Microprocessors and Microsystems 32 (2008) 244–253

Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

On-chip cache device scaling limits and effective fault repair techniques in future nanoscale technology

David Roberts [a,*], Nam Sung Kim [b], Trevor Mudge [a]

[a] University of Michigan, Ann Arbor, MI 48109, United States
[b] Intel Corporation, United States

ARTICLE INFO

Available online 26 April 2008

Keywords:
On-chip cache
Device scaling
Fault-tolerance
DVS

ABSTRACT

In this study, we investigate different cache fault-tolerance techniques to determine which will be most effective when on-chip memory cell defect probabilities exceed those of current technologies, which is highly anticipated in processor on-chip caches manufactured with future nanometer scale technologies. Our most significant finding from this study is that the devices in on-chip memory cells cannot be scaled at the same rate as devices in logic circuits due to the increasing number of erroneous memory cells with voltage scaling, requiring strong fault-tolerance techniques. Second, we propose a technique to minimize performance impacts under aggressive technology and voltage scaling. It works by merging pairs of faulty cache lines into good lines and performs better than TMR at high error rates. We also estimate up to 28% energy savings at low voltage, relative to a recent fault-tolerance scheme [A. Agarwal et al. A process-tolerant cache architecture for improved yield in nanoscale technologies. IEEE Trans. VLSI Syst. 13(1) (2005) 27–38].

1. Introduction

As microarchitects demand larger on-chip caches for higher performance, continuous device scaling has provided improved memory density for multi-megabyte upper-level on-chip caches at a reasonable die cost. However, the device scaling comes at a price. The reduced device feature size causes exponentially increasing subthreshold and gate-leakage power problems in on-chip caches fabricated with sub-90 nm process technology resulting in more static power consumption [2]. Furthermore, process parameter variations, e.g. random dopant fluctuations causing threshold voltage variations or mismatches across the devices used in a on-chip memory cell and more oxide defects in devices during the manufacturing process have worsened yield problems in on-chip caches manufactured with sub-90 nm technology [3].

To overcome low yield problems caused by scaling device sizes and integrating more on-chip memory cells, there have been several proposed techniques. One is to implement redundant memory columns; there are one or two redundant columns per memory sub-bank or sub-array. If a defective cell is found during the manufacturing test, the entire column containing the defective cell is replaced with a redundant column. This wastes many memory cells to fix one defective cell and requires fuses to replace the column containing the defective cell with the redundant column. The second technique is to use error correction codes (ECC). Currently, a single error correction (SEC) and double error detection (DED) technique is used. Even though this can fix one defective cell per sub-array row, the memory array is made more vulnerable to soft errors since the correction capability of the code has been used up by fixing defective memory cells. The third technique is to disable a part of the on-chip cache memory array resulting in a smaller size. An example is the Intel Celeron processor. It is very similar to the Pentium processor, but it has the entire or half of the L2 cache disabled as a result of memory sub-arrays containing defective cells that could not be fixed using the redundant columns in the disabled part of the on-chip cache memory block. All these techniques are only effective when there are a small number of defective cells in the on-chip cache. However, the number of defective cells in large on-chip caches will rise if we want to continue scaling memory cell size along with technology scaling.

Hard-wired redundancy is becoming a less attractive option due to limited area available for spare memory cells. In addition, it will no longer be possible to find a single set of cache blocks which consistently fail at each operating point [3]. Prior work suggests that avoiding defective cache memory cells at the block level can be very cost-effective in terms of both area and performance overheads. However, these studies were performed with either outdated cache hierarchies and benchmarks [4] or for direct-mapped caches only [1]. Under aggressive voltage scaling and on-chip memory cell sizing, we show that higher defect rates with existing

* Corresponding author.
E-mail addresses: daverobe@umich.edu (D. Roberts), nam.sung.kim@intel.com (N.S. Kim), tnm@eecs.umich.edu (T. Mudge).

0141-9331/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.micpro.2008.03.012

fault-tolerance schemes result in significant processor performance degradation. A dynamic voltage scaling (DVS) environment adds to the complexity of working with on-chip caches containing unpredictable defective memory cells; as the operating voltage changes, so does the number of defective cells.

In this paper we begin with an analysis of L2 cache activity in a modern processor architecture based on the Intel Pentium 4. Emphasis is placed on L2 caches because of their widespread use and relatively large area compared to L1 (L1 caches are also relevant, and the error analysis within this paper can also be applied to other levels besides L2). We show the impacts of defective cache blocks on performance and compare ways of addressing this problem. The major contributions of this paper are;

- Trade-off analysis between performance and area for different cell sizes and fault-tolerance techniques.
- A novel cache block grouping scheme for good performance at higher fault probabilities.

The rest of the paper is organized as follows. Section 2 presents related work and explains in detail the basic fault-tolerance scheme upon which this work is based. Section 3 explains the problems encountered with on-chip cache memory reliability in new processes and its impact on performance of set-associative caches when defects are present. Based upon this analysis, we show existing and proposed techniques of reducing performance impacts in the presence of defects in on-chip caches in Section 4. The techniques are compared in Section 5, and concluding remarks are presented in Section 6.

2. Related work

Pour and Hill [4] derive an analytical model of the performance loss of a set-associative cache given a set of defective blocks. They employ an extra "valid" bit per cache block to identify whether or not it is defective. Their key findings for caches of size up to 32 kB are that miss ratio increase is negligible unless a set is completely disabled by faults.

In [5] they present a model to estimate memory-failure probability using combined row and column redundancy. The Power4 architecture [6] employs parity on L1 caches and Hamming codes on L2. In addition, L1 and L2 have spare bits, while L3 has redundant cache lines. If correctable error thresholds are exceeded, a cache line delete function allows up to 2 deletions per L3 cache. For defects detected at power-on BIST that cannot be handled, the L3 cache is disabled.

The Nanobox [7] applies redundancy and other ECC codes to logic functions built using lookup tables.

A technique for memory self-repair at high defect densities is presented in [8]. It relies upon prior knowledge of the polarity of the error (i.e. faults are always stuck at 0 or 1). In our cache application, the scheme will not work because the value read from faulty bits is unpredictable and can change with operating point (e.g. voltage or temperature).

Agarwal et al. [1] noted that the number of defective cells and their location changes depending on operating voltage. In addition, they proposed a cache block re-mapping technique for direct-mapped caches. The technique relies on a defective block mapping table determined prior to execution using BIST. They consider the use of block re-mapping in conjunction with ECC and row redundancy. Because we often refer to this scheme in the paper, it is explained in more detail as follows. Fig. 1 is the same figure as Fig. 7 in [1] and illustrates the fault tolerance scheme presented in that work. It is based on a direct-mapped cache consisting of lines organized in rows and columns. Rows

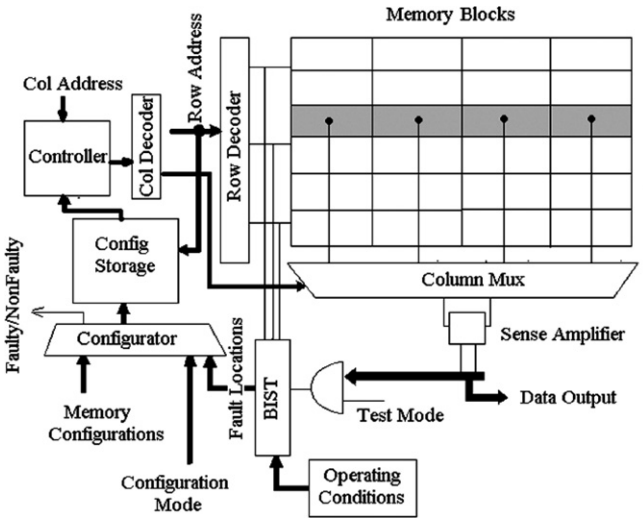

Fig. 1. The one bit implementation (OBI) technique.

are addressed as usual using part of the incoming address. However, the column address may be re-mapped to avoid a known faulty block. This is achieved by performing a look-up in the "config storage" which contains a map of defective block locations. In this instance, there is one bit per block (hence one bit implementation, or OBI) which is set to 1 if the corresponding block is defective. When the cache is accessed, the controller uses the OBI to select a non-defective column to store data to, using a fixed mapping. Additional bits are required in the tag to indicate the column in which data is stored. This prevents faulty blocks from being read.

3. Impact of on-chip cache failure rate on processor performance

3.1. On-chip cache device scaling and failure rate

Currently, the feature sizes (e.g. 45 nm) are so small that it is very difficult to control the uniformity of device parameters across dies and wafers. In particular, smaller devices that are extensively used in on-chip cache memory cells are increasingly sensitive to parameter variations. Furthermore, dynamic voltage scaling is very commonly used to reduce power consumption of the processors and their on-chip caches should be able to operate at the same voltage as the processor core, to avoid adding overhead to allow separate voltage domains. However, as the supply voltage of on-chip cache memory cells decreases, we find more memory cells failing due to increased sensitivity to process variation at lower supply voltage. Failure types are read failures (flipping of the stored state during read operations), write failures (inability to write a state during write operations), access time failures (an increase in the access time of the cell resulting in the violation of the delay requirement), and/or retention failure (losing the stored state in standby mode) [10,11]. As a result, the lowest operating voltage (called V_{cc}min) of processors employing DVS is usually determined by the lowest supply voltage that keeps all on-chip memory cells functional. However, lower V_{cc}min is desirable if static and dynamic power consumption are to be reduced. The best way to improve V_{cc}min is to increase memory cell size to reduce the process variation sensitivity of the memory cells. However, a larger memory cell size increases the area occupied by on-chip caches resulting in increasing die cost or decreasing the on-chip cache size at a given die size (e.g. 12 MB instead of 16 MB for a

L3 cache). Hence, the memory cell size must be balanced to give proper V_{cc}min, yield, and on-chip cache capacity.

Fig. 2 shows normalized on-chip memory cell failure rates for 3 different memory cell sizes (A, B and C with relative areas of 1, 1.25, and 1.5, respectively) as a function of memory cell supply voltage. Note that depending on how we tune the sizes of the six transistors in the memory cell, the result varies significantly. They are obtained using Monte-Carlo simulations on memory cells designed with a 45 nm technology and process parameter variations corresponding to the technology. The failure rate we assume is significantly higher than the data presented in other work [1], however the increased failure rates can be expected in future smaller semiconductor process technology (e.g., 32 nm technology). As shown in Fig. 2, as either voltage or cell size decreases, the failure rate starts to increase exponentially. In other words, a larger cell can achieve a much lower V_{cc}min at the same failure rate. Finally, defect rate is proportional to die size. Hence, when we integrate more on-chip memory cells on a die along with device scaling, there will be a much greater chance that some memory cells contain defects and fail during post-manufacturing tests resulting in poor yield. The next section examines the relationship between on-chip cache memory cell failure and performance impact to determine the number of tolerable, non-corrected faults.

3.2. Performance impact of set-associative cache defects

The performance impact of on-chip cache memory cell failure partly depends upon the fault-tolerance technique employed. While some impact arises from increased miss rates as faulty regions are disabled, others incur a performance penalty when error-correcting codes are decoded.

We perform an analysis of block level fault-tolerance schemes. All data was obtained using the M5 simulator [12]. The simulator was configured to represent a modern out-of-order pipeline with similar specifications to a Pentium 4 (Table 1). The memory latency is relatively low, although this will not significantly affect L2 miss rates.

Fig. 3a and b present the impact of defective blocks in L2 (each containing 1 or more defective cells) on miss rate and instructions per cycle (IPC), respectively. We use the SPLASH-2 benchmark suite [13] as a workload representative of both memory and compute-intensive applications. In these graphs, defective block locations are allocated randomly, but consistently between benchmarks. The LRU scheme was modified so that defective (non-correctable) blocks are not considered for replacement. If all ways in a set are defective, accesses to that set bypass the cache and are forwarded to the next level of the memory hierarchy. This has the effect of reducing the number of available ways in a cache set, while using the standard tag matching mechanism to determine hits or misses within the remaining good blocks, which is identical to [14]. The data in Fig. 3a and b confirm the previous study [4] in that high block failure rates are required before there is any significant performance penalty. With this in mind, the next section compares existing and our fault-tolerance techniques for their performance and area at significant failure rates.

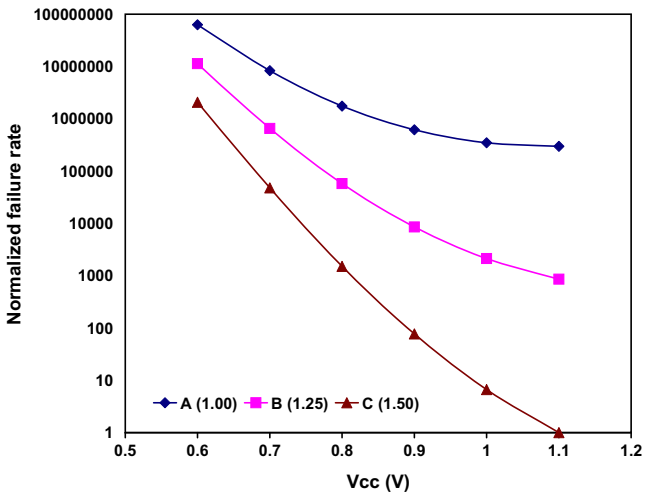

Fig. 2. Normalized cell failure rates as a function of voltage for three different memory cells.

Table 1
M5 CPU configuration (2 GHz clock)

Parameter	Value
Pipeline width	4
Branch prediction/BTB	Hybrid 4-way, 2 K entries
ROB/LSQ size	196/32 entries
INT ALUs/multi-divs/mem ports	6/2/4
FP ALUs/multi-divs	4/2
Functional unit latencies	INT: mul 3, div 20, all others 1 FP: adder 2, mul 4, div 12, sqrt 24
IL1 cache	16 kB, 2-way, 64B blocks, 1-cycle lat.
DL1 cache	16 kB, 2-way, 64B blocks, 3-cycle lat.
L2 cache	1MB, 8-way,64B blocks, 19-cycle lat.
Memory bus/latency	16 bytes with 6-cycle lat./100 cycles

4. Comparison of fault-tolerance techniques

The one bit implementation (OBI) mapping table model [1] is effective for low failure rates, but for higher rates we show that it rapidly becomes ineffective. We aim to allow more faults with stronger error correction, and observe the trade-off with area cost. In this section we derive an analytical model representing the fraction of good blocks remaining in the cache at different cell failure probabilities (and sizes). We consider tag bits as additional bits contained in each block.

We compared several cache fault-tolerance schemes in order to determine their area efficiency at different error rates (voltages and cell sizes). The model used represents the fraction of fault-free blocks available in the cache, denoted by F_{avail}. As a minimum, we decided to first apply the OBI scheme, followed by other error correction. From a storage standpoint, OBI provides the minimum data needed to identify where faulty blocks are, for avoidance. Since it has already been proven superior to redundant rows and SECDED ECC, all of our models build on the OBI baseline. OBI does not affect the cache access time and has minimum effect on processor performance [1]. Throughout the paper we refer to p_{fault} as the probability of failure of a single on-chip memory cell. Cell failures are assumed to be independent. As a first approximation this assumption is valid and has been widely used in other published cache-error related work ([1,5,8]).

To improve the effectiveness of schemes requiring an additional storage table (e.g. OBI) which must contain correct bits, we introduce a factor OBI_{ff}. This factor reduces the bit failure probability, representing larger size or higher voltage on-chip cache memory cells (see Fig. 2) used specifically for that table. We refer to this as "guaranteed correct" storage because one fault in this table could lead to bad cells being accessed. Using large cells is viable as long as the table does not contain too much data. In addition, delay does not vary significantly with cell size. "Basic storage" refers to the cells used in the cache itself.

The fault-tolerance schemes (Table 2) were chosen from a range of candidates, most of which are widely used today. We only model storage-related reliability while logic reliability is beyond the

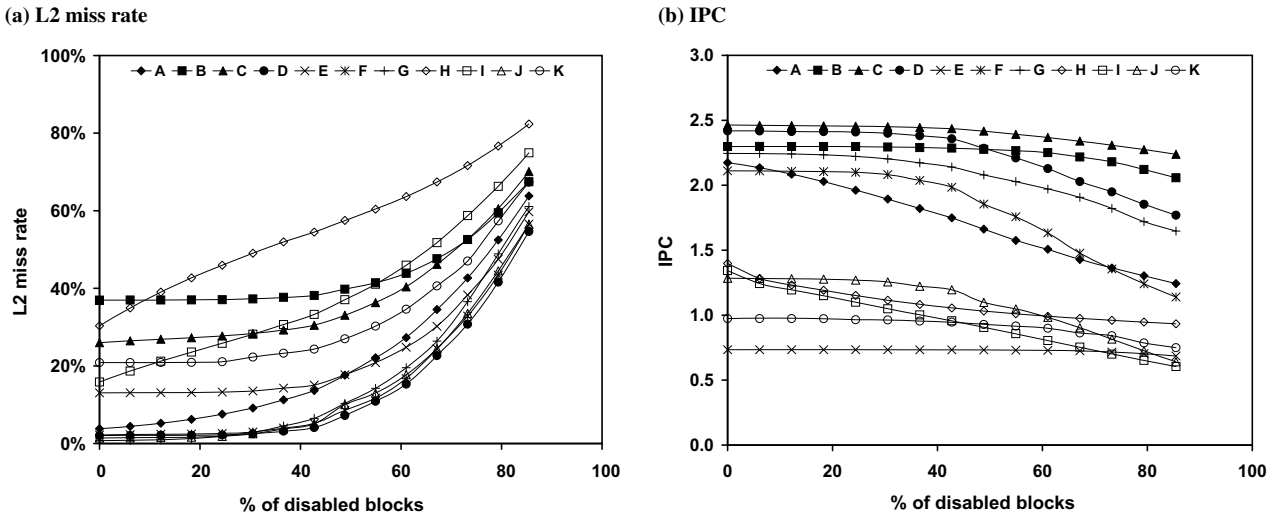

Fig. 3. L2 miss rate in (a) and IPC in (b) as a function of number of randomly disabled blocks. A-Cholesky, B-FFT, C-LUContig, D-LUNoncontig, E-Radix, F-Barnes, G-FMM, H-OceanContig, I-OceanNoncontig, J-Raytrace, and KWaterNSquared, respectively.

Table 2
Candidate cache fault-tolerance schemes and their storage overheads for an 1MB cache

Schemes	Storage	Description
OBI	1.002	One bit implementation from [1] indicates a faulty block with a single bit
Hamming SEC	1.020	Single error-correcting (SEC) hamming code
BCH DEC	1.037	Double error-correcting (DEC) Bose-Chaudhury-Hocquenghem
log(B)	1.022	Bad block table contains index of one bad bit and a spare cell to store the value of that faulty bit. If there is more than one defective bit, the block is disabled
Triple modular redundancy (TMR)	1.002	Faulty blocks are combined in groups of 3 inside the 1MB cache with a majority vote on each bit
Block grouping (GRP2)	Up to 1.528	Pairs of faulty blocks are combined to form single good blocks. A paired block is 'good' if there is only one faulty bit for each corresponding pair of 2 bits

scope of this paper. The following sections explain the fault model and storage overhead of each scheme.

Each cache consists of M sets and N ways where each block contains B bits (including tag bits).

We also derive an area efficiency E_{area} which takes into account the probability of failure of the "guaranteed correct" storage which includes the OBI table and any additional bits added by a scheme which must be correct for the cache to operate reliably. The fraction of available blocks is divided by die area consumed by all SRAM cells, then scaled by the probability of the guaranteed correct storage containing no faults Eq. (1).

$$E_{\text{area}} = \frac{F_{\text{avail}}}{\text{area}} \times p_{\text{non\_faulty\_GC}} \tag{1}$$

The $p_{\text{non\_faulty\_GC}}$ value is the probability that the guaranteed correct cells are fault-free, as a function of the probability of the large-size cell failure $p_{\text{fault\_GC}}$ and the number of guaranteed correct bits (GC_bits).

$$p_{\text{non\_faulty\_GC}} = (1 - p_{\text{fault\_GC}})^{\text{GC\_bits}} \tag{2}$$

In all of these schemes,

$$p_{\text{fault\_GC}} = \text{OBI}_{\text{ff}} \times p_{\text{fault}} = 10^{-5} \times p_{\text{fault}}$$

4.1. Existing fault-tolerance schemes

4.1.1. OBI

The "one bit implementation" consists of a table of bits, one per block, indicating whether or not each block contains 1 or more faulty bits. All of our schemes incur this storage overhead, because we use an OBI to indicate whether a block can be corrected or is unusable and cannot be accessed.

4.1.1.1. Storage overhead

GC_bits $= M \times N$ bits (guaranteed correct storage).

4.1.1.2. Fault model. The probability of a faulty bit is p_{fault}. The probability of a non-faulty block is the likelihood of every bit being fault-free in that block. We assume that this probability represents the fraction of non-faulty cache blocks, as follows;

$$F_{\text{avail}} = (1 - p_{\text{fault}})^{B} \tag{3}$$

4.1.2. SEC

Single error-correcting (SEC) codes were included due to their widespread use in existing devices.

4.1.2.1. Storage overhead

$$b = \lceil \log_2(B) \rceil \tag{4}$$

where b is the number of added ECC bits per cache block (basic storage).

4.1.2.2. Fault model. The model is modified to account for the increased block size (for check bit storage) and the ability to correct 0 or 1 bits.

$$F_{\text{avail}} = (1 - p_{\text{fault}})^{B+b} + \binom{B+b}{1} \times p_{\text{fault}} \times (1 - p_{\text{fault}})^{B+b-1} \tag{5}$$

4.1.3. BCH double error correction (DEC)

The Bose-Chaudhuri-Hocquenghem (BCH) error-correcting code was selected as a candidate DEC scheme. Alternatives such as Reed-Solomon and Golay codes are mentioned in [15]. BCH was chosen because of its low storage overhead. However, in

practise a less compute-intensive code can be used depending on the sensitivity of performance on L2 latency. We modeled a DEC BCH code storage overhead (with minimum distance $d_{\min} = 5$) based upon the equations in [16].

4.1.3.1. Storage overhead

$$b = 2 \times \lceil \log_2(B) \rceil \tag{6}$$

where b is the number of added ECC bits per cache block (basic storage).

4.1.3.2. Fault model. The probability of a faulty block is modified to account for the extra check bit storage and the ability to correct 2 bits.

$$F_{\text{avail}} = (1 - p_{\text{fault}})^{B+b} + \binom{B+b}{1} \times p_{\text{fault}} \times (1 - p_{\text{fault}})^{B+b-1} + \binom{B+b}{2} \times p_{\text{fault}}^2 \times (1 - p_{\text{fault}})^{B+b-2} \tag{7}$$

4.1.4. Log(B)

The log(B) scheme is an alternative single error-correcting scheme. A table stores the index of one faulty cell location per block, along with an additional bit to hold the correct state of that cell. This is equivalent to the distant repair scheme of [9] using one spare unit.

4.1.4.1. Storage overhead

$$b = \lceil \log_2(B) \rceil + 1 \tag{8}$$

These b bits per block are held in guaranteed correct storage.

4.1.4.2. Fault model. The fraction of available blocks is identical to that of SEC except that the additional bits are in guaranteed correct storage.

$$F_{\text{avail}} = (1 - p_{\text{fault}})^{B} + \binom{B}{1} \times p_{\text{fault}} \times (1 - p_{\text{fault}})^{B-1} \tag{9}$$

4.1.5. Triple modular redundancy (TMR)

Our triple modular redundancy implementation assigns faulty blocks to groups of three blocks with a majority vote on every bit (0 or 1 errors can be corrected per bit position). At most, 1/3 of logical bits can be recovered from the physical bits which are combined for a majority vote. In a hardware implementation, the bit comparison for the majority vote is performed at the final cache output stage. Therefore logic overhead will be small.

4.1.5.1. Storage overhead. No additional storage is allocated to identify which blocks are combined for TMR. For this typical case analysis, we assume that faulty blocks are combined with other arbitrarily located faulty blocks.

4.1.5.2. Fault model. We first consider each bit index as 3 bits which must have 0 or 1 faults to be corrected. This applies for all B bit indexes. However, because we only combine known faulty blocks after determining fault-free blocks, none of the three blocks are ever error-free and this probability (p_{good}) is subtracted from the main expression. The probability of a non-faulty block is $p_{\text{nfb}} = (1 - p_{\text{fault}})^B$.

$$p_{\text{good}} = \binom{3}{1} p_{\text{nfb}} \left[(1 - p_{\text{fault}}^2) + \binom{2}{1} (1 - p_{\text{fault}}) \times p_{\text{fault}} \right]^B - 3(1 - p_{\text{nfb}})(p_{\text{nfb}})^2 - 2(p_{\text{nfb}})^3 \tag{10}$$

$$\text{fraction\_repaired}_{\text{TMR}} = \frac{1}{3} \times \left(\left(\binom{3}{1} \times p_{\text{fault}} \times (1 - p_{\text{fault}})^2 + (1 - p_{\text{fault}}^3) \right)^B - p_{\text{good}} \right) \tag{11}$$

$$F_{\text{avail}} = p_{\text{nfb}} + (1 - p_{\text{nfb}}) \times \text{fraction\_repaired}_{\text{TMR}} \tag{12}$$

4.2. Proposed fault-tolerance scheme

4.2.1. Block grouping

For high error rates, we propose a new scheme (Fig. 4). Faulty physical blocks are grouped together (in groups of size G) to form a new, fully working logical block. In the rest of our analysis, we assume pairs (G = 2). Using larger groups is beneficial at extremely high error rates, but the analysis is beyond the scope of this paper. The concept is a similar to [8] except that knowledge of failure polarities is not required.

Compatible blocks have up to one faulty bit between them, at every corresponding bit index. This means that an additional "selector bit," which is known to be correct, can specify which bit contains a good value when reading data.

A "grouping table" is accessed as an additional step before a cache access, to identify the paired block.

- To read a grouped block, all blocks in the group are read. The selector bits then indicate which block in the pair contains good data, at each bit index. A single logical block is then returned to the processor.
- To write a grouped block, the same value is written to every component block.

4.2.1.1. Grouping table. This table is used to look-up the location of the other block in a group. If there is more freedom to combine faulty blocks that are compatible, more blocks can be recovered. They can be physically adjacent, in the same set or in any location inside the cache (depending on the desired complexity of block selection hardware). Each alternative has performance trade-offs, discussed later.

4.2.1.2. Selector bit table. Selection of the block which the data bit is read from is performed using a table of selector bits. These are stored in guaranteed correct cells, and can either cover every bit index in the block or a number of adjacent bits (e.g. two data bits per selector in Fig. 5). In this example, each selector bit indicates which block should be accessed for every pair of adjacent bits. Using fewer selector bits reduces the number of defects which can be tolerated but reduces storage overhead. For example, a single selector bit covering two adjacent data bits cannot handle the case where there is a fault in both blocks at that position. Later, we discuss off-chip caching of selector bits to reduce the die area of on-chip SRAM.

4.2.1.3. Table configuration. The tables are programmed during system start-up. Self-test routines determine whether cache cells are operating reliably at each voltage and frequency point, and the map is stored in main memory or on disk. When performance settings change, the cache is flushed and a new selector table loaded. These tables could also be hard-wired at manufacturing test.

4.2.1.4. Storage overhead. We call the first block to be accessed the "primary" block, and its paired compatible block the "secondary" block. For the grouping table, we first consider the most storage-intensive scenario where blocks are paired anywhere in the cache.

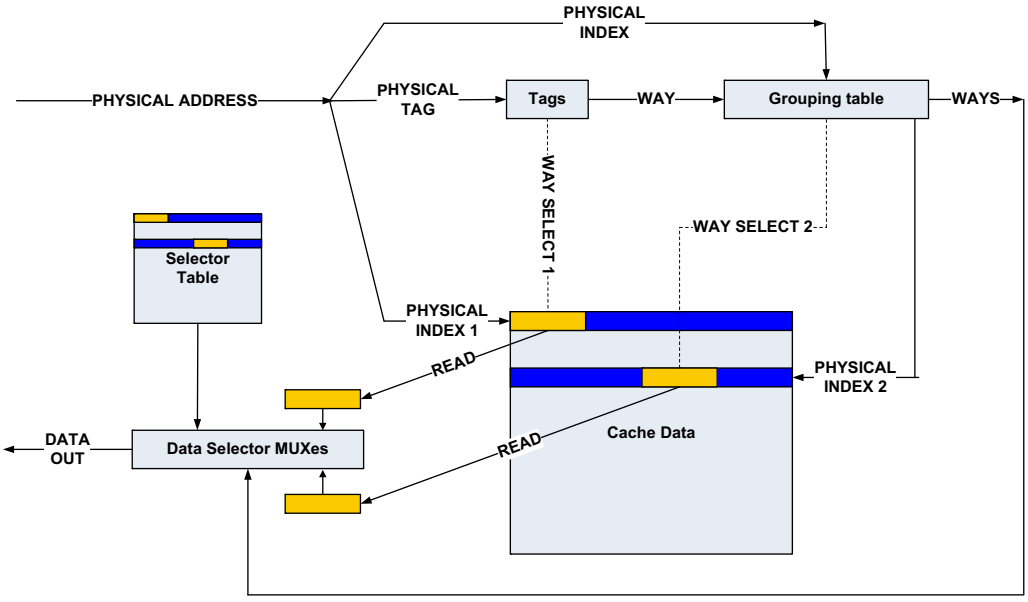

Fig. 4. The proposed block grouping scheme.

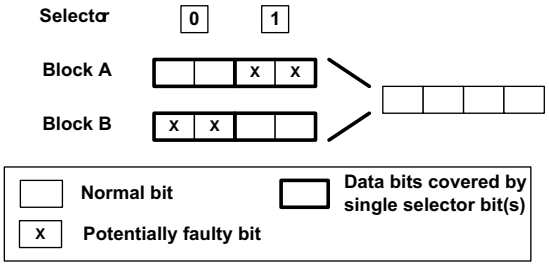

Fig. 5. Example of selector usage (2 bits/selector).

The grouping table has a number of entries equal to the number of blocks in the cache. Each entry stores the set and way index of a compatible block, to be looked up on a read access. The equation below assumes that there is an entry pointing to another block for every block position in the cache.

$$\text{group\_table\_size} = M \times N \times \log_2(M \times N) \tag{13}$$

As a lower-cost alternative, the storage requirement for pairs limited to the same set is given below. When implemented as an associative look-up, half of the blocks in a set have a pointer to another block in the same set.

$$\text{group\_table\_size} = \frac{N}{2} \times \log_2(N) \times M \tag{14}$$

Instead of using a grouping table, one could use the existing tag matching mechanism to simultaneously hit multiple blocks of the same group since their address tags are identical. It requires that the group resides in a single set so that address indexes are identical for each block. A banked cache design where ways are in different banks would allow fast parallel access to a group of blocks. A sequential access model is still feasible however, because an extra cycle to look-up a secondary block does not significantly impact performance for low-level caches (e.g. L2). Another, less effective zero-overhead alternative is to use a fixed grouping, for example, pairing together adjacent blocks.

The error-correcting ability of each variant is shown in Fig. 6. The results were derived from simulation, and pairs were formed using a greedy algorithm that allocates each faulty block with the next free compatible faulty block in sequential order. An optimal grouping will be even more effective.

Adjacent pairing is least effective and is not improved with associativity (Fig. 6a). It is clear that arbitrary pairing (Fig. 6b) is most effective at high error rates, although the per-set limitation (Fig. 6c) can be almost as effective. Increased associativity helps in this case by providing more pairing alternatives and can be seen in processors such as Niagara, with a 12-way L2 cache [17]. The final plot (Fig. 6d) represents blocks restricted to a set, and one selector bit is used for every two bits in the block. Fault tolerance is obviously reduced but there are now half as many selector bits.

The number of selector bits (used to choose one good bit from a group of G bits at each bit index) is the logarithm of the number of bits in the group. There are B selectors per cache block, and $(M \times N)/G$ logical blocks after grouping.

$$\text{selector\_bits} = \frac{M \times N}{G} \times (\log_2(G) \times B) \tag{15}$$

where G is the number of blocks in each group.

$$GC\_\text{bits} = MN + \text{group\_table\_size} + \text{selector\_bits (guaranteed correct storage)}.$$

4.2.1.5. Fault model. The parameter G can be varied, but all analysis in this paper uses block pairs ($G = 2$).

$$p_{\text{good}} = p_{\text{nfb}}^2 + \binom{2}{1} \times p_{\text{nfb}} \times (1 - p_{\text{nfb}}) \tag{16}$$

$$\text{fraction\_repaired}_{\text{GRP2}} = \frac{1}{2} \times ((1 - p_{\text{fault}}^2)^B - p_{\text{good}}) \tag{17}$$

$$F_{\text{avail}} = p_{\text{nfb}} + (1 - p_{\text{nfb}}) \times \text{fraction\_repaired}_{\text{GRP2}} \tag{18}$$

Each logical bit is formed from two bits. A set of two faulty physical blocks are compatible and can be recovered into a single logical block when;

- At most one of the grouped bits at each bit index are faulty, and
- This is true at every bit index in the block of B bits.

Note that we make an adjustment p_{good} to remove the impossible cases where any block contains no faults, as per TMR.

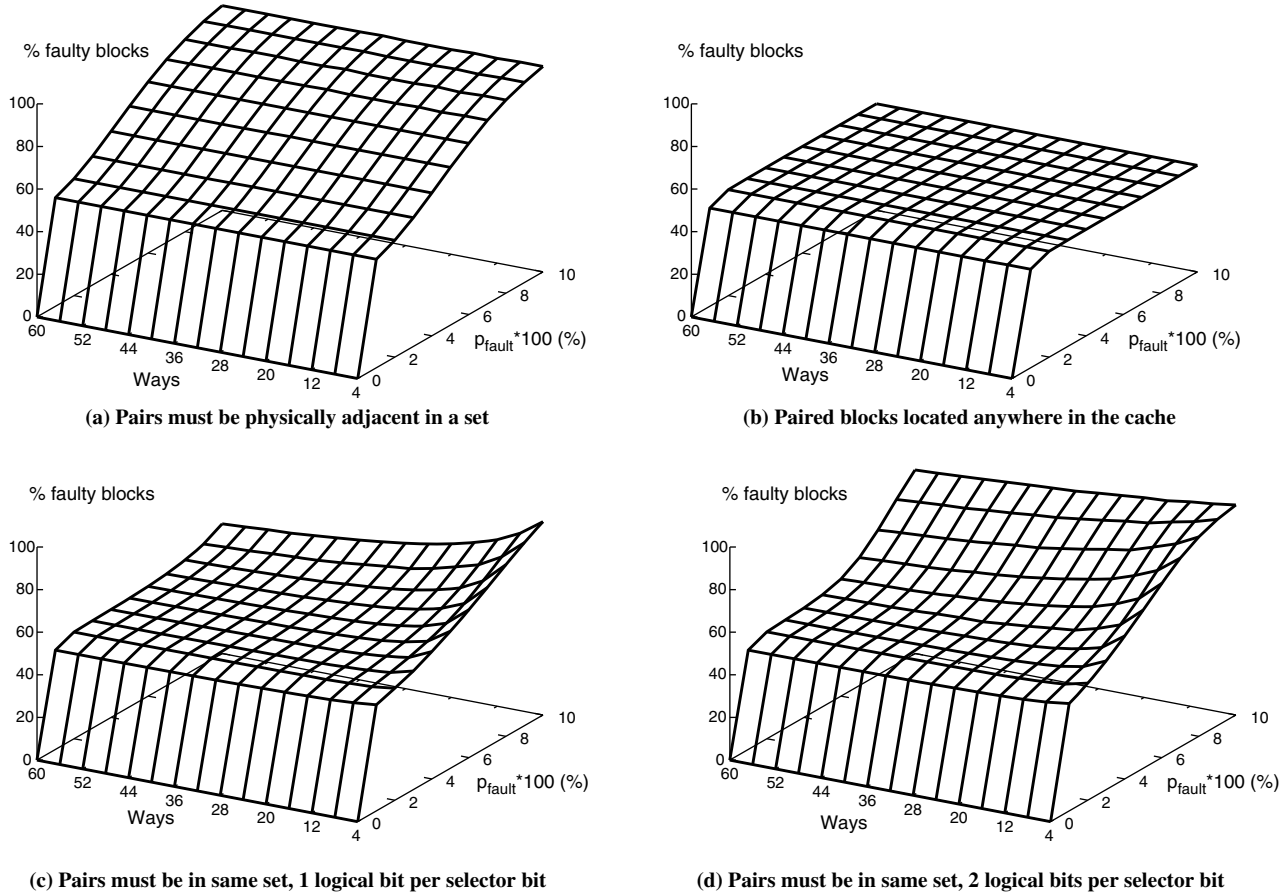

Fig. 6. Percentage of faulty blocks using block pairing ($G = 2$) for different group restrictions. Block size = 32 bytes.

4.2.2. Selector bit caching for block grouping

Compared with the other schemes, block grouping has strong fault tolerance characteristics but a potentially large storage overhead. By caching the working set of selector bits in on-chip SRAM and keeping less frequently used bits off-chip, area overheads can be reduced without significant performance impact. Although this appears only to move the reliability problem elsewhere, there are several benefits;

- Write accesses to the off-chip storage are infrequent (for example, at manufacturing test, or if BIST is re-run when there is a significant temperature change). Therefore the table could be stored in FLASH memory, for example. Off-chip DRAM is another possibility and in both cases there will be a reduction in total area and energy consumption. Although off-chip bus traffic when retrieving cached selector bits requires more energy than an on-chip access, this is an infrequent event.
- Moving the storage off-chip decouples most of the guaranteed correct storage from the CPU manufacturing process. This allows for a larger proportion of on-chip SRAM cells to be scaled down with the process technology. A different technology can be used off-chip.

We ran an initial study to see what performance and storage impact selector bit caching would have. The design in Fig. 4 was extended to have the working set of selector bit pages (stored in a parallel structure to the TLB) on-chip. On a TLB miss, we assume that the page table is accessed from main memory, so access latency to a small off-chip DRAM holding pages of selector bits is already accounted for. It is possible to use a small, standard cache rather than page-based method of finding the working set. This requires the use of tags, but makes more efficient use of storage space.

Using the same M5 configuration as before, we recorded TLB miss rates for varying numbers of TLB entries, then derived the performance hit for off-chip selector bit loading. For more realism and to support virtual memory, the simulator was run in full-system mode and benchmarks were run to completion under Linux. The results are given in Table 3. By keeping just the working set of cache pages in on-chip SRAM we have reduced the on-die storage overhead from 50% to less than 10%. Note that we used one selector bit per logical bit (see Fig. 6c). We opted to use 24 TLB entries because the performance data indicated much smaller

Table 3
Selector bit caching parameters and results

Parameters	Value
Cache size	1024 kB
Linux page size	8 kB
Logical bits per selector bit	1
Selector page size	4 kB
Full grouping table size/OBI table size	4 kB/2 kB
Set-restricted grouping table size	3 kB
On-chip SRAM selector bit storage	96 kB
Off-chip DRAM selector bit storage	512 kB
ITLB/DTLB entries	8/16
Selector bit table DRAM throughput	32 bytes/cycle
Total storage overhead of full group table	12.3%
Total storage overhead of set-limited group table	9.9%
Total storage overhead of tag-based grouping table	9.6%

slowdowns of 2% and 8% respectively for the Cholesky and Ocean-NonContig benchmarks. Due to the larger working set of Ocean-NonContig, increasing the number of data TLB entries does not significantly reduce miss rate.

5. Results

5.1. Performance and area under voltage scaling

We performed performance and power simulations using cell type C (1.5× the area of the smallest cell we considered in Fig. 2) and a rotating voltage schedule for the CPU. Every 20 ms (4×10^7 cycles @ 2 GHz), the voltage was changed to the next level in a in sequence from 0.7, 0.8, 0.9, 1.0 and 1.1 V. In reality, voltage changes will be less frequent in a DVS system. Because a large number of identical voltage changes occurred over the duration of each benchmark, a first-order approximation of energy consumption can be obtained by comparing overall execution times (assuming equivalent average power in each benchmark).

It is envisioned that each performance change will require the following additional steps;

- Invalidate dirty blocks and write back to the next level of the memory hierarchy (or high-speed local storage), if they are known to contain un-correctable faults.
- Enable the appropriate bad block map for the new performance level.
- Reinstate saved blocks, or allow to be fetched when next accessed, as usual.

Fig. 7a compares the L2 miss rate for each benchmark with the same voltage schedule using different fault-tolerance schemes. PAIR ADJ and PAIR ARB refer to 2-block grouping schemes where ADJ means only physically adjacent bad blocks are paired, and ARB means that blocks are paired arbitrarily throughout the cache using the same greedy algorithm used for Fig. 6. 5MOD and 7MOD implement 5- and 7-modular redundancy for each cell (representing 5 MB and 7 MB of physical storage), and NO FAULT refers to an ordinary error-free cache.

It is clear that execution time (and IPC) are relatively insensitive to L2 miss rates (see Fig. 7b). Note that arbitrary pairing generally performs better than 5-modular redundancy in most cases, without the large, fixed area overhead. Relative to OBI, PAIR ARB achieves an average 48% reduction in L2 miss rate and 13% reduction in execution time.

5.2. Performance under cell scaling

The E_{area} metric is shown in Fig. 8a. The most area efficient scheme is to use an OBI with TMR, as long as a cell is not scaled below size 1.4. This can be seen in the figure as the point of greatest E_{area} value. In fact, there is only a marginal improvement over using an OBI alone. Therefore, the area overhead of stronger error correction offsets the benefit of cell shrinking. The off-chip caching and arbitrary pairing variant was used for the grouping (GRP2) scheme. Therefore it initially has the worst E_{area} value due to the on-chip selector and grouping tables, but outperforms the others at smaller cell sizes due to superior fault-tolerance.

In Fig. 9 we examine E_{area} and as voltage is varied. This shows the same trends as Fig. 8. This means that cache performance will drop significantly depending on the fault-tolerance scheme and how far voltage is scaled in low-power (or low activity) modes. CPUs using DVS should dynamically select a fault-tolerance scheme with the highest F_{avail} at the operating voltage. For example, Fig. 9b indicates that DEC should be used down to 0.86 V and GRP2 below that.

5.3. Energy saving using block grouping at low V

An example of the energy benefits of block grouping is as follows. In an ultra-low voltage mode of 0.76 V (Fig. 9b) conventional SEC code has an F_{avail} of around 0.02 while GRP2 is approximately 0.45. This means that GRP2 provides much more L2 cache at that voltage, reducing miss rate and improving IPC. Considering the Barnes benchmark in Fig. 3b, IPC for 86% disabled blocks is at most 1.14 while for grouping the IPC is 1.76 (55% disabled blocks). This means that execution using grouping completes at least $(1 - 1.14/1.76) = 35\%$ sooner.

Energy savings are offset by the overhead of the selector bits and grouping table. If these are on-chip with one selector bit per data bit position, they will consume approximately 512 kB (or half the cache size). Despite this overhead, there will be a net energy saving. Using the power consumption of a 512 kB, 130 nm cache to represent the overhead of grouping [18], and that of a

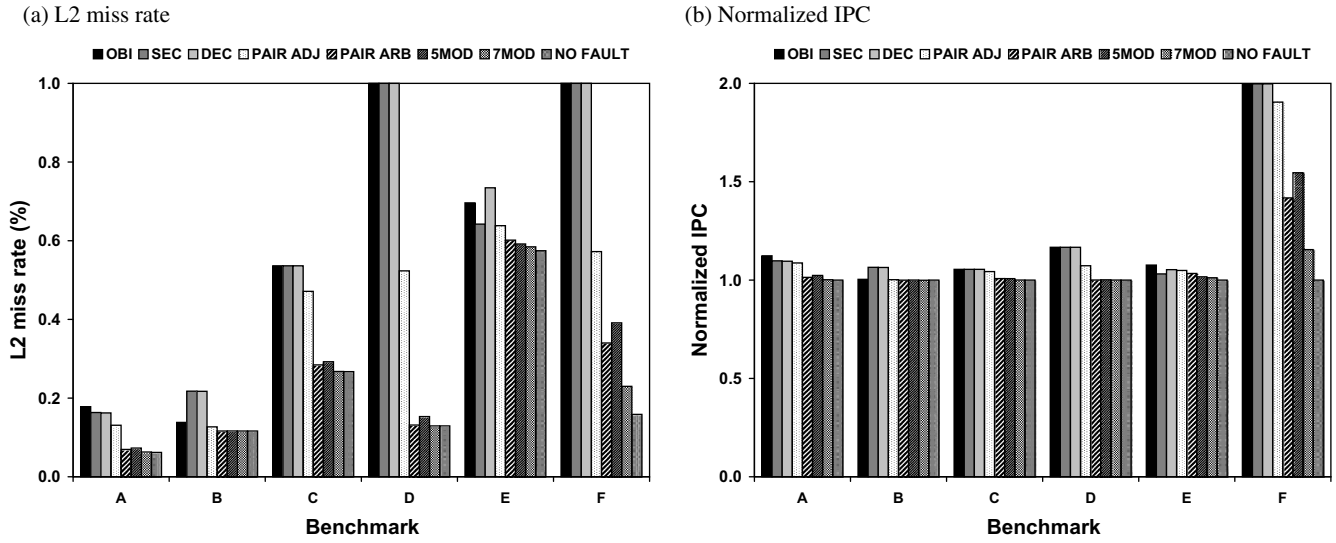

Fig. 7. L2 miss rate in (a) and normalized IPC in (b) for each technique. A-Cholesky, B-FFT, C-LUContig, D-Radix, E-OceanContig, and F-OceanNoncontig, respectively.

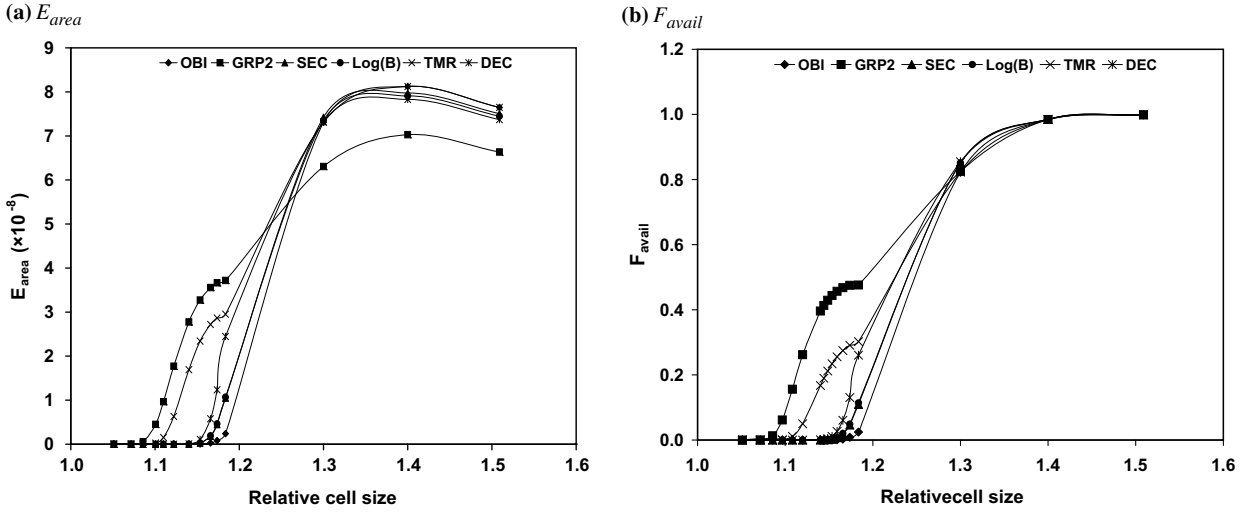

Fig. 8. E_{area} in (a) and F_{avail} in (b) with device scaling at 1.1 V (64-byte block size). Cell size is relative to the smallest considered size from Fig. 2.

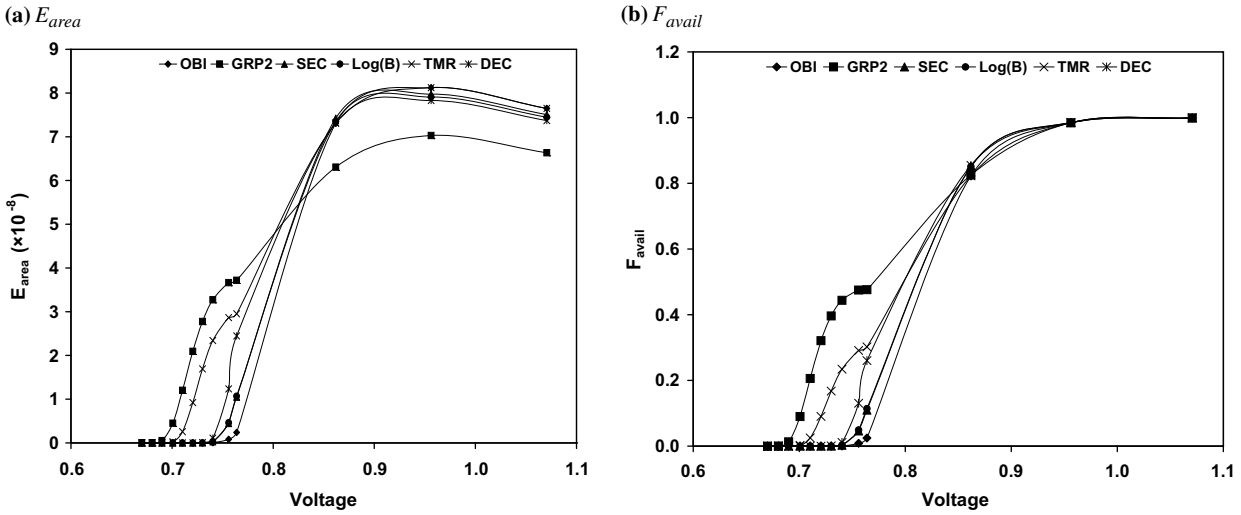

Fig. 9. E_{area} in (a) and F_{avail} in (b) with voltage scaling (64-byte block size, cell size C from Fig. 2).

Pentium M running at around the same frequency [19] we estimate power consumption for the SEC and grouping schemes (Table 4). We use the thermal design power of the processor which implies the CPU is 100% utilised by the workload. In addition, we assume that the processor uses SEC fault-tolerance. Even though we theoretically scale voltage down to 0.76 V, both the processor and selector table have their voltages scaled by the same factor, so the power consumption ratio between the two is approximately the same.

Table 4
Block grouping energy saving example

Parameters	Value
Pentium M thermal design power	24.5 W
Selector table overhead power	2.6 W
Speedup of grouping relative to SEC	35%
Power with SEC	24.5 W
Power with grouping table	24.5 + 2.6 = 27.1 W
Average power with grouping	27.1∗(1 − 0.35) = 17.6 W
Energy savings with grouping	28%

6. Conclusions

The analysis in the first part of this paper compared several ways of maximizing the number of usable cache lines in the presence of faults. These faults can be a combination of permanent manufacturing faults as well as ineffective operation at low voltages. Increasing the size of an SRAM cell increases this reliability at the cost of extra area.

Next, we proposed a novel fault-tolerance scheme that takes advantage of a region of larger or higher voltage SRAM cells to attain high reliability. The scheme works by grouping two or more cache lines divided into smaller regions. Selector bits in the high reliability memory cells are used to specify where the faulty bits are in the grouped blocks.

It was determined that the previously published OBI-based fault tolerance is the most area efficient scheme for fault-tolerance at a single voltage. However, as voltage is scaled down, maximum performance and energy savings are obtained by switching from DEC to our GRP2 scheme.

Instead of taking advantage of future scaling to reduce SRAM cell size, scaling should not go beyond a certain point. This is be-

cause the area overhead of trying to protect the smaller, but much less reliable cells is greater than that of not scaling the cells at all. However, error correction still has a place in cache design for low voltage performance and soft error tolerance. The strength of this error correction will be a function of expected soft error rates and how aggressively DVS is applied.

References

[1] A. Agarwal et al., A process-tolerant cache architecture for improved yield in nanoscale technologies, IEEE Trans. VLSI Syst. 13 (1) (2005) 27–38.

[2] N. Kim et al., Leakage current – Moore's law meets static power, IEEE Comput. 36 (12) (2003) 68–75.

[3] M. Agostinelli, Erratic fluctuations of SRAM cache Vmin at the 90 nm process technology node, in: IEEE International Electron Devices Meeting (IEDM), December 2005, pp. 655–658.

[4] A. Pour, M. Hill, Performance implications of tolerating cache faults, IEEE Trans. Comput. 42 (3) (1993) 257–267.

[5] S. Mukhopadhyay et al., Modeling of failure probability and statistical design of SRAM array for yield enhancement in nanoscaled CMOS, IEEE Trans. CAD 24 (12) (2005) 1859–1880.

[6] D. Bossen, J. Tendler, K. Reick, Power4 system design for high reliability, IEEE Micro. 22 (2) (2002) 16–24.

[7] A. KleinOsowski, D. Lilja, The NanoBox project: exploring fabrics of self-correcting logic blocks for high defect rate molecular device technologies, in: IEEE Computer Society Annual Symposium on VLSI, February 2004, pp. 19–24.

[8] M. Nicolaidis, A memory built-in self-repair for high defect densities based on error polarities, in: IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems, November 2003, pp. 459–466.

[9] M. Nicolaidis, N. Achouri, S. Boutobza, Dynamic data-bit memory built-in self-repair, in: International Conference on Computer Aided Design (ICCAD), November 2003, pp. 588–594.

[10] D. Burnett et al. Implementations of fundamental threshold voltage variations for high-density SRAM and logic circuits, in: IEEE International Symposium on VLSI Technology, June 1994, pp. 15–16.

[11] S. Mukhopadhyay et al., Modeling and estimation of failure probability due to parameter variation in nano-scale SRAMs for yield enhancement, in: IEEE International Symposium on VLSI Circuit, June 2004, pp. 64–67.

[12] N. Binkert, E. Hallnor, S. Reinhardt, Network-oriented full-system simulation using M5, Workshop on Computer Architecture Evaluation using Commercial Workloads, 2003, pp. 36–43.

[13] S. Woo et al. The SPLASH-2 programs: characterization and methodological considerations, in: International Symposium on Computer Architecture (ISCA), June 1995, pp. 24–36.

[14] D. Lamet, J. Frenzel, Defect-tolerant cache memory design, in: IEEE VLSI Test Symposium, April 1993, pp. 159–163.

[15] P. Mazumder, Design of a fault-tolerant three-dimensional dynamic random-access memory with on-chip error-correcting circuit, IEEE Trans. Comput 42 (12) (1993) 1453–1468.

[16] L. Joiner, J. Komo, Decoding binary BCH codes, IEEE SoutheastCon (March) (1995) 67–73.

[17] P. Kongetira, K. Aingaran, K. Olukotun, Niagara: a 32-way multithreaded Sparc processor, IEEE Micro. 25 (2) (2005) 21–29.

[18] J. Shin et al., Design and implementation of an embedded 512-kB level-2 cache subsystem, IEEE J. Solid-State Circuits (JSSC) (Sept.) (2005) 1815–1820.

[19] Intel Pentium M Power Data, 1.6 GHz, Technology, 1 MB L2 Cache <http://www.intel.com/design/intarch/pentiumm/pentiumm.htm/>.