing (e.g. FRAM, MRAM). However, this policy does not offer significant benefits when applied to Flash memory that requires erasing before writing. For this reason, we propose the *Multiple Allocated State Images (MASI)* policy as shown in Figure 6, based on concatenating a new image after the previous one, and only erasing when all NVM is filled.

Each image consists of the *.data*, *.bss* and *heap* segments, as well as the *stack* and a dedicated section, containing pointers and flags required for the restore. To identify these segments, the saving process needs to track the end of the heap segment and the top of the stack.

The proposed policy relies on the fact that the size of the NVM is normally multiple times larger compared to the used portion of main memory. For example, some microcontrollers (MCUs) offer up to 32 times more space in their NVM compared to their main memory [21].

To identify the location of the latest image that needs to be restored after a power outage, two variables need to be recorded on every image: (a) the size of the image and (b) a flag indicating whether this image has been previously restored. These variables can be saved at the beginning and the end of the image respectively, as shown in Figure 6. Consequently, during the restore phase, the system first reads the size of the image and then checks whether it has already been restored. If the flag bit is set, it means that a newer image has been stored below and the same procedure will be followed until a cleared flag is detected.

The average energy cost for saving the system state when the presented policy is applied to a system with a






Figure 6: State retention policy of *Multiple Allocated State Images*: to reduce unnecessary erase operations, multiple allocated images are concatenated into NVM until it is full

NVM technology with erase cost, is described by:

$$E_{\text{Save\_MASI}} = \frac{P_{\text{erase}} \cdot t_{\text{erase}}}{i} + P_{\text{track}} \cdot t_{\text{track}} + m \cdot (P_{\text{RVM}} \cdot t_{\text{RVM}} + P_{\text{WNVM}} \cdot t_{\text{WNVM}}) \quad (7)$$

where  $i$  represents the number of saving iterations the system can perform before erasing the NVM while the term  $\frac{P_{\text{erase}} \cdot t_{\text{erase}}}{i}$  describes the average erasing energy.

The energy savings that this policy offers are more evident as the number of images that can be saved in NVM before erasing increases. A smaller amount of allocated memory ( $m$ ) would result in more images being stored in NVM and, as a consequence, the average energy overhead for saving the system state can be effectively reduced.

The energy required to restore the system state is described by:

$$E_{\text{Res\_MASI}} = P_{\text{img}} \cdot t_{\text{img}} + m \cdot (P_{\text{RNVM}} \cdot t_{\text{RNVM}} + P_{\text{WVM}} \cdot t_{\text{WNVN}}) \quad (8)$$

where  $P_{\text{img}}$  and  $t_{\text{img}}$  refer to the power and time required to locate the latest image that needs to be restored, using one of the methods described earlier.

Figure 7 shows the modelled energy consumption for saving the system state when the *Multiple Allocated State Images* policy is applied to a system featuring an asymmetric NVM with erase cost such as Flash (using the values from Table 2). This figure is based on the assumption that the cost for tracking the end of the heap segment and the top of the stack is significantly lower compared to the other components of Equation (7) and can, therefore, be neglected. As the maximum number of iterations ( $i$ ) before erasing decreases, the average erasing energy increases, as shown by the “steps” in the average erase cost. Compared to saving the entire system state, energy savings of up to 95% can be achieved, when the size of allocated memory ( $m$ ) accounts for 10% of the total size of main memory ( $M$ ), according to the model.

A graph showing the energy requirements for restoring the system state can be plotted using Equation (8).




Figure 7: Modelled energy consumption for saving the system state using the *Multiple Allocated State Images* policy, applied to a system featuring NVM with erase cost (Equation (7)), using typical parameters of Flash from Table 2.

However, parameters  $P_{\text{img}}$  and  $t_{\text{img}}$  cannot be estimated accurately as they are also expected to account only for a small portion of the total restore energy, this expression becomes identical to Equation (6). For this reason, Figure 5 can be used as an estimate for the restore process when the *Multiple Allocated State Images* policy is applied.

Using this policy, the energy required for saving the system state can be effectively reduced depending on the size of allocated state which also dictates the number of saving iterations that can be performed without erasing, tackling the high energy cost of erasing the NVM.

The energy consumption of this policy is modelled using a Flash-based system as an example of a memory with a high erase cost. Even though Flash (which is an asymmetric technology) is currently the only available NVM technology exhibiting this property, this policy may also be beneficial on future symmetric memories which suffer from a high erase cost.

#### 4.2. Block-based Policies

The previous policy works efficiently with Flash memory only when the main application is not using a big part of the main memory. When this happens, a small number of images can be saved in the NVM, increasing the frequency that Flash memory needs to be erased. Moreover, when the entire main memory is used, this policy introduces an overhead, due to the time needed to identify the allocated space in the main memory (saving) and to detect the right image to be restored (restoring).

Furthermore, the main memory might contain data that has not been updated since the last restore. In this respect, data already stored in NVM is overwritten with its pre-existing content, resulting in unnecessary write operations. Depending on the NVM technology, these operations can be highly time and energy consuming. To address this challenge, we propose two different selective policies (*Updated Blocks* and *Multiple Updated Blocks*),




Figure 8: State retention policy of *Updated Blocks*: current state compared with previously saved image, and only blocks that have changed are updated

targeted for asymmetric memory technologies, which aim to reduce the amount of redundant writes.

#### 4.2.1. Updated Blocks

As shown in Figure 8, this selective policy is based on dividing the main memory into  $n+1$  number of  $b_n$  blocks ( $b_0$ - $b_n$ ) of size  $s$ . During the saving process, every block of the main memory is compared with the corresponding block previously saved in NVM. If a memory cell has changed since the last restore operation, the data of the whole block needs to be copied to the corresponding block on the NVM. For example, in Figure 8, only the highlighted blocks ( $b_0$ ,  $b_4$ ,  $b_{n-1}$  and  $b_n$ ) changed since the last restore and, therefore, are the only blocks that are updated in NVM during the saving. As a consequence, the unnecessary write operations are avoided and therefore, the time and energy overhead for saving the system state is reduced. However, this policy does not affect the restore process as the entire state needs to be restored to the main memory.

To get a better insight on the operation of the *Updated Blocks* policy, Equation (9) describes the maximum energy required to save the system state when this policy is applied to a system with a NVM without erase cost:

$$E_{\text{Save\_UBMAX}} = M \cdot P_{\text{comp}} \cdot t_{\text{comp}} + B \cdot s \cdot P_{\text{WNVM}} \cdot t_{\text{WNVM}} \quad (9)$$

where  $B$  represents the number of blocks that need to be updated in NVM and  $s$  represents the size of each block in bytes. The comparison between the corresponding blocks in VM and NVM includes a memory access for both memories, so that their values can be read and then compared. Consequently,  $P_{\text{comp}}$  is approximately equal to the sum of  $P_{\text{RVM}}$  and  $P_{\text{RNVM}}$ , while  $t_{\text{comp}}$  is approximately equal to the sum of  $t_{\text{RVM}}$  and  $t_{\text{RNVM}}$ . Equation 9 describes the maximum energy required to save the system state as it is considered that all memory cells in VM are compared with their corresponding cells in NVM. In practice, only a fraction of cells will be compared unless there are no memory changes between two consecutive power failures. This happens because once a cell in VM has been updated,

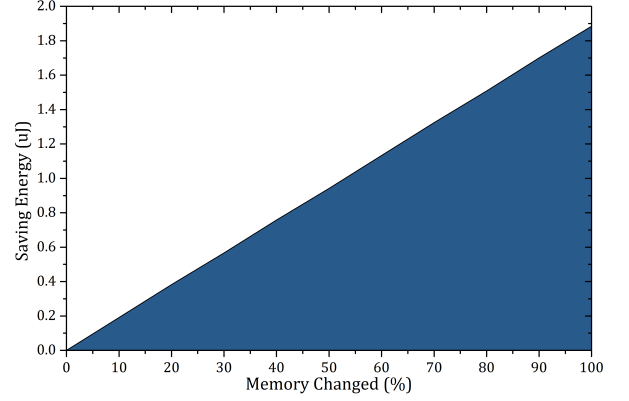


Figure 9: Modelled energy consumption for saving the system state using the *Updated Blocks* policy, applied to a system featuring an asymmetric NVM without erase cost (equation 9), using typical parameters of PCM from Table 2.

the entire block is saved in NVM without comparing the following cells of the block.

The modelled energy consumption for saving the system state using the *Updated Blocks* policy on a system with an asymmetric NVM without erase cost can be estimated using Equation (9). Figure 9 shows the relationship between the expected energy required for saving the system state and the percentage of changed memory, when the *Updated Blocks* policy is applied to a system featuring an asymmetric NVM without erase cost (such as PCM). Here, we conclude that the relationship between the required energy and the number of blocks that need to be updated is of linear nature.

This selective policy is expected to work efficiently with NVM technologies that are asymmetric and do not require erasing, such as PCM. However, it will not have a significant impact with NVM technologies that need erasing (i.e. Flash), due to the high cost of the erasing process. The following policy addresses selectivity with this type of memory.

#### 4.2.2. Multiple Updated Blocks

This policy is based on using the available free space in NVM memory to only save the parts (blocks) of main memory that have changed, using contiguous free space. Similarly to the *Updated Blocks* policy, the main memory is divided into  $n+1$  number of  $b_n$  blocks ( $b_0$ - $b_n$ ). As shown in Figure 10a, the first time a power failure occurs, the available NVM memory is erased and a system state (reference state) is saved. Upon subsequent power outages, each block of the main memory is compared with the corresponding block of the NVM. If a memory cell has changed since the last restore, its updated version is saved in the first available space, without replacing the previous version of the same block in NVM.

As an example, in Figure 10b, only blocks  $b_0$ ,  $b_1$  and  $b_3$  changed since the first power outage, while in Figure 10c, blocks  $b_0$ ,  $b_2$  and  $b_6$  have changed since the second power

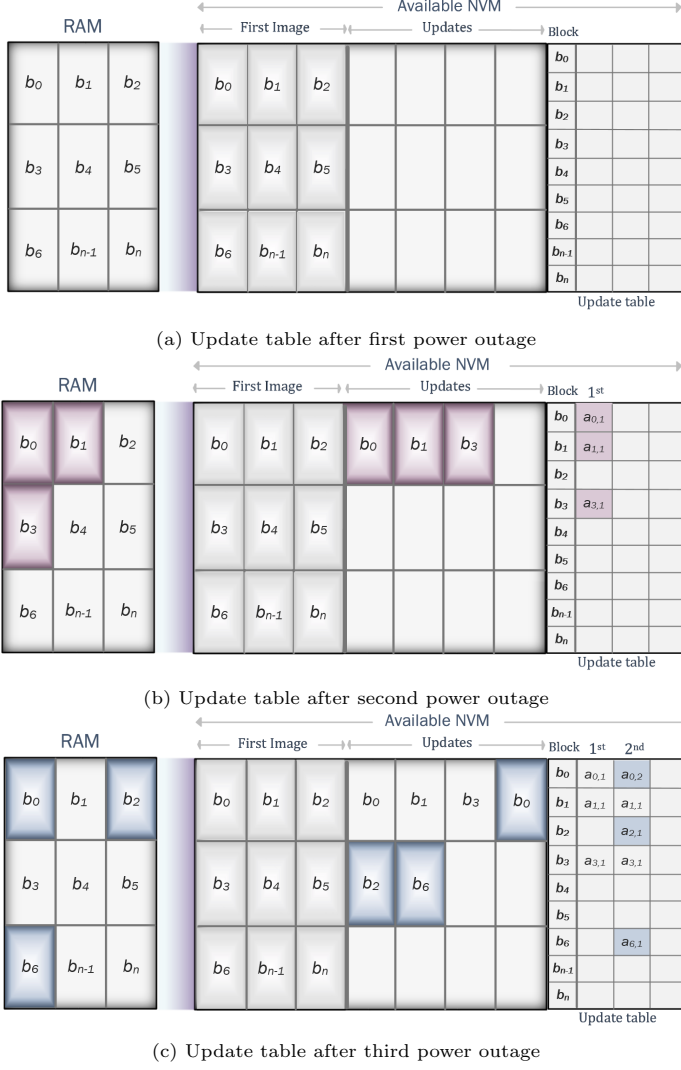


Figure 10: State retention policy of *Multiple Updated Blocks*

outage. If a previously saved block changes again (e.g.  $b_0$  in this case), it has to be saved in NVM. A table is created to keep track of the location in NVM of the latest version of each block ( $a_{n,m}$ , where  $n$  is the block number and  $m$  is the version). The content of this table is used during the restore process to locate the most recent version of each block to be restored to main memory. This table has a fixed number of rows, equal to the number  $n$  of blocks. To avoid erasing the table during every state saving, multiple versions of the table are saved contiguously. To identify the latest table, the area reserved for tables is swept through, until an unwritten cell is located which denotes the end of the most recent table. Once the available space for updates or tables has been filled, the entire NVM is erased and the same process is restarted.

The following equation describes the energy requirements of this policy when applied to a system with NVM

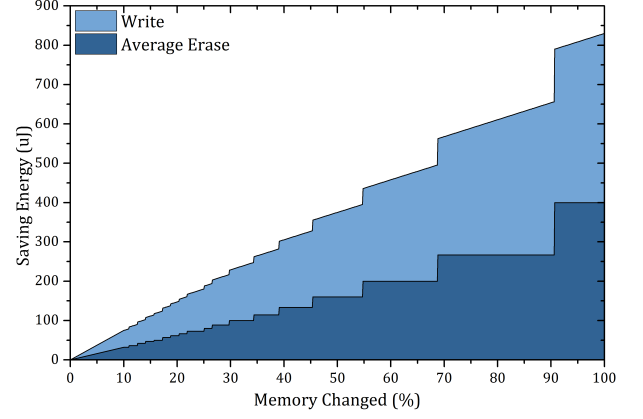


Figure 11: Modelled energy consumption for saving the system state using the *Multiple Updated Blocks* policy, applied to a system featuring an asymmetric NVM with erase cost (equation 11), using typical parameters of Flash from Table 2.

with erase cost:

$$E_{\text{Save\_MUBMAX}} = \frac{P_{\text{erase}} \cdot t_{\text{erase}}}{i} + M \cdot P_{\text{comp}} \cdot t_{\text{comp}} + B \cdot s \cdot P_{\text{WNVM}} \cdot t_{\text{WNVM}} + P_{\text{table}} \cdot t_{\text{table}} \quad (10)$$

where  $P_{\text{table}}$  and  $t_{\text{table}}$  describe the power and time required to create the latest version of the table.

Equation 10 can be useful for plotting a Figure to show the relationship between the energy required for saving the system state and the fraction of memory that has changed since the last restore. However, estimating the overhead for creating the table is difficult while it is expected to introduce only a small overhead compared to the other components of Equation (10). For this reason, Equation (11) is a simplified version of this equation which will be used for modelling purposes:

$$E_{\text{save\_MUB}} = \frac{P_{\text{erase}} \cdot t_{\text{erase}}}{i} + M \cdot P_{\text{comp}} \cdot t_{\text{comp}} + B \cdot s \cdot P_{\text{WNVM}} \cdot t_{\text{WNVM}} \quad (11)$$

Figure 11 shows the relationship between the expected energy required for saving the system state and the percentage of changed memory, when the *Multiple Updated Blocks* policy is applied to a system featuring an asymmetric NVM with erase cost such as Flash. To plot this graph, a main memory size of 4kB and a total NVM size of 16kB are considered. In addition, it is assumed that memory is changing in a contiguous way for simplicity reasons. It is shown that the erasing cost can be significantly reduced compared to erasing on every iteration (Figure 5) and energy savings of up to 94% can be achieved, when only 10% of the memory has changed. Finally, the instantaneous “steps” observed at the average erase energy are related to the decrease in the maximum number of iterations before erasing the NVM.

When the system state needs to be restored, the latest table needs to be identified so that the most recent version of each block can be located. The following equation can






Figure 12: Schematic of the experimental setup.

be used to calculate the expected energy requirements for the restore process:

$$E_{\text{Rest\_MUB}} = P_{\text{loc}} \cdot t_{\text{loc}} + M \cdot P_{\text{RNVM}} \cdot t_{\text{RNVM}} + M \cdot P_{\text{WVM}} \cdot t_{\text{WVM}} \quad (12)$$

where the product of  $P_{\text{loc}}$  and  $t_{\text{loc}}$  describe the energy required to locate the latest table upon restore.

The Allocated State policy performs well when applied on a system with a symmetric memory without erase cost such as FRAM, as explained in Section 2. We propose *Multiple Allocated State Images* and *Multiple Updated Blocks* which are designed to reduce the cost of erasing. The former can be applied both on symmetric and asymmetric memories with erase cost, while the latter is focused on asymmetric memories, which suffer from high erase cost. Finally, the *Updated Blocks* policy is designed to reduce redundant write operations, offering great potential for systems featuring an asymmetric memory without erase cost.

## 5. Experimental Design

The proposed policies were experimentally validated using two platforms, a Texas Instruments MSP430FR with FRAM memory [22], and a NXP LPC812 /colorredplatform [23] with Flash memory. The MSP430FR with FRAM allows read/write operations at the byte level, whereas the LPC812 with Flash memory is at the page level (64 bytes). The block diagram of the experimental setup is shown in Figure 12. Here, an energy harvester is used as the energy source of the system. The decoupling capacitance of the system (approximately  $16\mu\text{F}$ ) is adequate for saving/restoring the system state using the MSP430 platform with FRAM memory) while an additional capacitor (C1) is required for the Flash-based platform (LPC812) in order to provide enough energy/time for the state to be saved to Flash. The external low-power comparator used in [12] is used to enable interrupts to be triggered when the supply voltages surpasses a threshold ( $V_{\text{th}}$ ) set by the microcontroller. The microcontroller starts its state retention/restore operation when the corresponding interrupt has been triggered by the comparator.

To implement the *Multiple Allocated State Images* policy, the end of the heap segment and the top of the stack need to be identified. To make this feasible with the available hardware (LPC812), a combination of *malloc()* and




Figure 13: uBenchmark's configurable parameters

*free()* functions is used to locate the end of the heap segment which is preceded by the *.data* and *.bss* segments, as shown in Figure 2a. When the combination of these functions is executed, a general purpose register contains the address of the end of the heap segment and can, therefore, be copied to NVM. Subsequently, the top of the stack can be obtained in a similar manner, by saving the value of the stack pointer (SP).

In addition, identifying the latest image using the flagging methodology proposed in Section 4.1 is infeasible, as it only works for Flash memories that allow writing of a single byte. However, the minimum writeable size is equal to the page size for this platform (as well as many Flash-based microcontrollers). Consequently, a different method needs to be used as it would not be feasible to update the value of the flag once the image has been restored. In this case, the memory is swept through until an empty page is found, which denotes the latest image. This is done by exploiting the attribute of Flash memory which implies that when a page is erased, its cells are set to logic level "1". Consequently, as images are stored in a continuous way, the first empty page reveals the ending address of the most recent image. Once the latest image has been located, the pointers/flags section is used so that the restore process can be executed. This method, however, introduces a small overhead that gradually increases, depending on the total size of the previous images.

The proposed policies were evaluated using a custom application (uBenchmark), which allows us to define the percentage of allocated memory at compile time. Moreover, it enables to define a portion of the allocated memory where the data is randomly changed. As shown in Figure 13, three parameters are used to define this allocated space ( $\gamma$ ) as well as the boundaries of the randomly changed section ( $\alpha$  and  $\beta$ ).

For the results presented in Section 6, two different cases were considered. First, to allow plotting of the "Memory Used" (plotted on the x-axis of some results presented),  $\gamma$  is varied in order to change the portion of the allocated memory, while the values of  $\alpha$  and  $\beta$  are constant. When "Memory Changed" needs to be plotted,  $\alpha$  is fixed to 0 (start of the main memory),  $\gamma$  is equal to the size of the

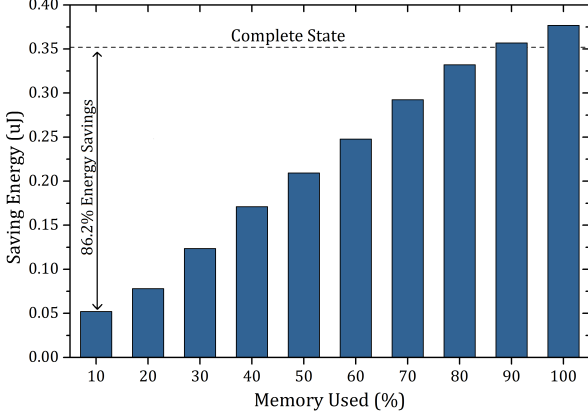


Figure 14: Experimental results showing the energy required to save the system state, when using the *Allocated State* policy on FRAM.

main memory and  $\beta$  is varied to adjust the percentage of the memory being changed.

The proposed selective policies were validated using HiBernus [11]. This transient approach was selected because it is energy efficient and, application and platform agnostic [15]. In the following section, these policies are validated considering symmetry and erase cost as the main properties of NVM technologies.

## 6. Experimental Results

### 6.1. Symmetric NVM Technologies without Erase Cost

For symmetric memory technologies (e.g. FRAM), comparing the already saved with the current system state would only have a negative impact on the efficiency of the saving process since the cost of reading from the NVM is equal to writing. Therefore, the *Updated Blocks* approach is unsuitable. In addition, the *Multiple Allocated State Images* and *Multiple Updated Blocks* policies would not offer any benefits as this type of NVM technology does not require erasing before writing. For this reason, only the *Allocated State* policy is experimentally validated.

Figure 14 shows the amount of energy required for saving the system state while changing the fraction of allocated memory, when the *Allocated State* policy is applied to the MSP430FR with FRAM. Depending on the fraction of memory used by the main application, energy savings of up to 86.2% can be achieved compared to saving the entire system state. As an example, when an application uses a small portion of the main memory such as 10%, the saving/restoring process requires 52nJ of energy to be completed. This is due to the overhead incurred by the tracking of the end of the heap segment and the top of the stack, as described in Section 3.2. This policy is more energy efficient compared to the *Complete State* policy as long as <88.4% of the memory is being used.

### 6.2. Asymmetric NVM Technologies with Erase Cost

This section presents the experimental results of various policies when applied to systems featuring asymmetric

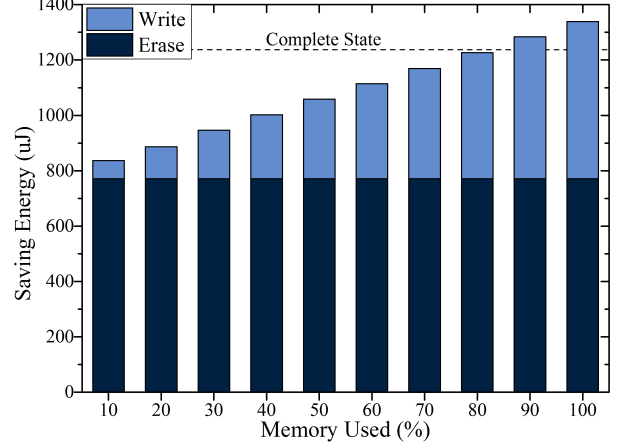


Figure 15: Experimental results showing the energy required to save the system state, when using the *Allocated State* policy on Flash memory.

NVM Technologies *with* erase cost and results in a discussion which compares the performance of the applied policies.

#### 6.2.1. Allocated State Policy

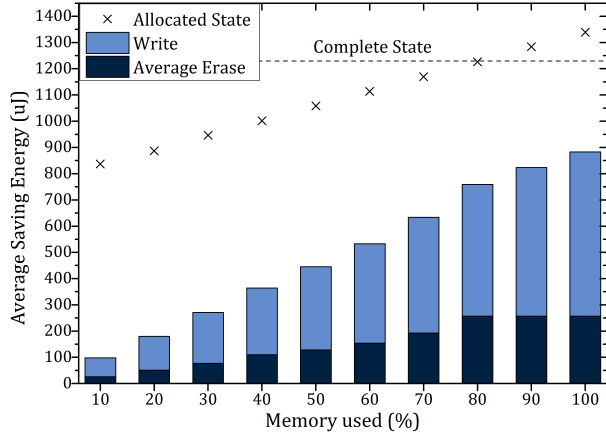
As shown in Figure 15, when the *Allocated State* policy is applied to the LPC812 with Flash memory, the energy overhead due to the erasing is dominant (771μJ).

However, this policy is still working efficiently compared to saving the complete state (shown by the dashed line), when the allocated memory is less than 82.3%. This is due to the overhead of having to track the end of the heap segment and the top of the stack so that only the allocated space is saved in NVM which counteracts with the benefits of this policy. Moreover, when the memory usage is small (e.g. 10%) the energy spent for the saving process alone (i.e. ignoring the erase overhead) is reduced by up to 84%.


In this case, the model for the *Allocated State* policy (presented in Section 3.2) is validated as Figures 5 and 15 show almost identical behaviour with an average error of less than 4%.

#### 6.2.2. MASI

Figure 16a shows the energy required by the *Multiple Allocated State Images* policy when applied to the same platform. This policy offers significantly higher energy efficiency as the number of saved images without erasing increases, depending on the percentage of allocated memory. This is due to the substantial energy reduction for erasing, considering that the cost of erasing is spread across the number of saving iterations that can be performed with a single erase. For ease of comparison, the X points show the energy requirements of the *Allocated State* policy while the dashed line shows the default *Complete State* policy. When using *Multiple Allocated State Images*, the energy overhead is reduced by 32.3-87.3% compared to *Allocated State*. Figure 16b shows the maximum number of images



(a) Energy requirement for saving the system state



(b) Number of saved images before erasing against percentage of used memory

Figure 16: Experimental results showing (a) the energy required to save the system state and (b) the number of iterations, when using the *Multiple Allocated State Images* policy on Flash memory.

that can be saved in NVM before erasing, depending on the memory usage. For this specific platform, the size of main memory is 4kB, while the available Flash memory is 12kB, as the .text segment has a static size of 4kB (total 16kB). As an example, for an application that has an allocated memory size equal to 10% of the total space, 30 images can be saved before erasing, leading to an average erasing energy overhead of  $26.7\mu\text{J}$ .

Comparing the experimental results (Figure 16a) with the modelled version of this policy (Figure 7), we observe that there is an average error of approximately 14%. As explained in Section 4.1, the cost for tracking the end of the heap segment and the top of the stack is not included and therefore, it is expected that the modelled energy consumption of this policy would be lower compared to the experimental results.

Figure 17 shows the energy needed to restore the image as a function of the number of images saved in NVM, for different values of memory usage (10-50%) when the *Multiple Allocated State Images* policy is applied. Here, a energy overhead can be seen that gradually increases with

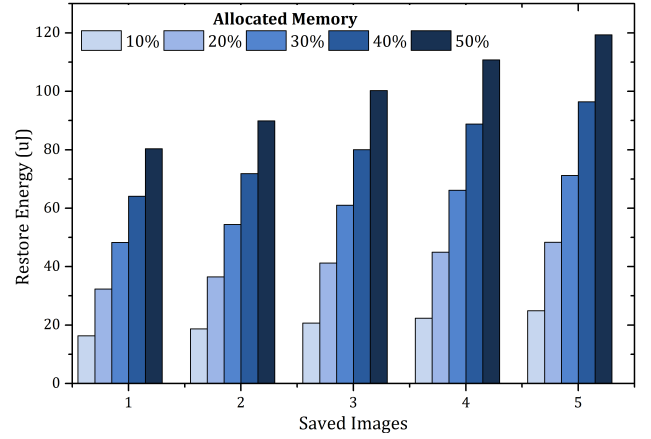


Figure 17: Experimental results showing the restore energy against the number of stored images in NVM, when using the *Multiple Allocated State Images* policy on Flash memory.

the total size of the memory occupied by the previously saved images. This is due to the process of locating the latest image to be restored, as described in Section 4.1.

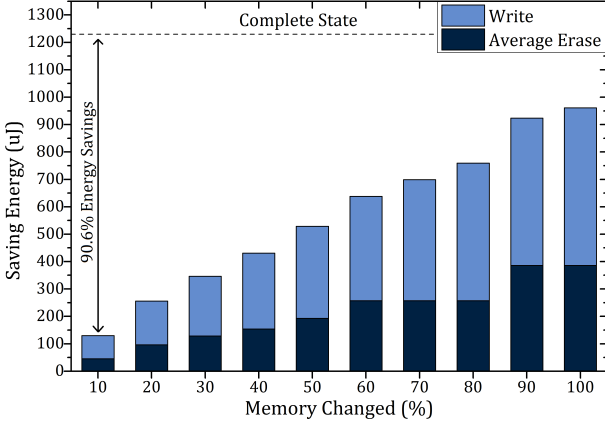
### 6.2.3. MUB

Figure 18a shows the energy overhead when the *Multiple Updated Blocks* policy is applied to this platform. In this case, the block size ( $b_n$ ) is equal to the page size (64 Bytes), which results in the main memory being divided into 64 blocks. This policy offers significantly higher time saving between 34.9% and 90.2%, when compared to saving the entire system state. Similarly, the energy overhead is reduced between 37.4% and 90.6% when 10% of the main memory is being used. This confirms that *Multiple Updated Blocks* is an energy efficient policy for Flash memory.

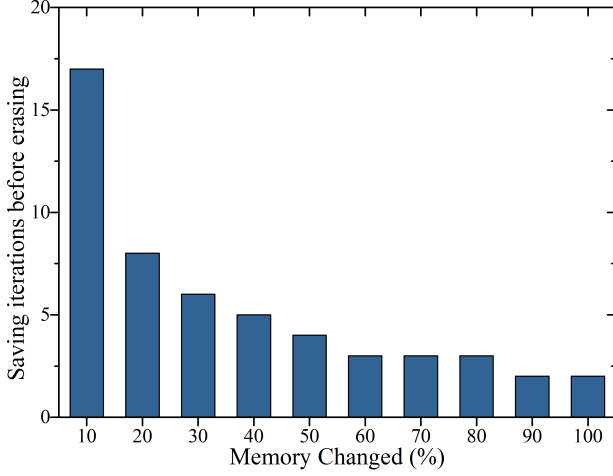
Figure 18b shows the number of state retention iterations that can be performed depending on the percentage of contiguously changed memory. For this platform, assuming a .text segment size of 4kB, the system is able to save the necessary data up to 17 times before erasing.

### 6.2.4. Discussion

Comparing the results presented in this section, we conclude that, while the *Allocated State* policy (Subsection 6.2.1) is the least efficient policy for a platform featuring an asymmetric NVM technology with erase cost such as Flash, it performs better than the *Complete State* policy as long as  $<82\%$  of main memory is allocated according to Figure 15. In addition, it is observed that unless the entire portion of allocated memory has changed between two power intermissions, the *Multiple Updated Blocks* policy (Subsection 6.2.3) offers higher energy efficiency compared to the *Multiple Allocated State Images* policy (Subsection 6.2.2). The only exception to this case is when the percentage of allocated memory is  $<10\%$ , where saving the allocated memory without considering the memory changes is more effective.



(a) Energy requirement for saving the system state



(b) Number of state retention iterations before erasing

Figure 18: Experimental results showing (a) the energy required to save the system state and (b) the number of iterations, when using the *Multiple Updated Blocks* policy on Flash memory.

### 6.3. Asymmetric NVM Technologies without Erase Cost

This section presents the experimental results of various policies when applied to systems featuring asymmetric NVM Technologies *without* erase cost and results in a discussion which compares the performance of the applied policies.

#### 6.3.1. Updated Blocks

The *Updated Blocks* policy is not appropriate for Flash memory, as the cost for erasing is prohibitive. However, for asymmetric memories without an erasing cost (e.g. PCM), it can be advantageous. However, currently there are no systems commercially available featuring a PCM memory to allow experimental evaluation of this policy. To illustrate this using the available hardware, we use Flash but negate the cost of erasing so that the benefits of this policy can be extracted. Figure 19 shows the energy required to perform a complete system state retention using the *Updated Blocks* policy. This policy is more efficient when compared to saving the entire system state, if  $<80.7\%$  of main memory has changed since the last restore, as shown




Figure 19: Experimental results showing the energy required to save the system state, when using the *Updated Blocks* policy on asymmetric NVM without erase cost. Policy validated using the available hardware as a proof of concept (Flash, negating erase cost).

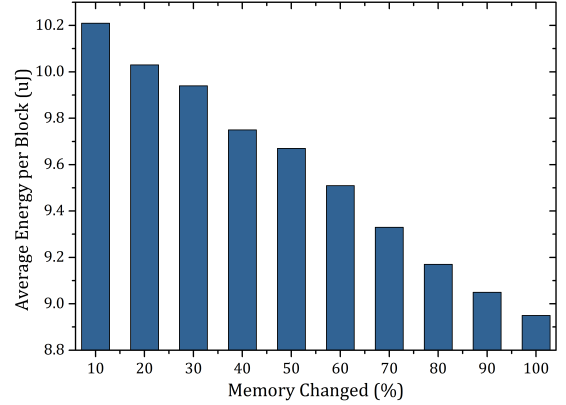


Figure 20: Experimental results showing the average energy required to save each block against percentage of memory contiguously changed. Policy validated using the available hardware as a proof of concept (Flash, negating erase cost).

by the dashed line in Figure 19. As these results have been obtained using a Flash-based system by negating the erase cost, they cannot be considered conclusive for PCM. However, comparing the properties in Table 2, we observe that the saving energy (y-axis) would scale by approximately two orders of magnitude, as PCM is more energy efficient compared to Flash (as confirmed by comparing Figure 9 and Figure 5). Also, the maximum energy savings would be lower, as PCM is a less asymmetric technology (the relative difference between read and write energy is lower), and therefore the effect of the policy is less evident. The overall behaviour, however, would still be observed and these results can be used as a proof of concept, rather than a quantitative result regarding the actual energy savings that this policy could offer to a system featuring a PCM memory.

Figure 20 shows the average energy required to save a block as a function of the number of updated blocks. As the amount of energy needed to write a block is constant ( $8.2\mu\text{J}$ ), the extra cost comes from the control mechanism needed for reading, comparing and updating each block,

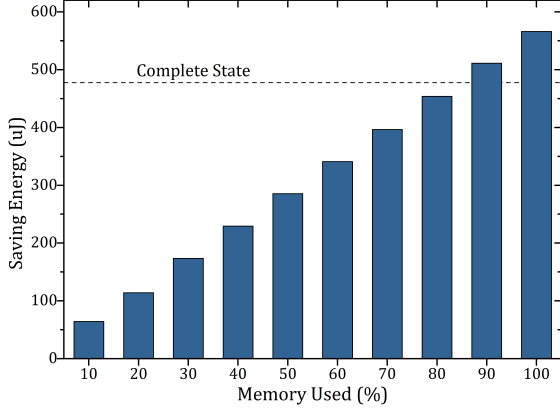


Figure 21: Experimental results showing the energy required to save the system state using the *Allocated State* policy on asymmetric NVM without erase cost. Policy validated using the available hardware as a proof of concept (Flash, negating erase cost).

as described in Section 4.2. Here, for a smaller number of updated blocks, the effect of the control mechanism is more evident whereas it is reduced when all blocks of the main memory have been updated and saved in NVM. The more blocks that need to be updated, the fewer comparisons need to be made between the current and previously saved system state.

### 6.3.2. Allocated State

As mentioned earlier, the *Allocated State* policy does not offer significant benefits when used with Flash. However, it can be a beneficial policy for asymmetric memories that do not require erasing. To demonstrate the effect of this policy on this type of NVM using the available platform (LPC812), we disregard the energy cost for erasing, focusing on the overhead for writing. Figure 21 shows the energy required to perform a system state retention using the *Allocated State* policy. When less than 84.7% of the main memory is used by the main application, this policy is more efficient when compared to saving the entire system state, as shown by the dashed line in Figure 21.

### 6.3.3. Discussion

Comparing Figures 19 and 21, we observe that unless the entire allocated state has changed since the last power failure, the *Updated Blocks* policy is more energy efficient compared to the *Allocated State* policy. However, the latter should be preferred when the frequency of power failures is low and, therefore, the likelihood that the vast majority of the memory cells have been updated is increased.

### 6.4. Summary

Table 3 summarises the different selective policies and their applicability to different NVM technologies. These have been extracted from the experimental results, considering the properties of the most important parameters of NVM technologies (symmetricity, erasing). For symmetric NVM technologies, only the *Allocated State* (for memories without erase cost) and *Multiple Allocated State*

Property / Policy	Allocated MASI State	Updated MUB Blocks
<b>Symmetric</b>		
No erase cost	✓	
Erase cost		✓
<b>Asymmetric</b>		
No erase cost	✓	✓
Erase cost		✓

Table 3: Comparison of proposed policies with regard to NVM properties (symmetricity, erase cost).

Memory Usage	Low Allocation	High Allocation
Few Updates	Upd.Blocks MUB	Upd.Blocks MUB
Entire Update	Alloc. State MASI	Alloc. State MASI

Table 4: Comparison of proposed policies with regard to memory usage (updates, allocation).

*Images* (for memories with erase cost) policies should be considered, as the overhead for comparing the VM with the NVM incurred by the other policies is prohibitive. For systems featuring an asymmetric NVM, the way the main memory is used by the application significantly affects the performance of each policy. Table 4 summarises the most suitable policies depending on the memory usage. When the entire allocated memory is updated, the *Allocated State* (asymmetric NVM without erase cost) and *Multiple Allocated State Images* (asymmetric NVM with erase cost) policies offer better energy efficiency, as the cost for comparing VM with its corresponding NVM blocks is eliminated. However, when fewer memory cells have been updated since the last power failure, the *Updated Blocks* (asymmetric NVM without erase cost) and *Multiple Updated Blocks* (asymmetric NVM with erase cost) policies perform better compared to the other policies.

## 7. Conclusions

In this paper, we have shown the inefficiency of current state retention policies in transiently-powered embedded systems when used on NVM technologies with certain properties. We presented novel selective policies for efficient state retention in order to identify the most energy efficient policy/platform combination. These software-based retention policies are targeted for different NVM technologies, exploiting their properties such as read/write symmetricity and the need for erasing. Unlike existing approaches, the proposed policies are based on the principles of (1) saving only the information that has changed since the last restore (*Updated Blocks* and *Multiple Updated Blocks*), and (2) avoiding the cost of erasing NVM by concatenating multiple images (*Multiple Allocated State*



*Images* and *Multiple Updated Blocks*). The existing and proposed policies were experimentally validated on two different and appropriate platforms (featuring Flash and FRAM memory). A comparison between different policies has been made, considering the effect of how the application is using memory. From this, we determine the most energy-efficient policy, based on the characteristics of the NVM technology of the system. Results show that using the appropriate policy/platform combination provides a reduction in the energy overhead of up to 86.2% for FRAM (using the *Allocated State* policy, Figure 14) and 90.6% for Flash memory (using the *Multiple Updated Blocks* policy, Figure 18a), compared to the typical approach of saving the entire system state.

## Acknowledgements

This work was supported in part by the UK Engineering and Physical Sciences Research Council (EPSRC) under Platform Grant EP/P010164/1 and the PRiME Programme Grant EP/K034448/1. Experimental data used in this paper can be found at DOI:10.5258/SOTON/D0590 (<https://doi.org/10.5258/SOTON/D0590>).

## References

- [1] J. Olivo, S. Carrara, and G. De Micheli. Energy harvesting and remote powering for implantable biosensors. *IEEE Sensors Journal*, 11(7):1573–1586, 2011.
- [2] P. D. Mitcheson. Energy harvesting for human wearable and implantable bio-sensors. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, pages 3432–3436. IEEE, 2010.
- [3] S. Kahrobaee and M. C. Vuran. Vibration energy harvesting for wireless underground sensor networks. In *Communications (ICC), 2013 IEEE International Conference on*, pages 1543–1548. IEEE, 2013.
- [4] D. Balsamo, A. Das, A. S. Weddell, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. Graceful performance modulation for power-neutral transient computing systems. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, 35(5):738–749, 2016.
- [5] S. Beeby and N. White. *Energy harvesting for autonomous systems*. Artech House, 2010.
- [6] F. Ongaro, S. Saggini, and P. Mattavelli. Li-ion battery-supercapacitor hybrid storage system for a long lifetime, photovoltaic-based wireless sensor network. *IEEE Transactions on Power Electronics*, 27(9):3944–3952, 2012.
- [7] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems (TECS)*, 6(4):32, 2007.
- [8] K. Ma, Y. Zheng, S. Li, K. Swaminathan, X. Li, Y. Liu, J. Sampson, Y. Xie, and V. Narayanan. Architecture exploration for ambient energy harvesting nonvolatile processors. In *High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium on*, pages 526–537. IEEE, 2015.
- [9] B. Ransford, J. Sorber, and K. Fu. Mementos: System support for long-running computation on rfid-scale devices. *ASPLOS’11*, 2011.
- [10] H. Jayakumar, A. Raha, and V. Raghunathan. Quickrecall: A low overhead hw/sw approach for enabling computations across power cycles in transiently powered computers. *27th International Conference on VLSI Design and 2014 13th International Conference on Embedded Systems*, 2014.
- [11] 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. *IEEE Embedded Syst. Lett.*, 2015.
- [12] D. Balsamo, A. S. Weddell, A. Das, A. R. Arreola, D. Brunelli, B. M. Al-Hashimi, G. V. Merrett, and L. Benini. Hibernus++: a self-calibrating and adaptive system for transiently-powered embedded devices. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(12):1968–1980, 2016.
- [13] Y. Liu, J. Yue, H. Li, Q. Zhao, M. Zhao, J. Xue, G. Sun, M.-F. Chang, and H. Yang. Data backup optimization for nonvolatile sram in energy harvesting sensor nodes. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2017.
- [14] N. A. Bhatti and L. Mottola. Efficient state retention for transiently-powered embedded sensing. In *EWSN*, pages 137–148, 2016.
- [15] A. Rodriguez, D. Balsamo, A. Das, A. S. Weddell, D. Brunelli, B. Al-Hashimi, and G. V. Merrett. Approaches to transient computing for energy harvesting systems: A quantitative evaluation. In *ENSSys 2015*, 2015.
- [16] P. Pavan, R. Bez, P. Olivo, and E. Zanoni. Flash memory cells—an overview. *Proceedings of the IEEE*, 85(8):1248–1271, 1997.
- [17] Texas Instruments. *FRAM New Generation of Non-Volatile Memory*, 9 2009.
- [18] J. S. Meena, S. M. Sze, U. Chand, and T.-Y. Tseng. Overview of emerging nonvolatile memory technologies. *Nanoscale research letters*, 9(1):526, 2014.
- [19] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger. Architecting phase change memory as a scalable dram alternative. In *ACM SIGARCH Computer Architecture News*, volume 37, pages 2–13. ACM, 2009.
- [20] A. J. Bhavnagarwala, X. Tang, and J. D. Meindl. The impact of intrinsic device fluctuations on cmos sram cell stability. *IEEE journal of Solid-state circuits*, 36(4):658–665, 2001.
- [21] Texas Instruments. *MSP430F15x, MSP430F16x, MSP430F161x Mixed Signal Microcontroller*, 3 2011. Rev. G.
- [22] Texas Instruments. *MSP430FR573x Mixed-Signal Microcontrollers*, 4 2016. Rev. K.
- [23] NXP Semiconductors. *LPC81x User manual*, 4 2014. Rev. 1.6.