

# Partitioning Techniques for Partially Protected Caches in Resource-Constrained Embedded Systems<sup>1</sup>

KYOUNGWOON LEE

University of California, Irvine

AVIRAL SHRIVASTAVA

Arizona State University

NIKIL DUTT

University of California, Irvine

and

NALINI VENKATASUBRAMANIAN

University of California, Irvine

---

---

<sup>1</sup>This is an expanded version of a conference paper published in the Proceedings of the IFIP Working Conference on Distributed and Parallel Embedded Systems (DIPES) 2008. The current manuscript extends the previous paper by: (i) exploring the design space for the instruction PPC (partially protected caches) in Section 6.3, (ii) proposing two more heuristic algorithms to efficiently find the best partitions with minimal vulnerability under the performance penalty in Section 4.3, (iii) presenting extensive experiments in Section 6 with the comprehensive energy consumption model in Section 5, and (iv) differentiating our work with comprehensive related work in Section 2 and preliminary experiments in Section 4.1.

---

Author's address: K. Lee, N. Dutt, and N. Venkatasubramanian, Department of Computer Science, University of California, Irvine, CA 92697.

A. Shrivastava, Department of Computer Science and Engineering, Arizona State University, Tempe, AZ 85281.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1084-4309/20YY/0400-0001 \$5.00

ACM Transactions on Design Automation of Electronic Systems, Vol. V, No. N, Month 20YY, Pages 1-0??.

Increasing exponentially with technology scaling, the soft error rate in even earth-bound embedded systems manufactured in deep sub-nanometer technology is projected become a serious design consideration. Partially Protected Cache (PPC) is a promising microarchitectural feature to mitigate failures due to soft errors in power, performance and cost sensitive embedded processors. A processor with PPC maintains two caches, one protected and the other unprotected, both at the same level of memory hierarchy. The intuition behind PPCs is that not all data in the application is equally prone to soft errors. By finding and mapping the data that is more prone to soft errors to the protected cache, and error-resilient data to the unprotected cache, failures induced by soft errors can be significantly reduced at a minimal power and performance penalty. Consequently, the effectiveness of PPCs is critically hinged on the compiler's ability to partition application data into error-prone and error-resilient data. The effectiveness of PPCs has previously been demonstrated on multimedia applications - where an obvious partitioning of data exists: the multimedia data is inherently resilient to soft errors, and the rest of the data and entire code is assumed to be error-prone. Since the amount of multimedia data is quite significant component of the entire application data, this obvious partitioning is quite effective. However, no such obvious data and code partitioning exists for general applications. This severely restricts the applicability of PPCs to data caches and instruction caches in general. This article investigates vulnerability-based partitioning schemes that are applicable to applications in general and effectively reduce failures due to soft errors at minimal power and performance overheads.

Our experimental results on HP iPAQ-like processor enhanced with PPC architecture, running benchmarks from MiBench suite demonstrate that our partitioning heuristics efficiently find page partitions for data PPCs that can reduce the failure rate by 48% at only 2% performance and 7% energy overhead, and find page partitions for instruction PPCs that reduce the failure rate by 50% at only 2% performance and 8% energy overhead, on average.

Categories and Subject Descriptors: B.3.4 [Memory Structures]: Reliability, Testing, and Fault-Tolerance

General Terms: Design, Experimentation, Performance, Reliability

Additional Key Words and Phrases: Page Partitioning Technique, Partially Protected Cache, Soft Error, Vulnerability, Embedded Systems

---

## 1. INTRODUCTION

System reliability is becoming the paramount concern in system design in the deep submicron era [ITRS 2005]. Four decades of technology scaling has brought us to a point where the transistors have become extremely susceptible to even small fluctuations in voltage levels, slight noises in power supply, signal interference., and even cosmic particle strike [Hazucha and Svensson 2000; Wrobel et al. 2001; Shivakumar et al. 2002; Baumann 2005]. All of these effects can temporarily toggle the logic value of a transistor, and it is therefore called a transient fault. Such transient faults are random and non-destructive, i.e., resetting the device restores normal behavior. Previous investigation finds cosmic radiation strikes to be responsible for more transient faults than all the other reasons combined [Baumann 2005]. A high energy radiation particle, e.g., an alpha particle, a neutron, or a free proton, may strike the diffusion region of a CMOS transistor and produce charge which can result in toggling the logic value of the gates or flip-flops. This phenomenon of change in the logic state of a transistor is called an *Upset*. An upset may have catastrophic consequences including the application generating incorrect results, accessing unauthorized memory regions, crashing, or going into an infinite loop. The incorrect or

erroneous behavior of an application due to upsets is called a *Failure*.

Not all upsets result in failures, and upsets can be masked due to masking effects such as electrical masking (upset is not strong enough to reach the next latching element), logical masking (upset on the input of the gate does not affect its output), latching-window masking (upset does not reach latch at the latching time), microarchitectural masking (upset happens on a variable which is not used any more), and software masking (upset happens in a function, whose output is not used) [Shivakumar et al. 2002]. An upset will become a failure if it is not masked by any of these effects.

Owing to the effectiveness of the latching-window masking, upsets in memory elements have significantly higher probability of causing a failure than upsets in combinational logic [Gaisler 1997; Liden et al. 1994]. In addition, since memory elements may occupy more than majority of the chip area, and the fact that they operate on lower voltages than combinational circuits, they are extremely vulnerable to radiation, and radiation-induced faults. In fact, according to [Mitra et al. 2005], more than 50% of soft errors occur in memories.

While it is possible to employ simple Error Correction Code (ECC) based techniques in off-chip and lower levels of memories, such solutions are not suitable for caches, as they are highly sensitive to the performance and power overheads of redundancy-based techniques. For example, using Single-bit Error Correction and Double-bit Error Detection (SEC-DED) codes may increase the cache access time by 95% [Li and Huang 2005], power consumption by 22% [Phelan 2003], and area cost by 25% [Krueger et al. 2008]. While it may be possible to hide the performance penalty, it is not possible to hide the power penalty. Consequently, novel techniques are required for caches that can eventually reduce failure rates while incurring minimal power and performance overheads.

Partially Protected Cache (PPC) architectures were proposed to mitigate failures due to soft errors on caches at minimal power, performance and area overheads [Lee et al. 2006]. A PPC architecture has two caches, one protected against soft errors, and the other unprotected, at the same level of memory hierarchy. The intuition behind PPC is that "not all data is equally prone to soft errors," and that most of the soft errors show up in a relatively small amount of application data (and code). Thus by protecting a small amount of data by a small protected cache, programs can be made robust at minimal power, performance, and area penalty.

Clearly the effectiveness of PPC architectures is critically hinged on the compiler's ability to partition the application data and code into error-prone and error-resilient data and code. At the microarchitecture level, we use *Vulnerability* as the measure of how error-prone a variable is [Mukherjee et al. 2003; Asadi et al. 2005]. Vulnerability of a program variable is essentially the probability that the occurrence of a soft error in the variable in the cache will affect the program state, possibly causing a program failure. Estimating the vulnerability of a program variable is an inherently difficult task, as several program and microarchitectural factors may have a significant effect on the vulnerability of the program variable. These include i) the access pattern of the variable, e.g., a variable that is not read by the processor, and will be overwritten is not vulnerable, ii) the access pattern of the other program variables, e.g., if the other data evicts this variable from the cache, then it will be in memory,

and will not be vulnerable, and iii) cache characteristics, e.g., the size, associativity, write and allocate policies of the cache can significantly affect the time a datum is vulnerable in the cache.

Previously, PPCs have been shown to be extremely cost effective in soft error protection for multimedia applications. An obvious partitioning of data into soft error-prone and soft error-resilient data exists in multimedia application: wherein, the multimedia data itself is quite soft error-resilient. For example, in an image or video processing application, a soft error in the image or video data itself only causes a slight degradation in the Quality of Service (QoS). In contrast, most other data, e.g., loop control variables, stack pointers, etc., are not error-resilient. This obvious partitioning is very effective since the size of the multimedia data is quite significant.

However, no obvious data partitioning exists for general applications. To use PPC architectures for applications in general, a data partitioning scheme that divides application data (and code) into soft error prone and soft error resilient is extremely critical. In this article, we examine application profile to partition the application data and code into error-prone and error-resilient, enabling the use of PPC architectures for several application specific embedded systems.

We find that Monte Carlo exploration is unable to find interesting data partitions. While Genetic Algorithm can efficiently search the exploration space, it does not achieve high reduction in vulnerability. Our partitioning heuristics are aware of runtime and vulnerability, and are therefore able to efficiently prune the search space and uncover Pareto-optimal partitions. We propose data partitioning schemes for both data PPC and instruction PPCs. Experimental results on the HP iPAQ h4600- like processor memory subsystem [Hewlett Packard ] running benchmarks from the MiBench suite [Guthaus et al. 2001] demonstrate that data PPC architectures can reduce the vulnerability by 48% with 2% performance and 7% energy penalty on average. In addition, instruction PPC architectures can reduce the program vulnerability by 50% while incurring 2% performance and 8% energy consumption overheads, on average.

## 2. RELATED WORK

Radiation-induced soft errors have been under investigation since late 1970s. Due to incessant technology scaling, soft error rate (SER) has exponentially increased [Hazucha and Svensson 2000], and now it has reached a point, where it becomes a real threat to system reliability.

### 2.1 Packaging Solutions

Radioactive substances such as alpha particles emitted by packaging and wafer processing materials are one of the major sources of radiations that cause soft errors in semiconductors. Thus, advances in process technology such as purification of packaging materials, radiation hardening, and elimination of Boron-10 ( $B^{10}$ ) impurities, are expected to mitigate the soft errors [Baze et al. 2000]. However, the effects of interactions between high energetic cosmic particles (e.g., neutrons) and radioactive materials cannot be prevented completely [Mastipuram and Wee 2004].

## 2.2 Process Technology Solutions

Process technology solutions such as SOI (Silicon On-Insulator) processes [Musseau 1996; Roche et al. 2003] have been proposed. In order to mitigate the soft errors, they extend the depletion region or raise the capacitance, which increases the critical charge of semiconducting devices. The critical charge is the least charge to be able to invert the bit value of the memory cell. However, process engineering technology may require the cost of additional process complexity, the loss of manufacturability, and extra substrate cost [Baumann 2005].

## 2.3 Microarchitectural Solutions

Microarchitectural solutions attempt to reduce the number of upsets that translate into errors and/or errors that result in failures. Solutions at the microarchitecture level can be categorized based on the components where they are applied: the combinational components, the sequential components, and the memory components.

### **Solutions for Combinational Logic**

Logic elements were considered more robust against soft errors than memory elements mainly due to the masking effects. However, many researchers predict that the logic soft errors will become one of main contributions to the system unreliability [Shivakumar et al. 2002; Baumann 2005; Nieuwland et al. 2006]. The simplest and most effective way to reduce failures due to soft errors in combinational logic is Triple Modular Redundancy (TMR) [Pradhan 1996], which typically uses three functionally equivalent replicas of a logic circuit and a majority voter. But the overheads of hardware and power for conventional TMR exceed 200% [Nieuwland et al. 2006]. Duplex redundancy [Mohanram and Touba 2003; Nieuwland et al. 2006] is also available but it requires more than 100% area and power overheads without any optimization techniques. In order to reduce the high overheads in conventional redundancy techniques, Mohanram et al. in [Mohanram and Touba 2003] presented a partial error masking by duplicating the most sensitive and critical nodes in a logic circuit based on the asymmetric susceptibility of nodes to soft errors. Nieuwland et al. in [Nieuwland et al. 2006] proposed a structural approach analyzing the soft error rate sensitivity of combinational logic to identify the critical components at circuits.

### **Solutions for Sequential Logic**

Temporal redundancy is another main approach that has been used to combat soft errors in circuits. In order to detect soft errors, Nicolaidis in [Nicolaidis 1999] applied fine time-grain redundancy within the clock cycle greater than the duration of transient faults by using the temporal nature of soft errors. Similarly, Anghel et al. in [Anghel and Nicolaidis 2000] exploited the temporal nature to detect timing errors and soft errors by means of time redundancy. Krishnamohan et al. in [Krishnamohan and Mahapatra 2004] proposed the time redundancy methodology by using the timing slack available in the propagation path from the input to the output in CMOS circuits. A Razor flip-flop was presented in [Ernst et al. 2003] to detect transient errors by sampling pipeline stage values with a fast clock and with a time-borrowing delayed clock.

### **Solutions for Memories**

By far, reducing soft errors in memories has been the most extensive research

topic. Error detection and correction codes (EDC and ECC) have been widely investigated and implemented as the most effective scheme to detect and correct soft errors in memory systems. However, an ECC system consists of an encoding block as well as a decoding block responsible for detection and correction, and of extra bits storing parity values. Thus, ECC-based techniques consume extra energy and incur performance delay as well as additional area cost [Pradhan 1996; Phelan 2003; Li and Huang 2005; Krueger et al. 2008], and are therefore not suitable for caches. Thus, only a few processors such as the Intel Itanium processor [Quach 2000] protect L2 and L3 caches with ECC [Stackhouse et al. 2008], but we are not aware of any processor employing ECC-based protection mechanism on L1 cache. This is mainly due to high overheads of ECC implementation [Kim 2006; Mohr and Clark 2006; Zorian et al. 2005]. [Zhang et al. 2003] proposed in-cache replication where the dead cache block space is recycled to hold replicas of the active cache block. Also, Zhang [Zhang 2005b] presented replication cache where a small fully associative cache is added to keep the replica of every write to the L1 data cache. However, these techniques incur overheads to maintain replicas. A cache scrubbing technique [Mukherjee et al. 2004] has been proposed, which can fix all single-bit errors periodically and prevent potential double-bit errors. Li et al. in [Li et al. 2004] evaluated the drowsy cache and the decay cache exploiting voltage scaling and shut-down schemes, respectively, in order to efficiently decrease the power leakage. They also proposed an adaptive error correcting scheme to different cache data blocks, which can save energy consumption by protecting clean data less than dirty data blocks. Kim in [Kim 2006] proposed the combined approach of parity and ECC codes to generate the reliable cache system in an area-efficient way. However, they all exploit expensive error correcting codes in order to protect all the data unnecessarily.

#### **Partially Protected Cache Architecture**

Lee et al. in [Lee et al. 2006] proposed PPC architecture and demonstrated the effectiveness in reducing the failure rate with minimal power and performance overheads. However, the effectiveness of PPCs has been limited only on multimedia applications, and there is no known approach to use PPCs for both data and instruction caches in general applications.

## **2.4 Software Solutions**

Software-only techniques have been studied to protect data and code from soft errors. Both software and hardware techniques have their own advantages and disadvantages in combating the impacts of soft errors. For example, hardware techniques increase the resource cost with high effectiveness to detect and even correct errors while software solutions mostly do not incur hardware costs with minimal coverage such as only error detection.

Reis et al. [Reis et al. 2005] presented the software-implemented fault tolerance (SWIFT) for soft error detection by exploiting unused resources and enhancing control-flow checking. Also, Luccetti et al. [Luccetti et al. 2005] proposed software mechanisms to tolerate soft errors by leveraging virtual machine and memory sharing techniques. However, they are limited only to detecting errors, and must be used in conjunction with recovery techniques. Through the user-specified annotations, the compiler can separate and map data elements in programs to reliable

domain which has protection techniques against soft errors, and to unreliable domain without protection [Chen et al. 2005]. But it requires the annotation for important data by user specification.

Soft error detection in software is extremely expensive in terms of delay, while it can be done without much overhead in hardware. In contrast, since the soft error rate is very low (as compared to processor clock cycle), soft error correction is efficient in software while it incurs too much overhead in hardware. Consequently, a combined approach that achieves the best of both hardware and software solutions will be very efficient. However, there is no hardware-software hybrid approach for soft error mitigation in resource-constrained embedded systems.

The PPC architecture with software page partitions is the promising one as a joint solution of hardware-software techniques. The compiler separates the failure critical and failure non-critical data and maps each of them into the two caches in a PPC for the selective data protection technique in multimedia applications [Lee et al. 2006]. However, there is no partitioning technique for general applications.

*Our Contribution – This article investigates the software challenges in using PPCs. The contribution of this article is in developing techniques to utilize PPC architectures for applications in general and establish PPC as an effective microarchitectural solution to mitigate failures due to soft errors not only for data caches but also for instruction caches.*

### 3. PARTIALLY PROTECTED CACHES AND PROBLEM DEFINITION




Fig. 1. Partially Protected Cache Architecture: one protected cache and the other unprotected cache at the same level of hierarchy

In a processor with *Partially Protected Cache* (PPC), the processor has two caches at the same level of memory hierarchy. As shown in Fig. 1, one of two caches is protected from soft errors, while the other one is unprotected. Any soft error protection mechanism can be implemented in the protected cache, e.g., increasing the thickness of oxide layer of the transistors in the cache, or adding redundancy logic like SECDED. To keep the access latencies of the protected cache and the unprotected cache the same, the protected cache is typically smaller than the unprotected cache.

Each page in the memory is mapped exclusively to one of the caches in a PPC architecture. The page mapping is set as a page attribute by the compiler. The mapping of the pages present in the cache resides in the Translation Lookaside Buffer (TLB). On a cache access, first a TLB lookup is performed to find out if the

page is present in the cache, and if so, in which one? Thus, only one cache lookup is performed per cache access.

While PPC architectures can be very effective in reducing the failure rate with minimal performance and power overheads, the effectiveness hinges on the ability to partition the application data and code between the two caches. To motivate for the need and effectiveness of page partitioning to reduce the failure rate, we performed a small experiment. First we map all the application pages to the unprotected cache, and then move the pages to the protected cache one by one. Fig. 2 plots the failure rate at each step of this exploration for benchmark *susan corners* from MiBench suite on a modified *sim-outorder* simulator from SimpleScalar [Burger and Austin 1997] to model HP-iPAQ like system. To estimate the failure rate, we injected soft errors on data caches for each execution of the benchmark, counted the number of failures out of a thousand executions. Each execution is defined as a success if it ends and returns the correct output. Otherwise, it is a failure. Fig. 2 shows that the failure rate of the application drops rapidly as pages are moved from the unprotected cache to the protected cache. Note that y-axis is logarithmic in Fig. 2. However, the pages have to be carefully moved to the small protected cache, as it is small; mapping too many pages to the small cache may increase the cache misses and result in a significant degradation of performance and increase in the energy consumption. Indeed, the performance can decrease by up to 27% for *susan corners* when all pages are mapped to the 256 byte protected cache as compared to all pages to the 4 KB unprotected cache in a PPC. So there is a definite need to study the tradeoff between the failure rate and performance (energy consumption) in finding the partitions for PPC architectures.

Therefore, the partitioning problem is a multi-objective optimization problem in which we need to reduce the failure rate, at minimal performance degradation, and minimal increase in the energy consumption. Since, even medium sized applications use a large number of data pages; our benchmarks selected from MiBench suite [Guthaus et al. 2001] access 27 - 95, on average 56 pages. Owing to their exponential complexity, enumerative techniques (e.g. trying all the possible page partitions and picking up the best one) do not work.




Fig. 2. Failure Rate Reduction by Moving Pages from the Unprotected Cache into the Protected Cache One by One in a PPC



We formulate our problem as: *Given an allowable performance degradation, determine the page partitioning to minimize the failure rate at minimal energy penalty.*

## 4. OUR APPROACH

### 4.1 Vulnerability: Microarchitectural Metric for Failure Rate

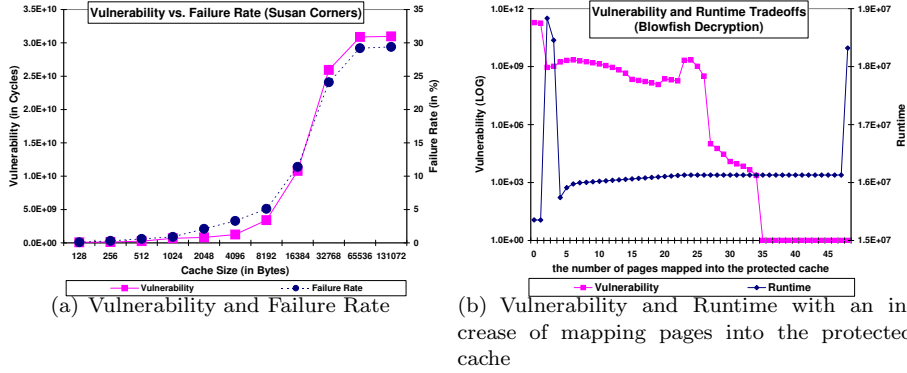


Fig. 3. Vulnerability is a good metric for estimating failure rate and partitioning pages to PPC architectures.

Several architectural techniques have been developed to improve the failure rates of applications. Developing any technique to reduce the failure rate requires estimating the failure rate. However, failure rate estimation is a very time intensive task. Recall that the occurrence of an error does not imply application failure. To estimate the failure rate, several simulations need to be performed, and the mean be reported as the failure rate.

Assuming that each simulation  $X$  is an independent event, and we perform  $n$  experiments, then if the probability of a failure is  $p$ , then  $X$  is a binomially distributed random variable, which follows the binomial distribution with parameters  $n$  and  $p$ . We write  $X \approx B(n, p)$ . Therefore, the mean of  $X$  will be  $\mu = np$ , and the variance of  $X$  will be  $\sigma^2 = np(1 - p)$ . By definition of Confidence Interval, in using  $\bar{X}$  to estimate  $\mu$ , the error  $E = |\bar{X} - \mu|$  is less than or equal to  $z_{\alpha/2}\sigma/\sqrt{n}$  with confidence  $100(1 - \alpha)\%$ , where  $\alpha$  is a confidence level. Therefore the sample size to be able to state with  $100(1 - \alpha)\%$  confidence that the error  $|\bar{X} - \mu|$  will not exceed a specified amount  $E$  is  $N = (\frac{z_{\alpha/2}\sigma/\sqrt{n}}{E})^2$ . Thus for 95% confidence,  $z_{\alpha/2} = 2$ , and confidence interval .01%, the sample size,  $N = 40,000p(1 - p)$ .

To estimate the probability of a failure, suppose the probability of a failure is 2%, then  $N = 40,000 \times 0.02 \times 0.98 = 784$ . If a single run of the benchmark takes an hour, finding out the failure rate of the benchmark with 95% confidence and 1% error will take a month. Clearly, the simulation time is far too much for design space exploration. Another factor that makes automated exploration difficult is that the number of simulations needed depends on the probability of a failure  $p$ . For example, if  $p$  is 2%, then we need about 784 simulations, while if  $p$  is 20%, then we need 6,400 simulations.

Thus, to efficiently choose pages to be mapped to the protected cache, we need a metric to quantitatively compare page partitions in terms of susceptibility to soft errors. We use the concept of vulnerability, proposed in [Mukherjee et al. 2003; Asadi et al. 2005; Zhang 2005a; Wang et al. 2006], to partition the data into the protected and unprotected caches in a PPC. We observe that if an error is injected in a variable that will not be used, the error does not matter. However, if the erroneous value will be used in the future, then it will result in a failure. Thus a data is defined to be **vulnerable** for the time it is in the unprotected cache until it is eventually read by the processor or written back to the memory. The vulnerability of the data in an application is just the summation of the individual data vulnerability, which is measured in cycles to present the vulnerable time of this data. Similarly, we define the vulnerability of instructions in an application as the summation of individual instruction vulnerability in the unprotected instruction cache.

To validate our idea using vulnerability as a failure rate metric, we simulated the *susan corners* benchmark for various L1 cache sizes. Fig. 3(a) plots the *vulnerability* and the *failure rate* obtained by simulations. The failure rate is calculated in % by multiplying 100 with the number of failures divided by the number of runs, and the vulnerability is measured in cycles from the modified SimpleScalar *sim-outorder* simulator for the benchmark. Fig. 3(a) shows that the shape of the *vulnerability* closely matches the failure rate curve. Other applications also show similar trends. On average, the error in predicting the failure rate using *vulnerability* metric is less than 5%. In this article, we use *vulnerability* as the metric to estimate the failure rate, and perform automated design space exploration to decide the page partitioning between the two caches of a PPC. Note that the main benefit from the vulnerability metric other than this design space automation is that we can save the exploration time by hundred to thousand times as compared to estimating the failure rate as shown in Section 4.1.

Then we need a strategy to find pages causing high vulnerability of an application and to map them into the protected cache in a PPC architecture. We have developed a simple strategy, which: (i) profiles a vulnerability for each page, and (ii) explores the partition by moving a page with the highest vulnerability among the remaining pages in the unprotected cache to the protected cache. And simple experiments have been performed to observe the impact of page mappings on the vulnerability and the runtime as shown in Fig. 3(b). Fig. 3(b) presents the vulnerability reduction when each page is mapped from the unprotected cache to the protected cache in a PPC in the descending order of page vulnerability as the failure rate reduction shown in Fig. 2.

However, reducing vulnerability can be contrary to performance improvement. For example, to reduce the vulnerability of data, data should not remain in the cache for long. It is better to evict and reload the reused data to reduce the vulnerability, but this may degrade performance. Fig. 3(b) shows the trade-off between the vulnerability and the runtime, and page partitions can incur significant runtime overheads while reducing the vulnerability. Therefore, there is a fundamental trade-off between performance improvement and vulnerability reduction in page partitions for PPC architectures. Thus, our partitioning heuristics will find

the interesting partitions while moving pages with higher vulnerability than the others from the unprotected cache to the protected cache one by one under the performance constraint.

## 4.2 Traditional Page Partitioning Techniques

Our first attempt was to apply generic search algorithms, i) Monte Carlo Method (MC), and ii) Genetic Algorithm (GA), to explore the solution space. We represent a page mapping by an  $N$  bit number, such that if the  $i^{th}$  bit of the page mapping is 1, then the  $i^{th}$  page is mapped to the protected cache in a PPC. Thus, the page mapping 00..0 represents the default case, when all pages are mapped to the unprotected cache, and the page mapping 11..1 represents the case when all pages are mapped to the protected cache.

**4.2.1 Monte Carlo Method.** Monte Carlo (MC) algorithms are non-deterministic simulation methods, usually by exploiting pseudo-random numbers. They are widely used in simulations with a large number of degrees of freedom and uncertainty. For the MC exploration, we generate each bit of the page mapping, i.e., 0 or 1, with pseudo-random numbers. Through the simulation, the page mapping is then evaluated with respect to vulnerability, performance, and energy consumption.

Note that MC ignores the impact of sequences on the vulnerability and performance when exploring the next sequences.

**4.2.2 Genetic Algorithm.** Genetic Algorithms (GA) are adaptive search algorithms using evolutionary ideas such as mutations and crossovers (recombination). Initially, we form a randomly generated sequence, representing a page mapping. At each successive generation, the superior sequences in terms of the vulnerability are selected as the evolutionary page mappings through the simulation, where vulnerability, performance, and energy consumption are evaluated. In order to generate the next sequence, we implemented two GA operations such as mutation and crossover operations. For the mutation operation, simply pseudo-random number tells whether each bit in a sequence is modified or not. For the crossover operation, one point is selected in current sequences and they are swapped on page mappings to the next generated sequences.

Note that GA ignores the performance impact of sequences explored.

## 4.3 Customized Page Partitioning Techniques

Our page partitioning techniques employ the vulnerability metric to estimate the failure rate, and they are customized to find a page partition with minimal vulnerability under the runtime constraint.

**4.3.1 PPExplore – Page Partitioning Exploration.** Fig. 4 outlines our PPExplore partitioning algorithm, which starts from the case when no page is mapped to the protected cache, i.e., all pages are mapped into the unprotected cache (default case). In each step, pages are moved from the unprotected cache to the protected cache, each partition is evaluated, and the best page partition in terms of the vulnerability reduction under the runtime penalty is selected to be mapped into the protected cache in a PPC. Our page partitioning algorithm takes two parameters: (i) allowable runtime penalty ( $rPenalty$ ), and (ii) exploration width

( $eWidth$ ), i.e., how many partitions are maintained as best configurations for the whole exploration. PPExplore uses  $pCount$ , the number of pages in a benchmark, and searches for page mappings that will suffer no more than the specified runtime penalty, while trying to minimize the vulnerability. PPExplore maintains a set of best page mappings found so far (line 05) in  $bestConfigs$ , sorted in the order of the vulnerability. After initialization, the algorithm goes into a forever loop in line 07. It takes each existing best solution and tries to improve it by mapping a page to the protected cache, and by evaluating a new page map in terms of runtime, power, and vulnerability (lines 11-12). If the new page mapping is better than the worst solution in terms of the vulnerability in the  $newBestConfigs$  with runtime satisfied, then the new page mapping is inserted in the ordered list ( $newBestConfigs$ ) according to the vulnerability (lines 13-18). The loop in lines 09-20 is one step of exploration. After each step, the new set of page mappings is trimmed down to the exploration width (lines 21-23). The termination criterion of the exploration is when an exploration step cannot find any better page mapping. In other words, no page can be mapped to the protected cache to improve vulnerability (lines 24, 26) under the runtime penalty. Otherwise, the global collection of the best page mappings are updated (line 25).

PPExplore is very effective to eventually find interesting partitions with minimal vulnerability since it explores all the possible page partitions by moving one page at each step from the beginning.

---

```

PPExplore(rPenalty, eWidth, pCount)
01:  $pageMap0 = 0..0$ 
02:  $\langle runtime, power, vulnerability \rangle = simulate(pageMap0)$ 
03:  $config0 = (pageMap0, runtime, power, vulnerability)$ 
04: for ( $k = 0; k < eWidth; k++$ )
05:    $bestConfigs.insert(config0)$ 
06: endFor
07: for (;)
08:    $newBestConfigs = bestConfigs$ 
09:   for ( $i = 0; i < eWidth; i++$ )
10:     for ( $j = 0; j < pCount; j++$ )
11:        $testConfig.pageMap = addPage(newBestConfigs[i].pageMap, j)$ 
12:        $\langle runtime, power, vulnerability \rangle = simulate(testConfig.pageMap)$ 
13:       if ( $runtime < config0.runtime \times \frac{100+rPenalty}{100}$ )
14:         if ( $vulnerability < newBestConfigs[0].vulnerability$ )
15:            $newBestConfigs.insert(testConfig.pageMap, runtime, power, vulnerability)$ 
16:         endif
17:       endif
18:     endFor
19:   endFor
20: endFor
21: for ( $i = newBestConfigs.length(); i > eWidth; i--$ )
22:    $newBestConfigs.delete[i - 1]$ 
23: endFor
24: if ( $newBestConfigs[0].vulnerability < bestConfigs[0].vulnerability$ )
25:    $bestConfigs = newBestConfigs$ 
26: else break;
27: endif
28: endFor

```

---

Fig. 4. PPExplore: an exploration algorithm for page partitioning

4.3.2 *qPPExplore – quick PPExplore*. Our PPExplore is effective to find the interesting partitions but its complexity is  $O(mN!)$ , where  $N$  is the number of

pages to be explored and  $O(m)$  is the complexity of a simulation to evaluate a page partition. PPExplore is expensive since at each step PPExplore tries all possible partitions by mapping a page from the remaining pages at the unprotected cache into the protected cache, and finds a page partition with minimal vulnerability among them. On the contrary, the complexity of qPPExplore is  $O(mN)$  as shown in Fig. 5 since it selects the partition with the least vulnerability among partitions explored by moving a page from the unprotected cache to the protected cache in the descending order of the page vulnerability, which satisfies the runtime constraint. Note that each page is queued in the descending order of page vulnerability and thus the page with the highest vulnerability is mapped into the protected cache at each step. *config0* keeps the runtime, power, and vulnerability when all pages are mapped into the protected cache in a PPC (lines 01-03). Then, qPPExplore explores a partition by mapping the page with the highest vulnerability, and selects this partition as the best (*bestConfig*) if it satisfies the runtime constraint and it has less vulnerability than the least vulnerability so far (lines 07-13). It repeats page partitioning and evaluation until all pages are mapped into the protected cache in a PPC (lines 06, 14).

qPPExplore is efficient in terms of the exploration speed to explore the large set of page partitions and also effective to find interesting partitions with minimal vulnerability in benchmarks we have studied as demonstrated in Section 6.

---

```

qPPExplore(rPenalty, pCount)
01: pageMap0 = 0..0
02:  $\langle \text{runtime}, \text{power}, \text{vulnerability} \rangle = \text{simulate}(\text{pageMap0})$ 
03: config0 = (pageMap0, runtime, power, vulnerability)
04: bestVulnerability = vulnerability
05: bestConfig = baseConfig = config0
06: for (j = (pCount - 1); j > -1; j --)
07:   baseConfig.pageMap = addPage(baseConfig.pageMap, j)
08:    $\langle \text{runtime}, \text{power}, \text{vulnerability} \rangle = \text{simulate}(\text{baseConfig.pageMap})$ 
09:   if (runtime < config0.runtime  $\times \frac{100+rPenalty}{100}$ )
10:     if (vulnerability < bestVulnerability)
11:       bestConfig = baseConfig
12:     endif
13:   endif
14: endFor

```

---

Fig. 5. qPPExplore: a quick exploration algorithm for page partitioning

4.3.3 *EPPExplore – Enhanced PPExplore*. EPPExplore enhances our exploration algorithms by combining PPExplore with qPPExplore. PPExplore begins with exploring partitions from the default case, i.e., mapping all pages into the unprotected cache, as shown in lines 01-06 in Fig. 4, and tries to improve the vulnerability by finding a page with the minimal vulnerability, which is effective but slow to explore a large set of possible partitions. However, at the initial step, EPPExplore applies qPPExplore to find the best partition in terms of vulnerability under the runtime constraint. From the partition discovered by qPPExplore, EPPExplore applies the algorithm of PPExplore to further reduce the vulnerability under the runtime constraint. Fig. 6 shows that lines 06 to 14 are from qPPExplore

and lines 15 to 38 are from PPExplore. EPPExplore can explore the page partitions which may not be explored by PPExplore since PPExplore stops exploring partitions further if it does not improve the vulnerability any longer. In fact, it is possible to reduce the vulnerability by mapping multiple pages into the protected cache without degrading performance significantly while mapping one page out of remaining pages from the unprotected cache to the protected cache in a PPC does not reduce the vulnerability. Fig. 3(b) shows this possible scenario and increasing the number of pages to be mapped into the protected cache does not keep reducing the vulnerability as shown around page 25 in Fig. 3(b). In particular, EPPExplore discovers the interesting page partitions first by qPPExplore, and further reduces the vulnerability by extensively exploring partitions with an algorithm in PPExplore. The complexity of EPPExplore is between the complexities of qPPExplore and PPExplore.

---

```

EPPExplore(rPenalty, eWidth, pCount)
01: pageMap0 = 0..0
02: < runtime, power, vulnerability >= simulate(pageMap0)
03: config0 = (pageMap0, runtime, power, vulnerability)
04: bestVulnerability = vulnerability
05: bestConfig = baseConfig = config0
06: for (j = (pCount - 1); j > -1; j --)
07:   baseConfig.pageMap = addPage(baseConfig.pageMap, j)
08:   < runtime, power, vulnerability >= simulate(baseConfig.pageMap)
09:   if (runtime < config0.runtime ×  $\frac{100+rPenalty}{100}$ )
10:     if (vulnerability < bestVulnerability)
11:       bestConfig = baseConfig
12:     endif
13:   endif
14: endFor
15: for (k = 0; k < eWidth; k ++ )
16:   bestConfigs.insert(bestConfig)
17: endFor
19: for (;)
19:   newBestConfigs = bestConfigs
20:   for (i = 0; i < eWidth; i ++ )
21:     for (j = 0; j < pCount; j ++ )
22:       testConfig.pageMap = addPage(newBestConfigs[i].pageMap, j)
23:       < runtime, power, vulnerability >= simulate(testConfig.pageMap)
24:       if (runtime < config0.runtime ×  $\frac{100+rPenalty}{100}$ )
25:         if (vulnerability < newBestConfigs[0].vulnerability)
26:           newBestConfigs.insert(testConfig.pageMap, runtime, power, vulnerability)
27:         endif
28:       endif
29:     endFor
30:   endFor
31:   for (i = newBestConfigs.length(); i > eWidth; i -- )
32:     newBestConfigs.delete[i - 1]
33:   endFor
34:   if (newBestConfigs[0].vulnerability < bestConfigs[0].vulnerability)
35:     bestConfigs = newBestConfigs
36:   else break;
37:   endif
38: endFor

```

---

Fig. 6. EPPExplore: a combined exploration algorithm of qPPExplore and PPExplore for page partitioning

## 5. SETUP

In order to demonstrate the effectiveness of our page partitioning heuristics in exploring and discovering the partition with minimal vulnerability at minimal power and runtime<sup>2</sup> penalty, we have built an extensive simulation framework. The application is first compiled to generate an executable. The application is then profiled, and the *Page Vulnerability Estimator* calculates the vulnerability of each page accessed by the application. The pages are then sorted according to their vulnerabilities, and then *Page Partitioning Heuristics* partitions and maps the pages to the two caches in the PPC architecture. Through the simulations, *Page Partitioning Heuristics* find out the page mapping with minimal vulnerability under the runtime constraint. Finally, the executable and the page mapping are provided to the platform, which runs the application and generates outputs such as runtime, energy consumption, and vulnerability.

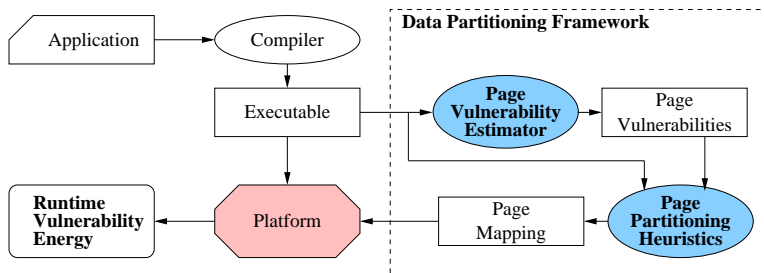


Fig. 7. Page Partitioning Exploration Framework for PPC Architectures

The platform is modeled using *sim-outorder* simulator from the SimpleScalar toolchain [Burger and Austin 1997]. The simulation parameters have been setup so as to model an HP iPAQ h4600 [Hewlett Packard] like processor memory system. We model a data PPC architecture consisting of a 4 KB unprotected cache and a 256 byte protected cache and an instruction PPC architecture of a 32 KB unprotected cache and a 2 KB protected cache. Cache parameters are set with the line size of 32 bytes, 4 way set-associativity, and FIFO (First-In First-Out) cache replacement policy. This model protects one small cache with an ECC-based technique such as a Hamming Code [Pradhan 1996]. The overheads of power and delay for ECC protected caches are estimated and synthesized using the CACTI [Shivakumar and Jouppi 2001] and the Synopsys Design Compiler [Synopsys Inc. 2001] as in [Lee et al. 2006]. And also SimpleScalar *sim-outorder* simulator has been modified to include the vulnerability computation. Thus, the modified *sim-outorder* returns the runtime and the vulnerability in cycles for each page partition. To estimate the system energy consumption, we consider the energy consumptions of the processor (including the processing pipeline and caches) and the energy consumption of the memory subsystem (including 2 off-chip SDRAMs and external buses). Thus, our model of system energy consumption ( $E$ ) consists of the energy

<sup>2</sup>Here runtime and performance are used interchangeably and represent the number of cycles for execution of an application

consumption of the processing core ( $E_{proc}$ ), that of the caches ( $E_{cache}$ ), and that of the memory subsystem ( $E_{mem}$ ).  $E_{proc}$  is estimated by multiplying the number of instructions to the power consumption per access, and  $E_{mem}$  is estimated by multiplying the number of cache misses to the power overhead per memory access. For the cache energy consumption ( $E_{cache}$ ), we detail the cache access and cache miss into read\_access\_hit, read\_access\_miss, write\_access\_hit, and write\_access\_miss since each operation results in different ECC events. For example, the energy consumption of read\_access\_hit is the sum of the access energy consumption and the energy consumption of ECC decoding while the energy consumption of read\_access\_miss is the sum of the access energy consumption and the energy consumptions of not only ECC decoding but also ECC encoding. So the energy consumption model for the cache is  $E_{cache} = RH \times d + RM \times (d + e) + WH \times e + WM \times (d + e) = (RH \times d + RM \times d + WH \times d + WM \times d - WH \times d) + (RM \times e + WM \times e + WH \times e) = A \times d + M \times e + WH \times (e - d)$  where  $RH$  is the number of read access and hit,  $RM$  is the number of read access but miss,  $WH$  is the number of write access and hit,  $WM$  is the number of write access but miss,  $A (= RH + RM + WH + WM)$  is the number of cache accesses,  $M (= RM + WM)$  is the number of cache misses, and  $d$  and  $e$  are energy consumption of ECC decoding and ECC encoding, respectively.

The HP iPAQ [Hewlett Packard] is a wireless handheld device, and MiBench [Guthaus et al. 2001] is the set of benchmarks that are representatives of applications that run on wireless handheld devices [Guthaus et al. 2001]. MiBench suite is therefore the right set of benchmarks that are supposed to run on the iPAQ, and we choose them. However, we pick only those benchmarks in which the runtime difference between the default case when all data is mapped to the 4 KB unprotected cache, and the case when all data is mapped to the 256 byte protected cache in the data PPC is more than 5%. Similarly, we select benchmarks for the instruction PPC. This is to avoid benchmarks for which only the small protected cache is enough. Note that although some of the benchmarks in MiBench are multimedia applications (for which an obvious data partitioning exists), we use our heuristics to partition the data and instructions of all applications in the selected benchmark suite.

## 6. EXPERIMENTAL RESULTS

### 6.1 Comparison of Exploration Algorithms

We bring out the details of exploration using MC (Monte Carlo), GA (Genetic Algorithm), and PPExplore with data PPCs over the *susan corners* benchmark, when PPExplore is configured for 5% performance penalty, and exploration width 2. Fig. 8(a) and Fig. 8(b) plot the runtime, energy consumption, and vulnerability of the page partitions searched by MC, GA, and PPExplore. Note that the y-axis in these graphs – the vulnerability scale – is logarithmic. The most important observation that we make from these graphs is that PPExplore searches much more useful page mappings (low vulnerability), as compared to MC and GA. We allow each technique to explore 1,900 page mappings. Thus, in total there are 5,700 page mappings. Out of them only 83 are Pareto-optimal. A page mapping is Pareto-optimal, if it is no worse than any other configuration in all the three dimensions, i.e., runtime, vulnerability and energy consumption. Out of these 83 Pareto-optimal page mappings, 68 were first drawn from PPExplore searches (82%),



12 came from GA (14%), and only 3 were discovered by MC (4%). This Pareto-optimal observation demonstrates the effectiveness of our algorithm as compared to MC and GA. The main reason for the effectiveness of PPEXplore as compared to MC and GA explorations is that MC and GA ignore the effects of partitioning on the performance and energy consumption.

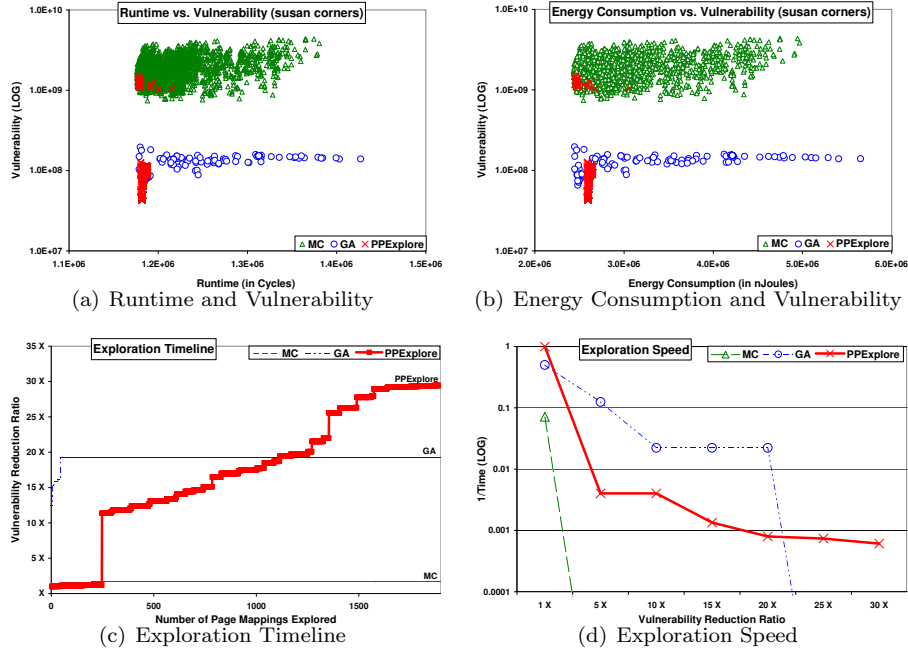


Fig. 8. Exploration by MC, GA and PPEXplore: PPEXplore effectively explores the design space

Fig. 8(c) plots the vulnerability reduction ratio as the exploration progresses for MC, GA, and PPEXplore. *Vulnerability Ratio* indicates the ratio of the vulnerability of the *default case*, i.e., all pages mapped into the unprotected cache, to the vulnerability of the page partition discovered by each exploration algorithm. A ratio greater than 1 implies the reduction of the vulnerability metric. The plot shows that while MC is ineffective, GA improves vulnerability by about 20 $\times$ , PPEXplore continuously finds better page mappings and is eventually able to reduce vulnerability by about 30 $\times$ . Since PPEXplore begins with the default page partition, PPEXplore improves the vulnerability gradually.

We also compare the speed of the various exploration algorithms. Fig. 8(d) plots the speed of exploration, i.e., inverse of the number of page partitions explored to achieve a required vulnerability reduction. The plot shows that MC is quite ineffective. Among GA and PPEXplore, GA is a faster approach when low reduction in vulnerability is required, but it is unable to achieve the high reduction in vulnerability. This is where, our approach is really effective in terms of the vulnerability reduction.

In summary, PPExplore is very effective to explore page partitions and to find out the interesting partitions to reduce the vulnerability with minimal power and performance overheads for PPCs, as compared to Genetic Algorithm and Monte Carlo method.

## 6.2 Effectiveness of Our Heuristics for Data PPC

We perform two kinds of experiments to demonstrate the effectiveness of our partitioning heuristics for page mappings into a data PPC for general applications. All the results in Fig. 9 show the reduction and overheads of page partitions discovered by our heuristics as compared to those of the default case where all data is mapped into the unprotected 4 KB cache, i.e., no protection on the data cache.

In the first set of experiments, we want to find the page partition with the least vulnerability without any performance loss (i.e.,  $pLoss = 0\%$ ) and the exploration width is set to 1 (i.e.,  $eWidth = 1$ ). Fig. 9(a) shows that PPExplore, qPPExplore, and EPPExplore find the page partitions for data PPCs, resulting in the vulnerability reduction by 16%, 25%, and 30%, respectively. Since the runtime penalty is set to 0%, the partitions we discovered incur no runtime overhead as shown in Fig. 9(c). Fig. 9(e) shows that the data partitions discovered by PPExplore, qPPExplore, and EPPExplore incur the overhead of the system energy consumption by less than 1%. Interestingly, we can find partitions for some benchmarks (e.g., *rijndael decryption* and *crc*), which not only reduce the vulnerability significantly (by about more than 95%) but also improve the performance and the system energy consumption. Note that qPPExplore only explores the number of data pages and its average over all benchmarks is about 56 while PPExplore explores 627 partitions and EPPExplore 401 partitions, on average. Thus, under no runtime penalty, qPPExplore is better than PPExplore in terms of the vulnerability reduction with minimal partitions explored, and EPPExplore is the best in terms of the vulnerability reduction and EPPExplore searches the less number of partitions than PPExplore.

In the next experiment, we allow 5% performance degradation and the exploration width of 1. Fig. 9(b) plots the vulnerability reduction, and Fig. 9(d) and Fig. 9(f) plot the increase in runtime and the increase in the system energy consumption of the least vulnerability page partitions obtained by PPExplore, qPPExplore, and EPPExplore. Vulnerability reductions achieved by PPExplore, qPPExplore, and EPPExplore are 56%, 48%, and 53%, respectively. Unfortunately, we observe very small amount of vulnerability reductions for benchmarks such as *rijndael encryption*, *blowfish decryption*, and *blowfish encryption*. This result is because some pages are very sensitive to the size of the cache and most pages do not cause any reduction of the vulnerability when they are moved from the unprotected cache to the protected cache in a data PPC. And no higher reduction in vulnerability has been observed with any other exploration algorithms such as MC and GA for those benchmarks. Note that qPPExplore only explores the number of data pages (56) while PPExplore explores 1,190 partitions and EPPExplore 401 partitions, on average over all benchmarks. Thus, qPPExplore and EPPExplore are very efficient in terms of the exploration speed compared to PPExplore while PPExplore is the best in terms of the vulnerability reduction in this set of experiments. Energy consumption overheads are 21%, 30%, and 34% for PPExplore, qPPExplore, and EPPExplore, respectively, while the runtime overheads are less than 5% over all the

benchmarks. Thus, even very small runtime degradation allows our heuristics to find page mappings that can significantly reduce the vulnerability. Note that they incur relatively high energy consumption overheads (in particular, the partitions for *fft* and *stringsearch* benchmarks incur more than 100%) since our heuristics do not restrict the energy consumption overhead. We can tradeoff the vulnerability for the reduction of the energy consumption. The experimental results with EPPEXplore when both the energy consumption (10%) and the runtime (5%) are limited show that the vulnerability reduction is decreased from 53% to 48% at 2% runtime overhead and 7% energy consumption overhead, on average.




Fig. 9. Evaluation of vulnerability, runtime, and energy consumption under different performance penalty with our partitioning heuristics for data PPCs

Table I. Runtime overheads and vulnerability reductions for different data inputs (benchmark - *susan corners*)

Data Input	Runtime Overhead (%)	Vulnerability Reduction (%)
small ( <i>default</i> )	4.4	52.3
balloons 1	0.7	35.6
balloons 2	5.4	51.3
columns 1	1.6	25.2
columns 2	3.6	42.6
feep 1	0.2	64.4
large 1	4.4	57.6
large 2	3.7	26.7
AVERAGE	2.8	43.3

Note that these experimental results are obtained with sample inputs coming with benchmarks and our profile-based page partitioning heuristics work well if the page mapping of applications codes and input data do not change. To observe the impact of different input data, we ran another set of experiments for a benchmark *susan corners* with the best page partition discovered by PPEXplore under the 5% runtime penalty for different data inputs. Table I shows the effectiveness of our exploration technique on different data inputs. Other than the default data input (*small*), the vulnerability reduction for 7 different data inputs ranges from 25% to 64% and the average reduction is about 43%, which is a little bit less than the vulnerability reduction (52%) with the default input in our experiments. These results show that the page partitions for data PPCs discovered by our heuristics can reduce the vulnerability not only with simulated data input but also with different data inputs. Under several different inputs for other benchmarks, we can observe the similar results and also observe less reduction of vulnerability in benchmarks with different data inputs.

In summary, the best page partitions discovered by our exploration heuristics show the vulnerability reduction by 56% with minimal overheads of runtime (by 2%) and energy consumption (by 21%) over benchmarks as compared to the default case when all data is mapped into the unprotected cache, and our exploration heuristics effectively expand the applicability of data PPCs in general.

### 6.3 Effectiveness of Our Heuristics for Instruction PPC

To explore the page partitions for instruction PPCs, similarly we employ PPEXplore, qPPEXplore, and EPPEXplore under no performance penalty and 5% performance penalty. For this set of experiments, PPEXplore and EPPEXplore are configured with exploration width of 2. In these experiments, we added a benchmark *dijkstra* instead of *djpeg*, *crc*, and *sha* since they show less than 5% runtime difference between the default case when mapping all instructions into the 32 KB unprotected cache and the case when mapping all instructions into the 2 KB protected cache in an instruction PPC. Under no runtime penalty, Fig. 10(a), Fig. 10(c), and Fig. 10(e) show that qPPEXplore discovers the page partitions for instruction PPCs with 13% reduction of vulnerability at the cost of 2% energy consumption with no runtime overhead as compared to the default case, PPEXplore discovers them with 22% reduction of vulnerability and 2% energy consumption overhead, and EPPEXplore discovers them with 21% reduction of vulnerability




Fig. 10. Evaluation of vulnerability, runtime, and energy consumption under different performance penalty with our partitioning heuristics for instruction PPCs

and 2% energy consumption overhead, on average. Under 5% runtime penalty, Fig. 10(b), Fig. 10(d), and Fig. 10(f) show that the page partitions discovered by qPPExplore reduce the vulnerability by 27% with 1% overhead of runtime and 3% overhead of energy consumption, the page partitions discovered by PPExplore reduce the vulnerability by 49.9% with 2.3% runtime overhead and 8.2% energy consumption overhead, the page partitions discovered by EPPExplore reduce the vulnerability by 49.6% with 2.9% runtime overhead and 12.8% energy consumption overhead. In terms of the efficiency (i.e., the number of partitions explored), while qPPExplore evaluates about 56 partitions (the average number of instruction pages

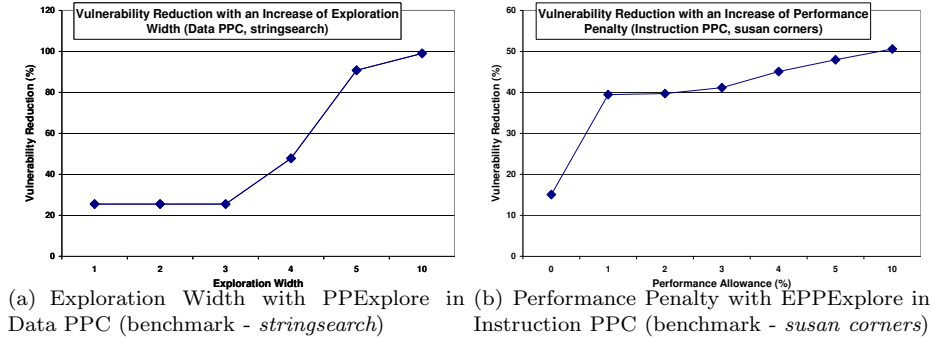


Fig. 11. Sensitivity of *exploration width* and *performance penalty* to vulnerability reduction

among benchmarks), PPExplore evaluates 681 partitions and 1,457 partitions, and EPPExplore explores 593 and 1,299 partitions, under no performance penalty and under 5% performance penalty, respectively. Thus, qPPExplore is the most efficient, and EPPExplore explores the less partitions than PPExplore while finding the effective pages in terms of the vulnerability reduction close to those of PPExplore. However, the partitions discovered by PPExplore are the most effective in terms of the vulnerability reduction with minimal overheads of runtime and energy consumption.

In summary, the results that our heuristics obtain are very effective since the instruction PPCs with page partitions discovered by our heuristics can reduce the vulnerability by more than half with less than 3% performance overhead and 8% to 13% overhead of the system energy consumption. Our page partitioning heuristics effectively expands the applicability of PPC architectures for instruction caches as well as for data caches.

#### 6.4 Sensitivity of Vulnerability Reduction

We also study the effectiveness of the vulnerability reduction with our heuristics by varying the allowable performance penalty and the exploration width. Fig. 11(a) shows that as we increase the exploration width from 1 to 10 with PPExplore, the vulnerability reduction increases with the benchmark, *stringsearch*, on data PPCs. Note that this experiment limits the runtime penalty as 0% and it is the very effective result (up to 99% vulnerability reduction under no performance penalty). However, increasing the exploration width can increase the number of partitions explored by up to *eWidth* times more than the case with the exploration width 1.

Fig. 11(b) shows that as we increase the allowable performance penalty, from 1% to 10%, the vulnerability reduction of benchmark, *susan corners*, on instruction PPC increases. Note that allowing the more performance penalty in our heuristics incurs more overheads in terms of performance and energy consumption. However, this performance penalty parameter can increase the effectiveness to find the interesting partitions with least vulnerability as shown in Fig. 11(b).

In summary, when increasing the exploration width and the allowable runtime penalty in our heuristic algorithms such as PPExplore and EPPExplore, the in-

crease of the vulnerability reduction has been observed in most benchmarks for instruction and data PPCs. Thus, we can definitely tradeoff the exploration time for the vulnerability reduction of applications.

## 7. SUMMARY

Owing to the incessant technology scaling, soft errors, especially in caches, are becoming a critical design concern for the reliability of embedded systems. Partially Protected Cache (PPC) architecture has been proposed as an effective microarchitectural means of improving the system reliability with minimal power and performance penalty for resource-constrained embedded systems. However, the challenge is in partitioning pages among the two caches in a PPC. While page partitioning schemes have been proposed for multimedia applications, there is no page partitioning scheme not only for data PPCs but also for instruction PPCs in general. The page partitioning space is huge, and existing random techniques are unable to identify and explore the page partitions that lead to low vulnerability. In this article, we develop page partitioning heuristics such as PPExplore, qPPExplore, and EPPExplore at design time that effectively and efficiently find page partitions resulting in, on average, 48% reduction in vulnerability (i.e., failure rate) at only 2% performance and 7% energy penalty for data PPCs, and 49% reduction in vulnerability at only 2% performance and 13% energy penalty for instruction PPCs over benchmarks.

The main contribution of our partitioning heuristics is that they increase the applicability of PPC architectures and establish PPC as the hardware/software hybrid solution of choice to improve the reliability of cache-based architectures. Our future work includes intelligent schemes to improve the page partitioning in PPCs and dynamic schemes to relocate page partitions at run time. Also, we plan to investigate partitioning techniques for more than two caches with different levels of protection for PPC architectures.

## REFERENCES

- ANGHEL, L. AND NICOLAIDIS, M. 2000. Cost reduction and evaluation of a temporary faults detecting technique. In *IEEE/ACM Design, Automation and Test in Europe Conference (DATE)*. 591–597.
- ASADI, G.-H., SRIDHARAN, V., TAHOORI, M. B., AND KAEI, D. 2005. Balancing performance and reliability in the memory hierarchy. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 269–279.
- BAUMANN, R. 2005. Soft errors in advanced computer systems. *IEEE Design and Test of Computers*, 258–266.
- BAZE, M. P., BUCHNER, S. P., AND MCMORROW, D. 2000. A digital CMOS design technique for SEU hardening. *IEEE Trans. on Nuclear Science* 47, 6 (Dec), 2603–2608.
- BURGER, D. AND AUSTIN, T. M. 1997. The SimpleScalar Tool Set, version 2.0. *SIGARCH Computer Architecture News* 25, 3, 13–25.
- CHEN, G., KANDEMIR, M., IRWIN, M. J., AND MEMIK, G. 2005. Compiler-directed selective data protection against soft errors. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*.
- ERNST, D., KIM, N. S., DAS, S., PANT, S., RAO, R., PHAM, T., ZIESLER, C., BLAAUW, D., AUSTIN, T., FLAUTNER, K., AND MUDGE, T. 2003. Razor: A low-power pipeline based on circuit-level timing speculation. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 7–13.

- GAISLER, J. 1997. Evaluation of a 32-bit microprocessor with builtin concurrent error-detection. In *IEEE International Symposium on Fault-Tolerant Computing (FTCS)*.
- GUTHAUS, M., RINGENBERG, J., ERNST, D., AUSTIN, T., MUDGE, T., AND BROWN, R. 2001. MiBench: A free, commercially representative embedded benchmark suite. In *IEEE Workshop on Workload Characterization*. 3–14.
- HAZUCHA, P. AND SVENSSON, C. 2000. Impact of CMOS technology scaling on the atmospheric neutron soft error rate. *IEEE Trans. on Nuclear Science* 47, 6, 2586–2594.
- Hewlett Packard. *HP iPAQ h4000 Series - System Specifications*. Hewlett Packard, <http://www.hp.com>.
- ITRS. 2005. *International Technology Roadmap for Semiconductors 2005 Executive Summary*. <http://www.itrs.net/Links/2005ITRS/ExecSum2005.pdf>.
- KIM, S. 2006. Area-efficient error protection for caches. In *IEEE/ACM Design, Automation and Test in Europe Conference (DATE)*. 1282–1287.
- KRISHNAMOHAN, S. AND MAHAPATRA, N. R. 2004. An efficient error-masking technique for improving the soft-error robustness of static CMOS circuits. In *IEEE International SOC Conference (SOCC)*. 227–230.
- KRUEGER, D., FRANCOM, E., AND LANGSDORF, J. 2008. Circuit design for voltage scaling and SER immunity on a quad-core Itanium processor. In *IEEE International Solid-State Circuits Conference (ISSCC)*. 94–95.
- LEE, K., SHRIVASTAVA, A., ISSENIN, I., DUTT, N., AND VENKATASUBRAMANIAN, N. 2006. Mitigating soft error failures for multimedia applications by selective data protection. In *International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES)*. 411–420.
- LI, J.-F. AND HUANG, Y.-J. 2005. An error detection and correction scheme for RAMs with partial-write function. In *IEEE International Workshop on Memory Technology, Design and Testing (MTDT)*. 115–120.
- LI, L., DEGALAHAL, V., VIJAYKRISHNAN, N., KANDEMIR, M., AND IRWIN, M. J. 2004. Soft error and energy consumption interactions: A data cache perspective. In *International Symposium on Low Power Electronics and Design (ISLPED)*. 132–137.
- LIDEN, P., DAHLGREN, P., JOHANSSON, R., AND KARLSSON, J. 1994. On latching probability of particle induced transients in combinational networks. In *IEEE International Symposium on Fault-Tolerant Computing (FTCS)*.
- LUCCHETTI, D., REINHARDT, S. K., AND CHEN, P. M. 2005. ExtraVirt: Detecting and recovering from transient processor faults. In *ACM Symposium on Operating Systems Principles (SOSP)*. ACM Press, New York, NY, USA, 1–8.
- MASTIPURAM, R. AND WEE, E. C. 2004. *Soft Errors' Impact on System Reliability*. <http://www.edn.com/article/CA454636>.
- MITRA, S., SEIFERT, N., ZHANG, M., SHI, Q., AND KIM, K. S. 2005. Robust system design with built-in soft-error resilience. *IEEE Computer* 38, 2 (Feb), 43–52.
- MOHANRAM, K. AND TOUBA, N. A. 2003. Partial error masking to reduce soft error failure rate in logic circuits. In *IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT)*. 433–440.
- MOHR, K. AND CLARK, L. 2006. Delay and area efficient first-level cache soft error detection and correction. In *IEEE International Conference on Computer Design (ICCD)*.
- MUKHERJEE, S. S., EMER, J., FOSSUM, T., AND REINHARDT, S. K. 2004. Cache scrubbing in microprocessors: Myth or necessity? In *IEEE Pacific Rim International Symposium on Dependable Computing (PRDC)*. 37–42.
- MUKHERJEE, S. S., WEAVER, C., EMER, J., REINHARDT, S. K., AND AUSTIN, T. 2003. A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 29–40.
- MUSSEAU, O. 1996. Single-event effects in SOI technologies and devices. *IEEE Trans. on Nuclear Science* 43, 2 (Apr), 603–613.
- NICOLAIDIS, M. 1999. Time redundancy based soft-error tolerance to rescue nanometer technologies. In *IEEE VLSI Test Symposium (VTS)*. 86.



- NIEUWLAND, A. K., JASAREVIC, S., AND JERIN, G. 2006. Combinational logic soft error analysis and protection. In *IEEE International Symposium on On-Line Testing (IOLTS)*. 99–104.
- PHELAN, R. 2003. Addressing soft errors in ARM core-based designs. Tech. rep., ARM.
- PRADHAN, D. K. 1996. *Fault-Tolerant Computer System Design*. Prentice Hall. ISBN 0-1305-7887-8.
- QUACH, N. 2000. High availability and reliability in the Itanium processor. *IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 61–69.
- REIS, G. A., CHANG, J., VACHHARAJANI, N., RANGAN, R., AND AUGUST, D. I. 2005. SWIFT: Software implemented fault tolerance. In *International Symposium on Code Generation and Optimization (CGO)*. IEEE Computer Society, Washington, DC, USA, 243–254.
- ROCHE, P., GASLOT, G., FORBES, K., OAPOS, SULLIVAN, V., AND FERLET, V. 2003. Comparisons of soft error rate for SRAMs in commercial SOI and bulk below the 130-nm technology node. *IEEE Trans. on Nuclear Science* 50, 6 (Dec).
- SHIVAKUMAR, P. AND JOUPPI, N. 2001. CACTI 3.0: An Integrated Cache Timing, Power, and Area Model. In *WRL Technical Report 2001/2*.
- SHIVAKUMAR, P., KISTLER, M., KECKLER, S., BURGER, D., AND ALVISI, L. 2002. Modeling the effect of technology trends on soft error rate of combinational logic. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 389–398.
- STACKHOUSE, B., CHERKAUER, B., GOWAN, M., GRONOWSKI, P., AND LYLES, C. 2008. A 65nm 2-billion-transistor quad-core Itanium processor. In *IEEE International Solid-State Circuits Conference (ISSCC)*. 92–93, 598.
- Synopsys Inc. 2001. *Design Compiler Reference Manual*. Synopsys Inc., Mountain View, CA, USA.
- WANG, S., HU, J., AND ZIAVRAS, S. G. 2006. On the characterization of data cache vulnerability in high-performance embedded microprocessors. In *IEEE International Conference on Embedded Computer Systems, Architecture, Modeling, and Simulation (SAMOS)*. 14–20.
- WROBEL, F., PALAU, J. M., CALVET, M. C., BERSILLON, O., AND DUARTE, H. 2001. Simulation of nucleon-induced nuclear reactions in a simplified SRAM structure: Scaling effects on SEU and MBU cross sections. *IEEE Trans. on Nuclear Science* 48, 6, 1946–1952.
- ZHANG, W. 2005a. Computing cache vulnerability to transient errors and its implication. In *IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT)*. 427–435.
- ZHANG, W. 2005b. Replication cache: A small fully associative cache to improve data cache reliability. *IEEE Computers* 54, 12 (Dec), 1547–1555.
- ZHANG, W., GURUMURTHI, S., KANDEMIR, M., AND SIVASUBRAMANIAM, A. 2003. ICR: In-cache replication for enhancing data cache reliability. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*.
- ZORIAN, Y., VARDANIAN, V. A., ALEKSANYAN, K., AND AMIRKHANYAN, K. 2005. Impact of soft error challenge on SoC design. In *IEEE International Symposium on On-Line Testing (IOLTS)*. 63–68.