Circuit Design of a Novel Adaptable and Reliable L1 Data Cache

Azam Seyedi[†‡], Gulay Yalcin[†‡], Osman S. Unsal[†], Adrian Cristal[†°]
[†]BSC-Microsoft Research Centre [‡]Universitat Politècnica de Catalunya
[°]IIIA - Artificial Intelligence Research Institute CSIC - Spanish National Research Council
{azam.seyedi, gulay.yalcin, osman.unsal, adrian.cristal }@bsc.es

ABSTRACT
This paper proposes a novel adaptable and reliable L1 data cache design (Adapcache) with the unique capability of automatically adapting itself for different supply voltage levels and providing the highest capacity. Depending on the supply voltage level, Adapcache defines three operating modes: In high supply voltages, Adapcache provides reliability through single-bit parity. In middle range of supply voltages, Adapcache writes data to two separate cache-lines simultaneously in order to use one line for error recovery when the other line is faulty. In near threshold supply voltages, Adapcache writes data to three separate cache-lines simultaneously in order to provide the correct data based on bitwise majority voter.

Categories and Subject Descriptors
B.3.2 [**Memory Structure**]: Design Styles

General Terms
Performance, Design, Verification, Reliability

Keywords
Cache Design, Low Power Design, Fault Tolerance

1. INTRODUCTION
A very effective approach in reducing the energy consumption is to reduce the supply voltage close to the transistor's threshold. However, the energy reduction in the low-power mode comes with a drastic increase in the number of memory cell failures especially in large memory structures such as on-chip SRAM memories. This motivates us to design a cache which is resilient to large number of cell failures and operates at lower supply voltages.

In this paper, we propose a novel highly reliable L1 data cache for three distinct supply voltage ranges. Our proposed cache, Adapcache, automatically adapts itself for different supply voltages in order to provide the most possible capacity with minimum energy consumption and access time. Adapcache operates in three states: (1) Single Copy State (SCS), (2) Double Copy State (DCS) and (3) Triple Copy State (TCS). When the supply voltage is relatively high, Adapcache operates in SCS to satisfy high performance execution. In this state, Adapcache provides reliability only based on single-bit interleaved parity.

GLSVLSI'13, May 2–3, 2013, Paris, France.

DCS is utilized in medium low supply voltage ranges in which error rate starts to increase. Adapcache writes data to two selected cache-lines simultaneously. At every read instance, Adapcache compares both those values through XOR circuits to check their equality. This approach is similar to Dual Modular Redundancy technique (DMR) but unlike DMR, the utilization of parity bits enables Adapcache to decide the correct value upon an error. When the supply voltage is at near-threshold and the error rate increases drastically, Adapcache operates in TCS. Adapcache writes data to three selected cache-lines simultaneously to provide very high reliability. In read, the correct value is decided through bit-wise majority voters similar to Triple Modular Redundancy technique (TMR) but the error correction capability is improved by parity bits participation in majority voting.

2. ADAPCACHE CICUIT DESIGN
Our sub-bank design is based on the structure that authors proposed in [1]. Figure 1 shows the block diagram of each sub-array structure. Whenever the supply voltage is high, Adapcache operates in SCS and only one cache-line is activated at each access time. For writing the selected cache-line, signal IEU1 is high and activates input buffers and data can transfer to the selected cache-line via Bus4 and Bus1; and similarly for reading the selected line, signal OEU1 is high and output buffers are active and data is transferred from Bus1 to Bus3. The error protection is based on bit-parity calculation in order to achieve higher speed but less accuracy.

In medium supply voltage range Adapcache operates in DCS where at each cache-access two cache-lines are activated simultaneously. At writing time data is written to two selected lines at the same time. Parity calculator circuits generate parity bits and write them in parity bit cells as well. For reading the two selected lines, Bus1 is divided into two parts and output buffers transfer two selected data each from separate sub-array slice groups (sub-array slices 0 to 3 and sub-array slices 4 to 7) to the XOR circuit to check their identically. If there is at least one bit flip, signal EN10 activates two parity calculator circuits to calculate parity bits of selected lines and compare with the original parity bits of each selected-lines. Whenever one of comparator shows equality, the related output buffer transfers its data to Bus3. In TCS, for writing the selected cache-lines signal IEU1 is high and data is transferred and written in three selected lines simultaneously. At the reading time, Bus1 is divided into three parts and three selected lines each from separate sub-array slice group (sub-array slices 0, 1, 2, and sub-array slices 3, 4, 5, and sub-array slices 6, 7 and extra slice) are transferred to a majority voter. The majority voter output for cache-lines is DataM and for their parity-bits is ParityM. The parity bits of DataM are calculated and compared with ParityM. If there are any differences, the parity bits of selected lines should be calculated and compared with their original parity bits. Whenever one of comparator shows equality, the related output buffer transfers its data to the output. We designed a new decoder circuit to activate

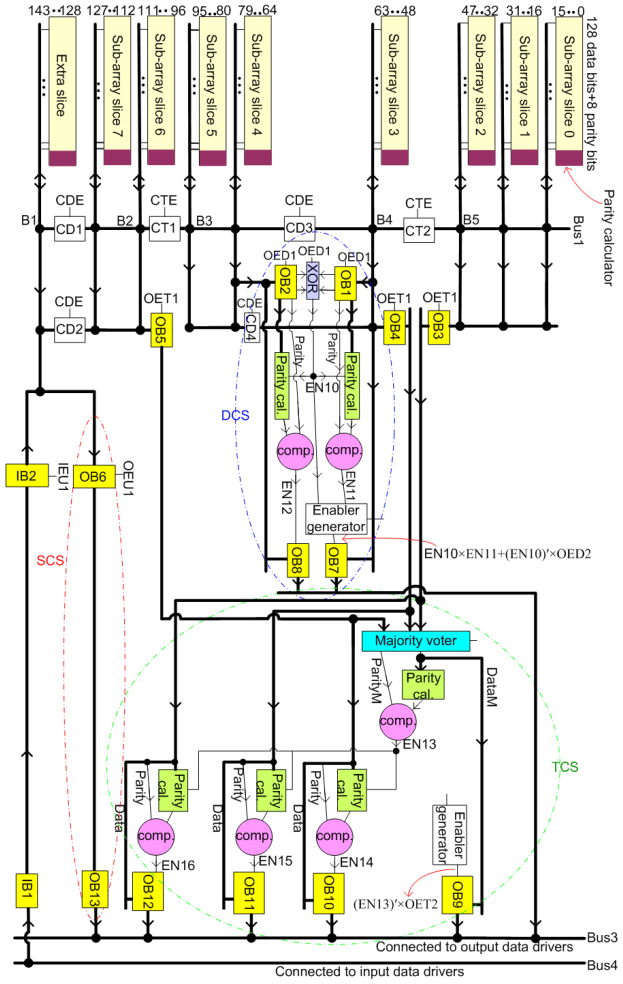

Figure 1: Sub-array block diagram.

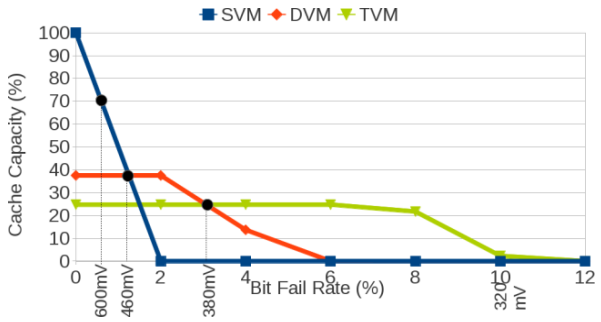

Figure 2: Comparison between SCS, DCS and TCS.

one, two or three word-line addresses at each access time and described its detail in [2].

3. EVALUATION

We evaluate useful cache capacity, access time and energy consumption of Adapcache and compare to DMR and TMR caches and also to Parichute [3]. We inject random persistent faults [4] according to bit failure rate in [0%-12%] by repeating each experiment 100 times. We calculate the useful cache capacity as the undisabled portion of the cache as can be seen in Figure 2. Bit failure rate for persistent failures in the given Vcc is examined by Miller et al [3] which we reference in our work. Figure 2 presents the Vcc intervals that SCS, DCS and TCS can be used

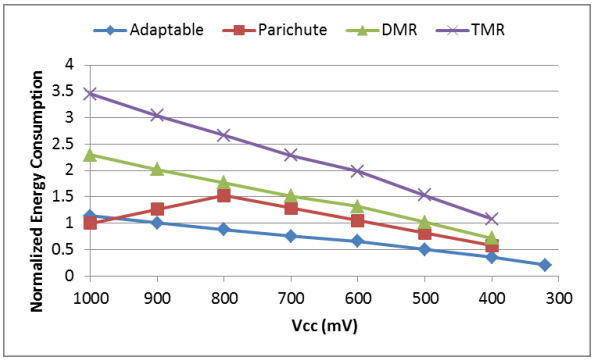

Figure 3: Normalized energy consumption of caches.

efficiently according to persistent bit failure rate. The intersections of the three lines indicate the optimum useful points of SCS, DCS and TCS according to Vcc. Thus, we decide using SCS when Vcc=[1V-600mV], DCS when Vcc=]600mV-400mV] and TCS when Vcc=]400mV-320mV].

We simulate a 64-KB L1 data Adapcache with Hspice 2003.03 using HP 45-nm Predictive Technology Model at 2GHz processor frequency for (1V-0.6V) supply voltages, at 900MHz for (0.6V-0.4V) supply voltages, and at 400MHz for (0.4V-0.32V) supply voltages. We calculate the energy consumption and latency of the Adapcache versus the Vcc levels and compare it with three 64-KB caches which deploy TMR, DMR and Parichute techniques. Parichute, can reduce baseline energy consumption by 64%. In comparison, Adapcache can reduce baseline energy by 74% with a 7X less access latency than Parichute. Figure 3 depicts the normalize energy consumption of each cache structure (We normalized based on a single parity protected cache in Vcc=1V).

We calculate the latency of all structures for a single cache-line. In Adapcache, if a line is faulty, it takes 3 clock cycles to read and correct it. For DMR case it takes 5 clock cycles to read selected lines and detect the errors. Similarly, it may take 7 clock cycles for TMR structure to read selected cache-lines and check and correct the errors. Finally for Parichute, it may take up to 20 cycles if the bit failure rate is very high. This latency is almost two times the L2 cache latency which may not be acceptable.

4. ACKNOWLEDGMENTS

This work was partially supported by the cooperation agreement between the Barcelona Supercomputing Center and Microsoft Research, by the Ministry of Science and Technology of Spain and the European Union (FEDER funds) under contracts TIN2007-60625 and TIN2008- 02055-E, and by the European Network of Excellence on High-Performance Embedded Architecture and Compilation (HiPEAC), and by the European Community´s Seventh Framework Programme [FP7/2007-2013] under the ParaDIME Project (www.paradime-project.eu), grant agreement no. 318693.

5. REFERENCES

[1] A. Seyedi, et al. Circuit design of a dual-versioning L1 data cache for optimistic concurrency. In GLSVLSI 2011

[2] A. Seyedi, et al. Circuit Design of a Novel Adaptable and Reliable L1 Data Cache, Technical Report UPC-DAC-RR-CAP-2013-3 , UPC, 2013

[3] T. N. Miller. Parichute: Generalized Turbocode-Based Error Correction for Near-Threshold Caches. In MICRO 2010.

[4] Z. Chishti et al. Improving Cache Lifetime Reliability at Ultra-Low Voltages. In MICRO 2009.