

HaVOC: A Hybrid Memory-aware Virtualization Layer for On-Chip Distributed ScratchPad and Non-Volatile Memories

Luis Angel Bathen; Nikil Dutt

Center for Embedded Computer Systems
University of California, Irvine
Irvine, CA 92697, USA
{lbathen, dutt}@uci.edu

ABSTRACT

Hybrid on-chip memories that combine Non-Volatile Memories (NVMs) with SRAMs promise to mitigate the increasing leakage power of traditional on-chip SRAMs. We present HaVOC: a run-time memory manager that virtualizes the hybrid on-chip memory space and supports efficient sharing of distributed ScratchPad Memories (SPMs) and NVMs. HaVOC allows programmers and the compiler to partition the application's address space and generate data/instruction block layouts considering virtualized hybrid address spaces. We define a data volatility metric used by our hybrid memory-aware compilation flow to generate memory allocation policies that are enforced at run-time by a filter-inspired dynamic memory algorithm. Our experimental results with a set of embedded benchmarks executing simultaneously on a Chip-Multiprocessor with hybrid NVM/SPMs show that HaVOC is able to reduce execution time and energy by 60.8% and 74.7% respectively with respect to traditional multitasking based SPM allocation policies.

Categories and Subject Descriptors

C.3 [Special-purpose and Application-based systems]: Real-time and embedded systems; B.3 [Design Styles]: Virtual Memory; D.4 [Storage Management]: Distributed memories

General Terms

Algorithms, Design, Management, Performance

1. INTRODUCTION

The ever increasing complexity of embedded software and adoption of open-environments (e.g., Android) is exacerbating the deployment of multi-core platforms with distributed on-chip memories [12, 18]. Traditional memory hierarchies consist of caches, however, it is known that caches may consume up to 50% of the processor's area and power [3]. As a result, ScratchPad Memories (SPMs) are rapidly being adopted and incorporated into multi-core platforms for their high predictability, low area and power consumption. Efficient SPM management can significantly reduce power con-

sumption [17, 29, 13], and may be a good alternative to caches for applications with high levels of regularity (e.g., Multimedia).

As sub-micron technology continues to scale, leakage power will overshadow dynamic power consumption [19, 14]. Since SRAM-based memories consume a large portion of the die, they are a major source of leakage in the system [2], which is a major issue for multi-core platforms. In order to reduce leakage power in SRAM-based memories, designers have proposed emerging Non-Volatile Memories (NVMs) as alternatives to SRAM for on-chip memories [27, 15, 24]. Typically, NVMs (e.g., PCRAM [20]) offer high densities, low leakage power, comparable read latencies and dynamic read power with respect to traditional embedded memories (SRAM/e-DRAW). One major drawback across NVMs is the expensive write operation (high latencies and dynamic energy). To mitigate the drawbacks of the write operation in NVMs, designers have made the case for deploying hybrid on-chip memory hierarchies (e.g., SRAM, NVM) [27], which have shown up to 37% reduction in leakage power [11], and increased IPC as a byproduct of the higher density provided by NVMs [24]. Orthogonal to traditional hybrid on-chip memory subsystems which have been predominately focused on caches, Hu et al. [11] showed the benefits of exploiting hybrid memory subsystems consisting of SPMs and NVMs.

In this paper, we present *HaVOC*, a system-level HW/SW solution to efficiently manage on-chip hybrid memories consisting distributed ScratchPad Memories (SRAM) and Non-Volatile Memories (e.g., MRAMs) to support multitasking Chip-Multiprocessors. *HaVOC* allows programmers to partition their application's address space into virtualized SRAM address space and virtualized NVM address space through a minimalistic API. Programmers (through annotations) and compilers (through static analysis) can then specify hybrid memory-aware allocation policies for their data/instruction blocks at compile-time, while *HaVOC* dynamically enforces them and adapts to the underlying memory subsystem. The **novel contributions** of our work are that we:

- Explore distributed shared on-chip hybrid memories consisting of SPMs and NVMs and virtualize their address spaces to facilitate the management of their physical address spaces
- Introduce the notion of data volatility analysis to drive efficient compilation and policy generation for hybrid on-chip memories
- Present a filter-driven dynamic allocation algorithm that exploits filtering and volatility to find the best memory placement

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2012 June 3-7, 2012, San Francisco, California, USA
Copyright 2012 ACM 978-1-4503-1199-1/12/06 ...\$10.00.




Figure 1: HaVOC-aware Policy Generation (a) and Enforcement (b).

2. MOTIVATION

Unlike caches, SPMs are software controlled memories as their management is completely left to the programmer and compiler. At first glance, we would need to take traditional SPM management schemes and adapt them to manage hybrid memories. However, SPM based allocation schemes (e.g., [29, 13]) assume physical access to the memory hierarchy; consequently, the traditional SPM based programming model would require extensive changes to account for the different characteristics of the NVMs. This motivates the need for a simplified address space to minimize changes to the SPM programming model.

The challenge of programming and managing SPM/NVM-based hybrid memories is aggravated by the adoption of open environments (e.g., Android OS), where users can download applications, install them, and run them on their devices. In these environments, it is possible that many of the running processes will require access to the physical SPMs, therefore, programmers and compilers can no longer assume that their applications are the only ones running on the system. Traditional SPM-sharing approaches [9, 26, 28] would either allocate part of the physical address space to each process (spatial allocation) or time-share the SPM space (temporal allocation). Once the entire SPM space has been allocated, all remaining data is then mapped to off-chip memory. In order to reduce the overheads of sharing the SPM/NVMs, our scheme exploits programmer/compiler-driven policies obtained through static analysis/annotations (Sec. 3.3) and uses the information to efficiently manage the memory resources at run-time (Sec. 3.4).

3. HAVOC OVERVIEW

Figure 1 (a) shows our proposed compilation flow, which takes annotated source code, and performs various types of SPM/NVM-aware static analysis techniques (e.g., instruction placement, data reuse analysis, data volatility analysis); the compiler then uses this information to generate allocation policies assuming the use of virtual SPMs (vSPMs) and virtual NVMs (vNVMs). Figure 1 (b) shows our proposed dynamic policy enforcement mechanism for multitasking CMPs. The *HaVOC* manager (black box) takes in the vSPM/vNVM allocation policies provided by each application (Application 1 & 2), and decides how to best utilize the underlying memory resources. The rest of this section will go over each of the different components at a high level, for more details please refer to our technical report [4]. In the following discussion we use data/instruction blocks interchangeably as our approach supports placement of both data and instructions onto the hybrid memories.

3.1 Target Platform and Assumptions

Figure 2 shows a high level diagram of our SPM/NVM




Figure 2: HaVOC-Enhanced CMP.

enhanced CMP, which consists of a number of OpenRISC-like cores, the *HaVOC* manager, a set of distributed SPMs and NVMs, a DRAM/NVM main memory hierarchy, and an AMBA AHB bus-based communication fabric.

We make the following assumptions: 1) The application can be statically analyzed/profiled so that data/instruction blocks can be mapped to SPMs [12, 13, 28]. 2) We operate over blocks of data (e.g., 1KB mini-pages). 3) We can map all data/instructions to on-chip/off-chip memory and do not use caches (e.g., [12]). 4) Part of off-chip memory can be locked in order to support the virtualization of the on-chip SPM/NVM memories.

3.2 Virtual Hybrid Memory Space

In order to present the compiler/programmer with an abstracted view of the hybrid memory hierarchy and minimize the complexity of our run-time system we propose the use of virtual SPMs and virtual NVMs. We leverage the concept of vSPMs [5], which enables a program to view and manage a set of vSPMs as if they were physical SPMs. In order to virtualize SPMs, a small part of main memory (DRAM) called protected evict memory (*PEM*) space is locked and used as extra storage. The run-time system would then prioritize the data mapping to SPM and PEM space based on a utilization metric. In this work we introduce the concept of virtual NVMs (vNVMs), which behave similarly to vSPMs, meaning that the run-time environment transparently allows each application to manage their own set of vNVMs. Management of virtual memories is done through a small set of APIs [4], which send management commands to the *HaVOC* manager. The *HaVOC* manager then presents each application with intermediate physical addresses (*IPAs*), which point to their virtual SPMs/NVMs. Traditional SPM-based memory management requires the data layout to use physical addresses by pointing to the base register of the SPMs, as a result, the same is expected of SPM/NVM-based memory hierarchies [11]. In our scheme, all policies use virtual SPM and NVM base addresses, so any run-time re-mapping of data will remain transparent to the initial allocation policies as the *IPAs* will not change.

3.3 Hybrid Memory-aware Policy Generation

The run-time system needs compile-time support in order to make efficient allocation decisions at run-time. In this paper we present various ways by which designers may gen-

erate policies (manual through annotations or through static analysis). These policies are then enforced (currently in best effort fashion) by the run-time system in order to prioritize the access to SPM/NVM space for the various applications running on the system. Each policy attempts to map data to virtual SPMs/NVMs, while the *HaVOC* manager dynamically maps the data to physical memories.

3.3.1 Volatility Analysis

$$Data_{lifetime} \leftarrow \bigcup_{i=0}^n ST_i \quad (1)$$

$$Write_{freq}^i \leftarrow Writes_i \div ST_i, i = 0 \dots n \quad (2)$$

$$Data_{volatility} \leftarrow STDEV(Write_{freq}^i), i = 0 \dots n \quad (3)$$

$$C(D_i) = C_{load}(D_i^f) + C_{evict}(D_i^t) + C_{util}(D_i^r, D_i^w) + \Delta_{leak}(D_i^r, D_i^w, D_i^{lifetime}) \quad (4)$$

We introduce a new metric, *data volatility*, to facilitate efficient loading of data on the hybrid on-chip memory configurations. Data volatility is defined as the write frequency of a piece of data over its accumulated lifetime. In order to estimate the volatility of a data block we first define a sampling time (ST_i), which can be in cycles, so that the union of all sample times equals the block's lifetime (Eq. 1). Next, we calculate the write frequency for each sample time (Eq. 2). Finally, we estimate the volatility of the data as the variation in its write frequency (Eq. 3). This metric is useful when deciding whether data is worth (cost effective) being mapped onto NVM. Highly volatile data implies that at some point the cost of keeping data in NVM during its entire lifetime might be greater than leaving it in main memory. As a result, when two competing applications request NVM space, the estimated cost function (e.g., energy savings) will be used to prioritize allocation of on-chip space, while volatility can be used as a tie breaker and prediction metric of cost fluctuation. Volatility may also be used to decide the granularity at which designers might do their data partitioning. We define the expected cost metric ($C(D_i)$) for a given data block (D_i) as shown in Eq. 4, which takes into account the cost of transferring data between memory type D_i^f and memory type D_i^t (C_{load} , C_{evict}), the utilization cost (C_{util}), and the extra leakage power consumed by mapping the given data to the preferred memory type. The cost represents *energy or latency*.




Figure 3: Data volatility across various lifetimes.

Figure 3 (a) shows sample JPEG [21] code and how partitioning its data's lifetime may affect its data's volatility. Figure 3 (b) shows the global life time of the data arrays (a , b , c , zt , qt), where the number of accesses to NVM would be (128 rd, 23K wr) for qt/zt if we map and keep them in NVM during the entire execution of the program. To accommodate other data structures onto SPM space, arrays qt/zt 's lifetime may be split, resulting in finer life-time granularities (Figure 3 (c-d)). Though the rd/wr ratio of data remains the same (qt/zt have 23K reads to 0 writes), finer granularity

lifetimes might yield higher volatility (qt/zt now have 23K writes to NVM since they are loaded every time *block_decode* executes), making qt/zt poor candidates for NVM.

3.3.2 Memory Allocation Policy Generation

Programmers can embed application-specific insights into source code through annotations [10] in order to guide the compilation process. Since we are working with virtualized address spaces, programmers can create hybrid memory-aware policies that define the placement for a given data structure by simply defining the following parameters: $\langle \text{preferred memory type, reads, writes, lifespan, volatility} \rangle$. These annotations are used at run-time by the *HaVOC* manager to allocate the data onto the *preferred* memory type.

Instruction blocks are very good candidates for mapping onto on-chip NVMs since their volatility is quite low (e.g., write once and use many times). In this work, we borrow traditional SPM-based instruction-placement schemes [16] and enhance them to account for the possibility of mapping the instructions to NVM memories by introducing volatility analysis into the flow. Like we discussed in Sec. 3.3.1, the granularity of the code partitioning (e.g, function, basic block, etc.) will affect how volatile the placement will become. As a result, when mapping a block of instructions onto vNVM/vSPM, we need to partition our code such that Eq. 5 is met, where $C(D_i)$ represents the cost in Eq. 4. Our goal is to partition the application such that we can minimize the number of instruction replacements ([16]) in order to minimize energy and execution time.

$$C(D_i)_{Off-Chip} \ll C(D_i)_{On-Chip} \quad (5)$$

Data placement candidates are obtained from static analysis (e.g., data reuse analysis [13]) or profiling. The idea is to map highly read-reused data with a long access distance onto vNVM to minimize number of fetches from off-chip memory, while highly reused read-data with short lifetimes will be mapped to vSPM preferably. Highly reused-read-modify data with low write-volatility should be mapped to vNVM, while highly reused write-data and read-modify data with high write-volatility should go to vSPM.

The last step in our flow involves the generation of near-optimal hybrid memory layouts we define as *allocation policies*. This process takes as input data/instruction blocks with various pre-computed costs (e.g., Eq. 4) obtained from static analysis, profiling, and annotations, which are combined and fed as inputs to an enhanced hybrid memory-aware allocator based on [11].

3.4 HaVoC Manager




Figure 4: *HaVOC* Manager

The *HaVOC* manager may be implemented in software as an extended memory manager embedded within a hypervisor/OS or as a hardware module (e.g., [9, 5]). The software implementation is quite flexible and does not require modifying existing platforms. The hardware version requires additional hardware and the necessary run-time support, but

the run-time overheads will be much lower than the software version. In this work, we present a proof-of-concept embedded hardware implementation (Figure 2). Figure 4 shows a block diagram of the *HaVOC* manager. It consists of a memory-mapped slave interface, which is used by the system’s masters (e.g., CPUs) and handles the read/write/configuration requests. The address translation layer module converts IPAs to physical addresses (PAs) in one cycle [7, 5]). The manager consists of 1KB to 256KB of configuration memory used to keep block metadata information (e.g., volatility, # accesses, etc.). The allocation/eviction logic uses the cost estimation (e.g., efficiency) logic to prioritize access to on-chip storage. Finally, the internal DMA (iSDMA) allows the manager to asynchronously transfer data between on-chip and off-chip memory. In order to use the *HaVOC* manager, the compiler generates two things: 1) the creation of the required virtual SPM/NVMs through the use of our APIs and 2) The use of IPAs instead of PAs during the data/instruction layout stage (e.g., memory maps using purely virtual addresses for SPMs/NVMs). Any read/write to a given IPA is translated and routed to the right memory. Any write to *HaVOC* configuration memory space is then used to manage the virtualized address space. The goal is to allow each application to manage its virtual on-chip memories as if it had full access to the on-chip real-estate.

3.5 HaVOC’s Dynamic Policy Enforcement

Table 1: Filter Inequalities and Preferred Memory Type

Filter	Pref.	Inequalities
F1	sram	$E(D_i^{spm}) > E(D_i^{dram}) \wedge E(D_i^{nvm}) < E(D_i^{dram}) \wedge V > T_{vol}$
F2	nvm	$E(D_i^{nvm}) > E(D_i^{dram}) \wedge V < T_{vol}$
F3	either	$E(D_i^{spm}) > E(D_i^{dram}) \wedge E(D_i^{nvm}) > E(D_i^{dram})$
F4	dram	$E(D_i^{spm}) < E(D_i^{dram}) \wedge E(D_i^{nvm}) < E(D_i^{dram})$

Algorithm 1 FilterDynamic Allocation Algorithm

```

Require: req{size, cost, volatility}
pref_mem ← filter(req, volatility)
2: if allocatable(req, pref_mem) then
    return ipa ← update_alloc_table(req)
4: end if
min_set ← sortH2LowEff(alloc_table, size)
6: if  $E(min\_set) < E_{evict}(min\_set) + E(req)$  then
    evict(min_set)
8: return ipa ← update_alloc_table(req)
else
10: return ipa ← mm_malloc(req)
end if

```




Figure 5: Dynamic Hybrid Memory Allocation Policies

We define *three* block-based allocation policies: 1) *Temporal* allocation, which combines temporal SPM allocation ([28]) and hybrid memory allocation ([11]), and adheres to the initial layout obtained through static analysis (Sec. 3.3); however, the application’s SPM and NVM contents must be swapped on a context-switch to avoid conflicts with other tasks (Fig. 5 (a)). 2) *FixedDynamic* allocation, which combines dynamic-spatial SPM allocation ([9]) and hybrid memory allocation [11], and maps the data block to the preferred memory type (adhering to the initial layout) as long as there is space, otherwise, data is mapped to DRAM (Fig. 5 (b)). 3) *FilterDynamic* allocation (Alg. 1), which exploits the

concept of filtering and volatility to find the best placement. Each request is filtered according to a set of inequalities (shown in Table 1) which determine the preferred memory type (Fig. 5 (d)). The volatility of the data block (V) and its mapping efficiency ($E(D_i) = C(D_i)/|D_i|$) are used to determine what memory type would minimize the block’s energy (or access latency). For instance, data with low volatility and high energy efficiency could potentially benefit more from being mapped to NVM than SRAM (e.g., filter $F2$ in Table 1). If there is enough preferred memory space (e.g., SPM or NVM), the dynamic allocator adheres to Eq. 4 prior to loading the data. If there is not enough space, then the allocator follows Alg. 1 and sorts the allocated blocks from highest to lowest efficiency (e.g., energy per bit). It then compares the cost of evicting the least important blocks (MIN_{Set}) with the cost of dedicating the space to the new block. If the efficiency of bringing the new block offsets the eviction cost and efficiency of the data already mapped to the preferred memory type ($E(MIN_{Set})$), then *HaVOC* evicts the MIN_{Set} blocks and updates the allocation table with the new block ($|new\ block| \leq |MIN_{Set}|$). In the event the preferred memory type is either NVM or SPM (filter $F3$ in Table 1), the allocator evicts the $\min(MIN_{Set}^{spm}, MIN_{Set}^{nvm})$. At the end, *HaVOC* allocates either on-chip space or off-chip space (unified PEM space ($sPEM$)), resulting in the allocation shown in Fig. 5 (c), where a data block originally intended for SPM is mapped to NVM.

4. RELATED WORK

Most efforts have focused on replacing and/or complementing main memory (DRAM) or caches (SRAM) with a combination of various NVMs to reduce leakage power and increase throughput. Joo et al. [15] proposed PCM as an alternative to SRAM for on-chip caches. Sun et al. [27] introduced MRAM into the cache hierarchy of a NUCA-based 3D stacked multi-core platform. Mishra et al. [24] followed up by introducing STT-RAM as an alternative MRAM memory and hid the overheads in access latencies by customizing how accesses are prioritized by the the interconnect network. Wu et al. [31] presented a hybrid cache architecture consisting of SRAM (fast L1/L2 accesses), eDRAM/MRAM (slow L2 accesses), and PCRAM (L3). Hybrid main memory has also been studied [32]. Mogul et al. [25] and Wongchaowart et al. [30] attempted to reduce write overheads in main memory by exploiting page migration and block content signatures. Ferreira et al. [8] introduced a memory management module for hybrid DRAM/PCM main memories. Static analysis has been explored to efficiently map application data on off-chip hybrid main memories [23, 22].

HaVOC is different from approaches that address hybrid cache/main memories in that we primarily focus on hybrid SPM/NVM-based hierarchies, however, our scheme can be complemented by both existing schemes that leverage the benefits of both on-chip SRAMs and NVMs as well as hardware/system-level solutions that address hybrid off-chip memories (e.g., DRAM/eDRAM and PCM). Our work is different from [11] in that we consider: 1) shared distributed on-chip SPMs and 2) dynamic support for multitasking systems, however, our scheme can benefit from compile-time analysis schemes that provide our runtime system with allocation hints (e.g., [22, 23, 11]). Like [5], we use part of off-chip memory to virtualize on-chip memories, however, our approach differs in that our primary focus is the ef-

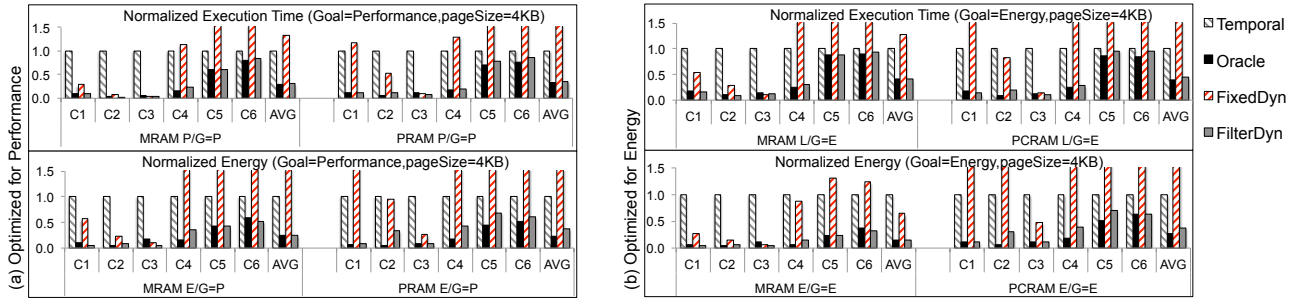


Figure 6: Normalized Execution Time and Energy for Performance Optimized (a) and Energy Optimized (b) Policies

efficient management of hybrid on-chip memories, and as a result, *HaVOC*'s programming model and run-time environment account for the different physical characteristics of SRAMs, DRAMs and NVMs (MRAM and PCRAM). Moreover, we believe that we can complement our static-analysis/allocation policy generation with other SPM management techniques [17, 29, 9, 26, 28].

5. EXPERIMENTAL RESULTS

Table 2: Configurations

Config.	Applications	CPUs	vSPM/vNVM Space	SPM/NVM Space
C1	adpcm,aes	1	32/128 KB	16/64 KB
C2	adpcm,aes,blowfish,gsm	1	64/256 KB	16/64 KB
C3	C2 & h263,jpeg,motion,sha	1	128/512 KB	16/64 KB
C4	same as C2	2	64/256 KB	32/128 KB
C5	same as C3	2	128/512 KB	32/128 KB
C6	same as C3	4	128/512 KB	64/256 KB

5.1 Experimental Setup and Goals

Our goal is to show that *HaVOC* is able maximize energy savings and increase application throughput in a multitasking environment under various scenarios. First, we generate two sets of hybrid memory aware allocation policies (Sec. 3.3.2), one set of policies minimizes execution time (Sec. 5.2) and the other minimizes energy (Sec. 5.3). These policies are generated at compile-time and enforced at run-time by a set of dynamic allocation policies (Sec. 3.5) under various system configurations (Table 2). Next, we show the effects of the allocation policy's block-size on execution time (Sec. 5.4). We built a trace-driven simulator that models a light-weight OS, with a round-robin scheduler and context-switching enabled (window = 50K cycles). We model CMPs consisting of an AMBA AHB bus, OpenRISC-like in-order cores, distributed SPMs and NVMs, and the *HaVOC* manager (Fig. 2). We bypassed the cache and mapped all data to either SPM, NVM, or main memory (see Sec. 3.3.2). We obtained traces from Mediabench II [21] by using SimpleScalar [1]. We model on-chip SPMs (SRAMs), MRAMs and PCRAM by interfacing our simulator with NVSim [6] and set leakage power as the optimization goal. To virtualize SPMs/NVMs we use the unified PEM space model discussed in Sec. 3.2 (sPEM). The *HaVOC* manager consists of 4KB low power SRAM memory.

5.2 Enforcing Performance Optimized Policies

For this experiment we generated allocation policies that minimized execution time for each application. We then executed each application on top of our simulated RTOS/CMP. Table 2 shows each configuration (C1-6), which has a set of applications running concurrently over a number of CPUs, and a predefined hybrid memory physical space. To show the benefit of our approach we implemented four policies:

the three described in Sec. 3.5 (*Temporal*, *FixedDynamic*, *FilterDynamic*), and a policy we call *Oracle* (black bar in Fig. 6), which is a near-optimal policy because on every block-allocation request, it feeds the entire memory map to the same policy generator the compiler uses to generate policies statically (see Sec. 3.3.2). The idea is to show that our *FilterDynamic* policy (backward-slashed bars in Fig. 6) achieves competitive quality allocation solutions as the more complex *Oracle* policy. Fig. 6 (a) shows the normalized execution time and energy for each of the different configurations (C1-6, Goal=Min Execution Time denoted as $G=P$) using 4KB blocks and different memory types with respect to the *Temporal* policy. The *FixedDynamic* policy (forward-slashed bars in Fig. 6 (a)) suffers the greatest impact on energy and execution time as memory space increases (C4-6) since it adheres to the decisions made at compile-time and does not efficiently allocate memory blocks at run-time. In general, we see that the *FilterDynamic* policy performs almost as good as the *Oracle* policy (within 8.45% execution time). Compared with the *Temporal* policy, *HaVOC*'s *FilterDynamic* policy is able to reduce execution time and energy by an average 75.42% and 62.88% respectively when the initial application policies have been optimized for execution time minimization.

5.3 Enforcing Energy Optimized Policies

Fig. 6 (b) shows the the normalized execution time and energy for each of the different configurations and memory types (Goal=Min Energy denoted as $G=E$) with respect to the *Temporal* policy. Similar to the case of $G=P$, both the *Temporal* and *FixedDynamic* policies are unable to efficiently manage the on-chip real-estate. The *FilterDynamic* and *Oracle* policies are able to greatly reduce execution time and energy, with the *FilterDynamic* within 3.54% of the execution time achieved by the *Oracle* policy. Compared with the *Temporal* policy, *HaVOC*'s *FilterDynamic* policy is able to reduce execution time and energy by an average 85.58% and 61.94% respectively when the initial application policies have been optimized for energy minimization. The goal of this experiment was to show that regardless of the initial optimization goal, *HaVOC*'s *FilterDynamic* policy is able to achieve as good results as the *Oracle* policy.

5.4 Block Size Effect on Allocation Policies

So far we have seen that the *Oracle* policy appears to be a feasible dynamic allocation solution, which would potentially enhance *HaVOC*'s virtualization engine. However, the *Oracle* policy is very complex as it runs in $O(Blks * spm_{size} * nvms_{size})$. The *FixedDynamic* policy on the other extreme runs in $O(Blks)$, however, its efficiency may be even worse than the *Temporal* policy. *HaVOC*'s *FilterDy-*




Figure 7: Effects of Varying Block Size.

dynamic policy on the other hand, keeps a semi-sorted list of data blocks, as a result it can be $O(Blks)$ best case or $O(Blks \log Blks)$ worst case for sorting, and the final filtering results runs in $O(Blks_{spm} + Blks_{nvm} + Blks_{dram})$, which results in $O(Blks \log Blks) + O(Blks_{spm} + Blks_{nvm} + Blks_{dram})$. Thus, the complexity and execution time of the *Oracle* policy will increase orders of magnitude as the number of data/instruction blocks to allocate increases ($Blks$) or as the available resources increases (spm_{size}, nvm_{size}). This is validated in Fig. 7, where we have varying block size (as a result, number of blocks to allocate increases) and increase in available resources (*C3-C6*). As we can see, for block size = 1KB, where the number of blocks to allocate is in the hundreds, the allocation time of the *Oracle* is orders of magnitude greater than the *FilterDynamic* policy. For block size = 2KB, we see that as resources increase, the *Oracle* allocation time once again prevents it from being a feasible solution. On average, we observe that *HaVOC*'s *FilterDynamic* is capable of achieving as good solutions as the *Oracle* policy (within 10% margin) with much lower complexity. On average, across all test scenarios (varying page sizes, different optimization goals, and different configurations) we see that the *FilterDynamic* is able to reduce execution time and energy by 60.8% and 74.7% respectively.

6. CONCLUSION

We presented *HaVOC*, a hybrid memory aware virtualization layer for dynamic memory management of applications executing on CMP platforms with hybrid on-chip NVM/SPMs. We introduced the notion of data volatility analysis and proposed a dynamic filter-based memory allocation scheme to efficiently manage the hybrid on-chip memory space. Our experimental results for embedded benchmarks executing on a hybrid memory-enhanced CMP show that our dynamic filter-based allocator greatly minimizes execution time and energy by 60.8% and 74.7% respectively with respect to traditional multitasking-aware SPM allocation policies. Future work includes: 1) Integrating *HaVOC*'s concepts in a hypervisor to support full on-chip hybrid memory virtualization, 2) Adding off-chip NVM memories to support the virtualization of on-chip NVMs (*nPEM*), and 3) Designing a scalable hybrid-memory aware virtualization layer.

7. ACKNOWLEDGMENTS

This work was partially supported by NSF Variability Expedition Grant Number CCF-1029783.

8. REFERENCES

- [1] T. Austin et al. SimpleScalar: an infrastructure for computer system modeling. *Computer*, 35(2), Feb. 2002.
- [2] N. Azizi et al. Low-leakage asymmetric-cell sram. *TVLSI*, Vol. 11, aug. 2003.
- [3] R. Banakar et al. Scratchpad memory: design alternative for cache on-chip memory in embedded systems. In *CODES '02*, 2002.
- [4] L. Bathen and N. Dutt. A hybrid-memory-aware virtualization layer for cmps. In *UCI TR #12-03*, 2011.
- [5] L. Bathen et al. SPMVisor: dynamic scratchpad memory virtualization for secure, low power, and high performance distributed on-chip memories. In *CODES+ISSS '11*, 2011.
- [6] X. Dong et al. Pcrsim: system-level performance, energy, and area modeling for phase-change ram. In *ICCAD '09*, 2009.
- [7] B. Egger et al. Dynamic scratchpad memory management for code in portable systems with an mmu. *ACM TECS*, 7, January 2008.
- [8] A. Ferreira et al. Using pcm in next-generation embedded space applications. In *RTAS '10*, april 2010.
- [9] P. Francesco et al. An integrated hardware/software approach for run-time scratchpad management. In *DAC '04*, 2004.
- [10] S. Guyer et al. An annotation language for optimizing software libraries. In *DSL '99*, 1999.
- [11] J. Hu et al. Towards energy efficient hybrid on-chip scratch pad memory with non-volatile memory. In *DATE '11*, 2011.
- [12] IBM. Cell. <http://www.research.ibm.com/cell/>, 2005.
- [13] I. Issenin et al. Multiprocessor system-on-chip data reuse analysis for exploring customized memory hierarchies. In *DAC '06*, 2006.
- [14] ITRS. Process integration, device and structures. <http://www.itrs.net/>, 2007.
- [15] Y. Joo et al. Energy- and endurance-aware design of phase change memory caches. In *DATE '10*, march 2010.
- [16] S. Jung et al. Dynamic code mapping for limited local memory systems. In *ASAP '10*, july 2010.
- [17] M. Kandemir et al. Dynamic management of scratch-pad memory space. In *DAC '01.*, 2001.
- [18] S. Kaneko et al. A 600mhz single-chip multiprocessor with 4.8gb/s internal shared pipelined bus and 512kb internal memory. In *SSCC '03*, 2003.
- [19] N. Kim et al. Leakage current: Moore's law meets static power. *Computer*, 36(12), dec. 2003.
- [20] B. Lee et al. Architecting phase change memory as a scalable dram alternative. *SIGARCH Comput. Archit. News*, 37, June 2009.
- [21] C. Lee et al. Mediabench: a tool for evaluating and synthesizing multimedia and communications systems. In *MICRO '97*.
- [22] K. Lee et al. Application specific non-volatile primary memory for embedded systems. In *CODES+ISSS '08*, 2008.
- [23] T. Liu et al. Power-aware variable partitioning for dsps with hybrid pram and dram main memory. In *DAC '11*, 2011.
- [24] A. Mishra et al. Architecting on-chip interconnects for stacked 3d stt-ram caches in cmps. In *ISCA '11*, 2011.
- [25] J. Mogul et al. Operating system support for nvm+dram hybrid main memory. In *HotOS'09*, 2009.
- [26] V. Suhendra et al. Scratchpad allocation for concurrent embedded software. In *CODES+ISSS '08*, 2008.
- [27] G. Sun et al. A novel architecture of the 3D stacked MRAM L2 cache for CMPs. In *HPCA '09*, feb. 2009.
- [28] H. Takase et al. Partitioning and allocation of scratch-pad memory for priority-based preemptive multi-task systems. In *DATE '10*, 2010.
- [29] M. Verma et al. Data partitioning for maximal scratchpad usage. In *ASP-DAC '03*, 2003.
- [30] B. Wongchaowart et al. A content-aware block placement alg. for reducing pram storage bit writes. In *MSST '10*.
- [31] X. Wu et al. Hybrid cache architecture with disparate memory technologies. In *ISCA '09*, 2009.
- [32] P. Zhou et al. A durable and energy efficient main memory using phase change memory technology. In *ISCA '09*, 2009.