



Reconfiguring Cache Associativity : Adaptive Cache Design for Wide-Range Reliable Low-Voltage Operation Using 7T/14T SRAM

Jung, Jinwook
Nakata, Yohei
Okumura, Shunsuke
Kawaguchi, Hiroshi
Yoshimoto, Masahiko

(Citation)

IEICE Transactions on Electronics, 96(4):528-537

(Issue Date)

2013-04-01

(Resource Type)

journal article

(Version)

Version of Record

(Rights)

copyright©2013 IEICE

(URL)

<https://hdl.handle.net/20.500.14094/90002980>



Reconfiguring Cache Associativity: Adaptive Cache Design for Wide-Range Reliable Low-Voltage Operation Using 7T/14T SRAM*

Jinwook JUNG^{†a)}, Nonmember, Yohei NAKATA[†], Shunsuke OKUMURA[†], Student Members, Hiroshi KAWAGUCHI[†], and Masahiko YOSHIMOTO^{†,††}, Members

SUMMARY This paper presents an adaptive cache architecture for wide-range reliable low-voltage operations. The proposed associativity-reconfigurable cache consists of pairs of cache ways so that it can exploit the recovery feature of the novel 7T/14T SRAM cell. Each pair has two operating modes that can be selected based upon the required voltage level of current operating conditions: normal mode for high performance and dependable mode for reliable low-voltage operations. We can obtain reliable low-voltage operations by application of the dependable mode to weaker pairs that cannot operate reliably at low voltages. Meanwhile leaving stronger pairs in the normal mode, we can minimize performance losses. Our chip measurement results show that the proposed cache can trade off its associativity with the minimum operating voltage. Moreover, it can decrease the minimum operating voltage by 140 mV achieving 67.48% and 26.70% reduction of the power dissipation and energy per instruction. Processor simulation results show that designing the on-chip caches using the proposed scheme results in 2.95% maximum IPC losses, but it can be chosen various performance levels. Area estimation results show that the proposed cache adds area overhead of 1.61% and 5.49% in 32-KB and 256-KB caches, respectively.

key words: low-voltage adaptive cache design, reconfiguring associativity, dynamic voltage frequency scaling, 7T/14T SRAM

1. Introduction

Feature size in transistors continues to shrink along with the advance of process technology, achieving higher integration density and lower per-transistor cost. The transistor count of a current state-of-the-art microprocessor reaches several billion. In such a situation, power efficiency has become one of the most important design decisions even for high-performance system on chip (SoC) designs. In particular, large on-chip caches account for a great fraction of total processor power dissipation [2], [3].

Reducing the power supply voltage has been an extremely efficient technique to reduce the power dissipation. It is for this reason that the power management features based on Dynamic Voltage Frequency Scaling (DVFS), which reduce the supply voltage to various levels depending on current workload requirements, are extensively used [4]–[7]. However, ongoing technology scaling increases the variations in different device parameters and creates a con-

siderable spread in a transistor threshold voltage, mainly because of random dopant fluctuation (RDF), which has deviation that is inversely proportional to the square of a channel area [8]. Such variations strongly affect functionality and reliability at low voltages [9], [10].

This situation yields severe problems, particularly in SRAM, because minimum-sized transistors are used in its design. Device mismatching between transistors in SRAM caused by the process variations make the memory cell unreliable, which results in increased cell failures (read failure, write failure, and access time failure) [11]. In addition, the minimum operating voltage (V_{\min}) tends to increase according to technology scaling [2]: a low V_{\min} decreases read/write margins in SRAM and reliability deterioration arises. To make matters worse, SRAMs occupy a substantial fraction of the total die area and transistor count in processors [16]. Consequently, a large SRAM block such as the last level cache (LLC) determines the V_{\min} of the whole processor and restricts voltage scaling. This in turn decreases the range of supply voltage and precludes the exploitation of aggressive DVFS.

Several studies have been conducted for the reliable low-voltage cache operations [12]–[16]. Shirvani et al. proposed the PADed cache in [12], which uses a programmable address decoder and programs it to substitute erratic cache lines with non-faulty cache lines. Agarwal et al. presented a similar technique in [13]. Their architecture uses an online built-in-self-test (BIST) to specify the location of faults at low voltages and then programs column multiplexers not to select the faulty cache lines. However, these techniques are unable to deal with a great deal of erratic cache lines at low voltages.

Makhzan et al. presented the low voltage cache architecture which adds two spare SRAM blocks, one for locating the defective cache line and one for redundant cache line to substitute the defective cache line [14]. Nevertheless, their technique would need very large spare blocks to cover the excessively increased faulty cache lines at low voltages. Moreover, a large proportion of the redundant cache lines in the spare blocks are inefficient in low failure rate situation.

Kim et al. proposed a scalable multi-bit error protecting scheme which uses spare SRAM leverages two-dimensional error coding [15]. However, the capability of their architecture to improve cache reliability is strongly restricted the number and locations of erratic bits, which are randomly distributed in the entire cache because of RDF. This tech-

Manuscript received August 7, 2012.

Manuscript revised November 3, 2012.

[†]The authors are with Kobe University, Kobe-shi, 657-8501 Japan.

^{††}The author is with JST CREST, Tokyo, 102-0076 Japan.

*This paper is the extended version of the original paper presented at the 18th IEEE ICECS [1].

a) E-mail: jung@cs28.cs.kobe-u.ac.jp

DOI: 10.1587/transele.E96.C.528

nique also necessitates large spare SRAM blocks and additional cache access cycles.

Wilkerson et al. suggested two low voltage cache schemes, which are named as the Word-disable (WDIS) and the Bit-fix (BFI), in [16]. WDIS scheme produces a non-faulty cache line combining two consecutive cache lines in the low voltage operations, while BFI scheme leverages a quarter of cache lines to recover the erratic bits in other cache lines. However, these schemes incur cycle penalties and need dedicated circuitries which consume large amount of energy. In addition, because both the WDIS and BFI have only two operating modes, it cannot fulfill the various requirements of the power and performance level in a variety of workloads.

In this paper, we propose the associativity-reconfigurable cache using the 7-Transistor/14-Transistor SRAM (7T/14T SRAM) which has the operating margin enhancing feature for unreliable SRAM cells at low voltages. Our proposed cache reduces the V_{\min} of the entire cache by trading-off its associativity (the number of cache ways) and capacity. The proposed N -way set-associative cache can be reconfigured depending on the required degree of low-voltage reliability and processor performance. Its associativity can be chosen between $N/2$ and N achieving the desired V_{\min} for end performance. This reconfigurability in the cache associativity therefore provides reliable low-voltage operations for required performance and voltage levels, delivering the reduced power dissipation and improved energy efficiency. Furthermore, the proposed cache makes it possible to leverage optimal cache configuration in DVFS.

The remainder of the paper is organized as follows: Sect. 2 provides basic information about the 7T/14T SRAM cell, of which recovery features are utilized in our proposed cache. Sections 3 and 4 present details of the proposed associativity-reconfigurable cache. Section 5 presents our experimental methodologies and evaluation results of the proposed cache. Finally we give some concluding remarks in Sect. 6.

2. 7T/14T SRAM

We have proposed the 7T/14T SRAM cell [17]. Figure 1 shows a schematic view of 7T/14T SRAM: two pMOS transistors (M20 and M21) are connected to internal nodes (N00 and N10, N01 and N11) in a pair of the conventional 6-Transistor (6T) memory cells. The structure of 7T/14T SRAM thereby achieves an additional operating mode which is designated as the dependable mode along with the typical operating mode, the normal mode. In the dependable mode, 7T/14T SRAM features margin enhancements by combining two memory cells, especially in a low-voltage region. Two modes in 7T/14T SRAM can be summarized as shown in Table 1.

In the normal mode, a one-bit data is stored in one memory cell, which is more area-efficient. In the dependable read mode, only one wordline is asserted to gain a large β ratio (a ratio of two driver transistors' total size to one

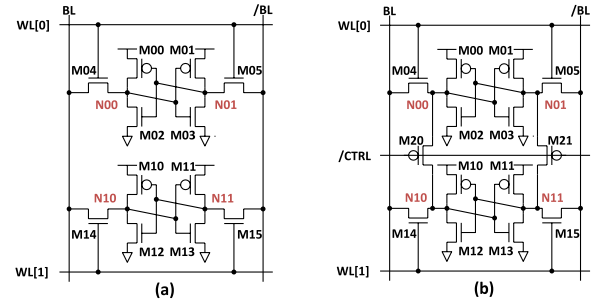


Fig. 1 Schematic of SRAM cell pairs: (a) conventional 6T SRAM and (b) 7T/14T SRAM.

Table 1 Two Modes in 7T/14T SRAM.

Mode	# of memory cells comparing 1 bit	# of WL drivers	CTRL
Normal	1 (7 Transistors/bit)	1	Off ("H")
Dependable (write)	2 (14 Transistors/bit)	2	On ("L")
Dependable (read)	2 (14 Transistors/bit)	1	On ("L")

access transistor size). An SRAM cell with no static noise margin [18] is recovered by the other SRAM cell through the two connecting pMOS transistors. In the dependable write mode, a data is written into a pair of memory cells by asserting both wordlines, which averages and mitigates the write margin degradation. Not only these margin enhancing features, 7T/14T SRAM in the dependable mode also shows lower bit error rates (BER), which is a metric of the SRAM failure rate, than traditional approaches to improve SRAM reliability such as error correcting codes (ECC) or hardware redundancy based techniques [17]. In addition, the dependable mode has better soft-error tolerance because its internal node has more capacitance [19].

The normal mode and the dependable mode of 7T/14T SRAM can be switched according to the operating voltage, power limit, and required voltage margin in the current operating condition. If 7T/14T SRAM has sufficient operating margins or higher performance is needed, then its recovery feature can be inactivated by negating two additional pMOS transistors. If there are needs for low-voltage reliability to reduce power dissipation or higher dependability, we can change its operating mode to the dependable mode. This mode transition of 7T/14T SRAM can be conducted dynamically by appropriate control of CTRL line in Fig. 1 without rebooting the entire system. Therefore the operating mode to suit the requirements of application can be chosen according to current operating conditions, such as operating voltage and power dissipation.

As mentioned above, a one-bit data is stored in two memory cells in the dependable mode. Although the area efficiency is decreased, the quality of the information is improved from that of the normal mode. We designate this concept as 'quality of a bit (QoB)', in which the operating voltage, power, and BER are controlled as attributes of one-

bit information [20].

3. Associativity-Reconfigurable Cache

In this section, we describe the proposed associativity-reconfigurable cache scheme. Many modern microprocessors have DVFS features and multiple operating modes such as high-performance/high-voltage modes and low-power/low-voltage modes [4]–[7]. These multiple operating modes provide optimal performance and power consumption with respect to current operating workloads. However, V_{\min} of large caches in microprocessors limit the use of wide range of voltage levels. This limitation makes it difficult to realize optimal DVFS control.

In the proposed associativity-reconfigurable cache, consecutive odd-even cache ways pair up by exploiting the structure of 7T/14T SRAM cell. This way pair organization enables that the proposed cache can decrease its operating voltage to various voltage levels resulting in the reduction of power dissipation by changing its current cache configuration. Switching operating modes in 7T/14T SRAM is conducted with respect to these way pairs.

Figure 2 shows the voltage reduction mechanism of the proposed associativity-reconfigurable cache. The V_{\min} of the entire cache is determined by the way pair which has the highest value of V_{\min} [8]. In the proposed cache, Cache V_{\min} reduction is achieved by application of the dependable mode of 7T/14T SRAM to the way pair with the highest value of V_{\min} . If all the way pairs enter the dependable mode, the proposed cache can fully exploit the margin enhancement feature of 7T/14T SRAM, resulting in the lowest value of the V_{\min} it can achieve.

To ascertain the V_{\min} of each way pair, we use a boot-time low-voltage memory test [21]. The memory test is performed with respect to the way pair of the proposed cache decreasing the testing voltage in the normal mode and the dependable mode, respectively. We assume that the memory test is performed using 10 mV voltage steps [22]. The testing voltage on which the defective cache lines are detected at first is the V_{\min} of way pair. Once defective cache lines are detected, then the current testing voltage and the way pair number are stored in cache controller. Next, testing is executed in the dependable mode and V_{\min} of each way pair in the dependable mode is also stored. Based on these




Fig. 2 Conceptual view of the V_{\min} reduction by the proposed associativity-reconfigurable cache.

stored value of each way pair's V_{\min} , the cache controller selects way pairs to which the dependable mode is applied according to the desired voltage and performance level.

3.1 Organization of the Proposed Cache Using 7T/14T SRAM

The N -way set associative cache using the proposed scheme can change its associativity between $N/2$ and N . Once the operating system and DVFS controllers choose the power mode based on the current operating conditions and thereby the operating voltage for the current workload, the proposed cache reconfigures its associativity to accommodate the selected voltage level. Figures 3 and 4 describe the organization of the proposed cache. Adjacent two odd-even cache ways compose a pair of odd-even ways: a way pair. These way pairs enable exploitation of the dependent feature of 7T/14T SRAM for adaptive cache design with a wide range of operating voltages. If no need exists to leverage the enhanced reliability of the dependable mode (for instance, when the operating voltage margin is sufficient to operate properly or application software requires large cache capacity for high performance), the two ways in a pair operate separately in the normal mode, as shown in Fig. 3. Figure 4 illustrates the case of the low-voltage operation mode. Odd-even ways in a way pair are logically bound together and constitute the dependable cache way that features enhanced operating margin, thereby enabling reliable low-voltage operations.

Figure 5 illustrates detailed views of the dependable cache ways interleaving odd-even cache lines. A one-bit




Fig. 3 An organization of the associativity-reconfigurable cache. All odd-even way pairs operate in the normal mode using 7T SRAM.




Fig. 4 An organization of the associativity-reconfigurable cache. Pair 0 operates in the dependable mode making one dependable way, whereas pair 1 operates in the normal mode; the entire cache operates as a 3-way set-associative cache in this example.

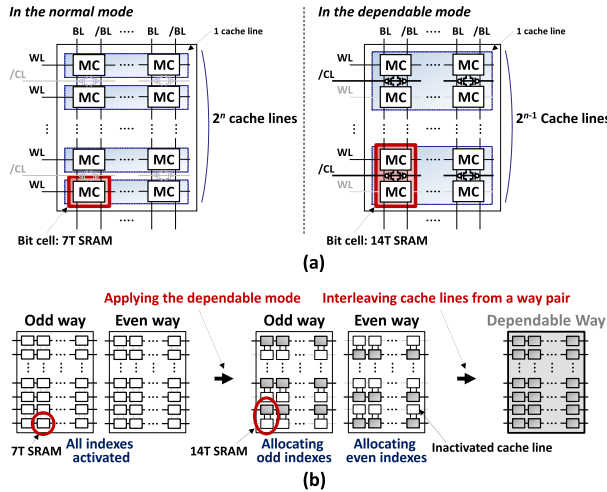


Fig. 5 Composition of dependable cache way in the dependable mode: (a) physical view and (b) logical allocation of cache lines in a way pair.




Fig. 6 Implementation of cache decoders.

data in the dependable way is made up of a pair of memory cells. Therefore, the capacity is halved and the associativity is decreased by one in a way pair, but improved reliability is obtainable. If the operation margins are sufficient, alternatively if high performance is a critical issue, then each way pair might be detached and be operated separately as described above.

The configuration of the proposed cache can be determined arbitrarily by activation or inactivation of corresponding control lines (CTRLs in Fig. 1). Which way pairs and how many way pairs to switch their mode to dependable mode can be parameterized. Appropriate associativity can be chosen by applying the dependable mode to some pairs of ways selectively. Therefore the desired voltage mode is obtainable.

In order to implement the proposed cache, it is necessary to use extended decoders as shown in Fig. 6; these are one n -to- 2^n decoder and one $n-1$ -to- 2^{n-1} decoder (where n is a bit width of the cache index). In the normal mode, the upper n -to- 2^n decoder is activated. It drives each cache way independently. On the other hand, in the dependable mode, the $n-1$ -to- 2^{n-1} decoder is asserted in the dependable mode; a pair of odd-even ways comprises one dependable way. Operating as one dependable way, it is necessary to decide either the odd way or the even way in the way pair is chosen because the capacity is halved. The least significant bit (LSB) in the cache index is used to do so. If the LSB is 0, then the $n-1$ -to- 2^{n-1} decoder's output connected

to the even way in the way pair, otherwise (if the LSB is 1) $n-1$ -to- 2^{n-1} decoder's output connected to the odd one.

3.2 Tag Array Organization

Cache tag arrays have relatively small capacity and therefore occupy a smaller fraction of the entire transistor budget compared with cache data arrays. For example, a simple tag array with valid and dirty bits is only 19 KB in the case of a 256 KB 8-way set-associative cache with 32-byte cache lines. This makes cache tag arrays typically have lower V_{min} than data arrays have.

However, because of the ever-increasing process variation, SRAM cell reliability is significantly degraded. In addition, data corruption in a cache tag because of unreliable low-voltage operations has severe negative impacts on the processor behavior. For example, a tag bit failure can cause cache hits to the wrong cache way, resulting in erratic read or write operations. Therefore we must handle even small SRAM arrays such as cache tag arrays.

In the proposed associativity-reconfigurable cache, the tag array is also implemented with 7T/14T SRAM. In the low operating-voltage mode, it pairs up odd-even ways and makes dependable ways along with its associated data array. Therefore the same control explained above can be carried out for mode transitions in the tag array.

3.3 Comparison with Traditional Approach

Traditionally, hardware redundancy based techniques have been widely used to improve yield and reliability at low voltages in on-chip cache design. For instance, many on-chip caches include redundant columns and/or rows, spare SRAM blocks to replace the faulty cache lines [10], [23]. Error coding techniques are also widely used to improve and correct erratic bits in cache lines and to improve reliability at low voltages [15], [24], [25].

However, these traditional schemes cannot tolerate at high failure rate at low voltages due to the ever-increasing process variations. In addition, these schemes necessitate large spare SRAM blocks and/or some cycle penalties which have influence on the entire processor performance.

On the other hand, In virtue of the lower BER in the dependable mode of 7T/14T SRAM than that of the traditional schemes [17], the proposed cache can achieve lower V_{min} of the entire cache than the traditional approaches. Although the proposed cache downsizes its capacity and associativity for reliable low voltage operations, it cache can reconfigure its capacity and associativity when higher performance and large cache is needed. The proposed cache can easily change its configuration to suit the requirement of current operating conditions. In addition, the proposed cache does not incur the cycle overhead. Furthermore, the proposed scheme is transparent to other traditional schemes. We can easily use the proposed cache with combination of the traditional approaches to obtain much higher reliability improvements at low voltages.

4. Associativity Reconfiguration Mechanism

In this section, we describe the mode transition in the proposed cache and its associativity reconfiguration process. In the proposed cache, the way pairs change their operating mode to the dependable mode and produce dependable cache ways according to the desired performance and voltage level of current workloads. Once the operating system or DVFS controller decide to change the processor operating voltage, an appropriate cache configuration is chosen to achieve required performance voltage level. The cache controller then selects way pairs the dependable mode will be applied to for satisfying the desired operating voltage.

In typical write-back cache implementations, cache lines may be dirty and thereby include data that are not written back to next level in the processor's memory hierarchy. Therefore, it is necessary to write back dirty cache lines in the way pairs to the next level cache or main memory before the transitions to the low-voltage mode.

Although the transition target way pairs cannot be accessed during the reconfiguration process, there are still cache lines in the other way. Therefore the entire cache can operate and serve cache accesses during the transition process. Note that a set-associative cache has multiple cache lines in a set.

4.1 Transition to the Low-Voltage Mode

In a write-back cache, there might be cache lines marked as dirty in the target way pairs. We must address dirty cache lines in a way pair that is a target of the mode transition. To preserve data in these dirty cache lines, we must write them back prior to the cache line duplication, which will be described in Sect. 4.2.

Figure 7 illustrates how to address these dirty cache lines in mode transition sequences. Note that we must write back only dirty cache lines in odd (even) indexed cache lines of even (odd) ways because of the opposite one will be copied and will therefore still be used after the mode transition. At first, we decide to apply the dependable mode for which way pairs are necessary to be switched to achieve the required voltage level. After the determination of way pairs, we examine whether the target way pair includes dirty cache lines in its cache ways. When target way pair includes

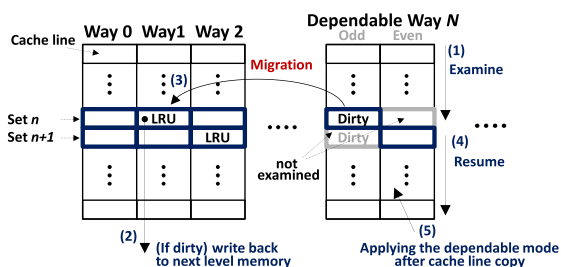


Fig. 7 Illustration of the mode transition process.

dirty cache lines, these dirty cache lines are migrated into the least recently used (LRU) cache lines in the same set. If the LRU cache lines are also dirty, these cache lines are written back to the write back queue preceding the migration process. If the dirty cache line in the target way pair itself is in the LRU state, then it will not be migrated but merely written back to the next level. Because the observation that dirty cache lines that reach the least recently used state hardly overwritten again [26], the migration process imposes negligible performance overhead.

The transition time to the dependable mode depends on the number of dirty cache lines in the way pair. If all the cache lines are clean, then there is no need to write back cache lines to the next level in the memory hierarchy. However, it is necessary to write back all the cache lines in the worst case if every cache line in the target way pair is dirty. In addition, the number of dirty cache lines and their distribution depend on the current processor state and workloads. Therefore the mode transition time cannot be determined. To reduce the number of dirty cache lines, we can use the eager write-back techniques proposed in [26]. Additionally, for the predictable transition time, we may determine the hard limit of the number of dirty cache lines.

4.2 Simultaneous Cache Line Copy in Way Pairs

Once the transition to the dependable mode finishes, the way pair composes one dependable way. However, all the cache lines in the dependable way are invalid at this time because they include only irregular data, even though SRAM cell margin is enhanced. It is unpredictable what data are stored in the cache lines after the transition process to the dependable mode. Such unpredictability makes the entire associativity virtually decreased by 2, immediately after the mode transition.

To solve this problem, we leverage the simultaneous block-level copy feature of 7T/14T SRAM [27]. As the 7T/14T SRAM cell connects its internal nodes using additional pMOS transistors, it is possible to transfer data between two SRAM cells through its additional transistors. By appropriately controlling supply voltage rails, word lines and the CTRL signals in Fig. 1, a block-level simultaneous copy function is realized. Entire copy sequences comprises four clock cycles [27].

Figure 8 shows how to use the simultaneous copy fea-




Fig. 8 Cache line simultaneous copy in a way pair. In the odd way of the way pair, cache lines from odd indexes are copied to even indexes, and vice versa in case of the even way.

ture of the 7T/14T SRAM in the proposed cache. As described above, a way pair comprises a dependable way at low voltages by interleaving its cache lines. Therefore, to solve the irregular data problem after the mode transition, even indexed cache lines in the odd numbered way are copied to odd indexed cache lines in it, and vice versa in the case of the even numbered way. The cache line duplication prevents the unpredictability of the data after the transition to dependable mode.

4.3 Transition to the High-Performance Mode

When the processor operating mode changes from low-voltage mode to high-performance mode, the proposed cache increases its associativity. It can be conducted easily by inactivating the CTRL signal in Fig. 1 and changing the SRAM cell operating mode to the normal mode. Note that it is unnecessary to write back dirty cache lines in this case.

After transition to the high-performance mode, a dependable way produces two odd-even ways of which cache lines operates in the normal mode. At this time each consecutive cache line includes the same data, because two SRAM cells constituting the 7T/14T SRAM cell have the identical data. Because consecutive indexed cache lines are associated different memory addresses, we must flush this half of cache lines. It can be done by setting cache status bit to the “invalid” state.

5. Experimental Evaluations

In this section, we describe our experimental evaluations of the proposed associativity-reconfigurable cache. The baseline cache system configuration is presented in Table 2. The minimum operating voltage (V_{\min}) improvement of the proposed cache is evaluated based on measurement results. Then we estimate the power dissipation and the energy efficiency of the each operating modes in the proposed cache. We also analyze impacts on the processor’s overall performance caused by the reconfiguration of associativity in the proposed cache. The mode transition time of the proposed cache is also evaluated. Finally, we estimate the area overhead of the proposed cache.

5.1 Minimum Operating Voltage Evaluation

We manufactured a 512-Kb 7T/14T SRAM macro using 65-nm CMOS technology. It consists of 32 16-Kb SRAM blocks. Figure 9 shows a layout of 16-Kb SRAM block and a die photograph of 512-Kb 7T/14T SRAM.

To evaluate the V_{\min} improvements in our proposed

cache, we measured these 512-Kb 7T/14T SRAM macros. It can be regarded as two ways of 8-way 256-KB L2 cache: i.e., a pair of odd-even ways in the proposed associativity-reconfigurable cache. Therefore, the V_{\min} improvement in the 8-way 256-KB L2 cache is evaluated with measuring four 512-Kb 7T/14T SRAM macros. A 512-Kb 7T/14T SRAM can also be matched to a 32-KB L1 cache. Therefore, we can evaluate the V_{\min} improvements in the 8-way 32-KB L1 caches by measuring one 512-Kb SRAM macro. We measured one more 512-Kb SRAM macro to evaluate the tag arrays in the L1 and L2 caches.

Figure 10 shows the measured V_{\min} s of each way pair in the 256-KB L2 cache tag and data arrays. In the normal mode, the measured V_{\min} s are 0.7 V, 0.66 V, 0.63 V and 0.59 V, respectively, whereas in the dependable mode, the respective values are reduced to 0.56 V, 0.54 V, 0.56 V and 0.53 V. As an 8-way L2 cache, the proposed cache can operate at 0.7 V. If Pair 0 changes its mode to the dependable mode, the proposed cache can operate at 0.66 V as 7-way cache. Similarly, the V_{\min} can be scaled by trading off the associativity.

Figure 11 summarizes the V_{\min} scalability with respect to the operating associativity in the proposed cache. If all four way pairs enter the dependable mode, then the proposed cache can operate as a 4-way 128-KB cache at 0.56 V, achieving a 140 mV lower V_{\min} than the 8-way 256-KB in the normal mode. Applying the dependable mode in one way pair reduces V_{\min} by 30–40 mV. We also evaluated the V_{\min} of 32 KB 8-way set-associative cache in the same way. The 32-KB 8-way cache can operate in the normal mode at 0.62 V. This value is larger than the 6-way L2 cache’s V_{\min} . Consequently, in the case in which the L2 cache operates as a 4-way 128-KB, 32-KB 8-way L1 caches enter the depend-

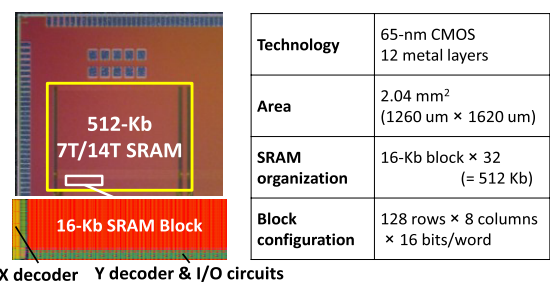


Fig. 9 512-Kb 7T/14T SRAM die photograph and 16-kb block layout.




Fig. 10 Measured V_{\min} s of each way pairs in the 256 KB cache.

Table 2 Baseline cache system configuration.

Level-1 Cache	32-KB 8-way set-associative cache (with 2.75 KB tag array)
Level-2 Cache	256 KB 8-way set associative cache (with 19 KB tag array)




Fig. 11 Measured V_{min} of the proposed caches in 256 KB and 32 KB cache with respect to each operating associativity.




Fig. 12 Power dissipation with respect to operating associativity in the proposed cache.

able mode for achieving minimum V_{min} . In this case, the L1 caches operate as 4-way 16-KB caches when the L2 cache operates as a 4-way 128-KB cache.

5.2 Power and Energy Efficiency

Based on the results of minimum operating voltage evaluations described in Sect. 5.1, we evaluated the power dissipation and the energy per instruction (EPI) of the proposed cache using a 65-nm process technology. We used our appropriately modified version of CACTI [29], a tool widely used for cache modeling, to evaluate the power dissipation of the proposed cache. For estimation of dynamic power and EPI, we also performed SPICE simulations to determine the operating frequencies in each operating associativity of the proposed cache, under the assumption of 24 fanout-of-4 (FO4) inverter chain delay per cycle [28].

Figure 12 shows our evaluation results of power dissipation. By reconfiguring the operating associativity from 8 to 4, the power dissipation and EPI are reduced by 67.48% and 26.70%, respectively. It can be observed that the power dissipation and EPI of the proposed cache is reduced as the operating associativity is decreased. In other words, the proposed cache can trade off the power dissipation and the energy efficiency with the operating associativity.

Table 3 Baseline processor configuration.

Parameter	Value
Processor frequency	1 GHz
L1 Instruction Cache	32 KB, 8-way, 32-byte line, 3-cycle access time
L1 Data Cache	32 KB, 8-way, 32-byte lines, 3-cycle access time
Unified L2 Cache	256 KB, 8-way, 32-byte lines, 10-cycle access time
Cache Replacement Policy	LRU
External DRAM latency	100 cycles

5.3 Processor Performance Evaluation

Since the proposed cache reduces the associativity and capacity in the low-operating voltage mode, it affects a cache hit rate and therefore the processor performance; it is necessary to evaluate the impact on the performance. To evaluate the impacts on the processor's overall performances, we conducted processor simulations using the proposed cache architecture with respect to various L2 cache configurations. We used the cycle accurate gem5 simulator [30] and choose 9 CINT and 14 CFP benchmarks from SPEC 2006 benchmark suite [31]. All simulations were executed for 2 billion instructions (with 1 billion warm-up period). We chose instructions per cycle (IPC) as the indicator of the processor performance. To investigate and quantify the impacts on the processor's overall performance caused by the reconfiguration of associativity and capacity in the proposed cache, we fixed the processor frequency in our performance evaluations. Table 3 shows baseline processor configuration used in our processor simulations.

Figure 13 shows normalized instructions per cycle (IPCs) for each benchmark in SPEC2006 CINT and CFP relative to the normal mode, the 256 KB/8-way case. The IL1 and DL1 cache operate as 32 KB 8-way set-associative cache. The maximal IPC loss is 12.91% (a 4-way/128 KB case in gromacs). The IPC degradation is 0.72% on average when the dependable mode is applied to a single way pair and the cache associativity decreases by one. The average IPC degradation is 2.95% in the 128-KB 4-way L2 cache (in the case in which all the pairs operate in the dependable mode).

5.4 Transition Time Estimation

In this section, we evaluate the mode transition time of the proposed cache. As described in Sect. 4, the proposed cache migrates its dirty cache lines from the target way pair to the LRU way in the set when the transitions to the low-voltage operating modes occur. To estimate the mode transition time of the proposed cache, we first estimated the cache cycle time for the mode transition by CACTI [29]. In this estimation, we assumed that the proposed cache set the hard limit of its dirty cache lines on each way pair. We added four




Fig. 13 Normalized IPCs of SPEC2006 CINT and CFP benchmarks with respect to each operating cache associativity.

Table 4 Transition time estimation results.

Core Freq.	Migration cycles	Dirty cache line limit	Transition cycles
1 GHz	14 cycles	16 lines	228 cycles
1 GHz	14 cycles	32 lines	452 cycles
1 GHz	14 cycles	64 lines	900 cycles

Table 5 Area estimation results in 65-nm CMOS process.

Scheme	Area (mm ²)	Norm. area	Scheme	Area (mm ²)	Norm. area
6T SRAM 32-KB	1.29616	1	6T SRAM 256-KB	3.14161	1
Proposed 32-KB	1.31703	1.01610	Proposed 256-KB	3.31412	1.05491

cycles to the entire migration cycles because the proposed cache duplicates the cache lines after the dirty cache line migrations.

We estimate the transition time with respect to the several number of dirty cache lines. The estimation results are shown in Table 4. The estimation results show that up to 64 dirty cache line limits, the mode transition of the proposed cache can be executed within 1 μ s.

5.5 Area Estimation

In this section, we present our evaluation results of the area overhead of the proposed cache architecture. The 7T memory cell area is 11% greater than that of the conventional 6T memory cell [17]. We referred [27] for peripheral circuitry overhead for simultaneous block level copy. We assumed that the dedicated decoders have a negligible impact on the overall area because they occupy a much smaller proportion than SRAM arrays in the entire transistor number.

We estimated the area overhead of the proposed cache in the 32-KB and 256-KB 8-way set-associative cache. We used our modified version of CACTI [29] for these estimations. The results are shown in Table 5. In 65-nm process technology, the proposed cache imposes a 1.61% and 5.49% area overhead in 32-KB L1 and 256-KB L2 cache, respectively. Note that, for a small cache, the proposed cache adds a small area overhead because the fraction of the indispensable peripherals in area is relatively large.

6. Conclusion

In this paper, we proposed a novel adaptive cache design designated as the associativity-reconfigurable cache, for reliable wide-range low-voltage operations. The proposed cache has ability to reconfigure its associativity and capacity. This reconfigurability makes it possible to achieve optimal cache configuration which is suitable for the desired performance and voltage level of current workloads. We described details of its organization, reconfiguration mechanisms and several experimental results. Measurement-based evaluation results show that the proposed cache possesses the scalable characteristic of low-voltage reliability and that it can decrease V_{min} by 140 mV. We also evaluated the power dissipation and the energy efficiency, and the evaluation results show that the proposed cache can reduce its power dissipation and energy efficiency by 67.48% and 26.70%, respectively. Our processor performance evaluation results shows that applying the proposed cache architecture results in 2.95% maximum IPC loss but it can choose various performance levels. The area overhead of the proposed cache is only 1.61% and 5.49% in 32-KB and 256-KB cache respectively.

Acknowledgments

This work was supported by VLSI Design and Education Center (VDEC), The University of Tokyo in collaboration with the Semiconductor Technology Academic Research Center (STARC), e-Shuttle Inc., and Fujitsu Ltd.

References

- [1] J. Jung, Y. Nakata, S. Okumura, H. Kawaguchi, and M. Yoshimoto, "256-KB associativity-reconfigurable cache with 7T/14T SRAM for aggressive DVS down to 0.57 V," Proc. IEEE International Conference on Electronics, Circuits, and Systems, pp.524–527, Dec. 2011.
- [2] B. Rogers, A. Krishna, G. Bell, K. Vu, X. Jiang, and Y. Solihin, "Scaling the bandwidth wall: Challenges in and avenues for CMP scaling," Proc. International Symposium on Computer Architecture, pp.371–382, June 2009.
- [3] N. Kim, T. Austin, D. Blaauw, T. Mudge, K. Flautner, M. Irwin, M. Kandemir, and V. Narayanan, "Leakage current: Moore's law meets static power," Computer, vol.36, no.12, pp.68–75, Dec. 2003.
- [4] B. Stackhouse, S. Bhimji, C. Bostak, D. Bradley, B. Cherkauer, J. Desai, E. Francom, M. Gowan, P. Gronowski, D. Krueger, C.

- Morganti, and S. Troyer, "A 65 nm 2-billion transistor quad-core Itanium processor," *IEEE J. Solid-State Circuits*, vol.44, no.1, pp.18–31, Jan. 2009.
- [5] M. Floyd, M. Allen-Ware, K. Rajamani, B. Brock, C. Lefurgy, A.J. Drake, L. Pesantez, T. Gloekler, J.A. Tierno, P. Bose, and A. Buyuktosunoglu, "Introducing the adaptive energy management features of the Power7 chip," *IEEE Micro*, vol.31, no.2, pp.60–75, March 2011.
- [6] A. Branover, D. Foley, and M. Steinman, "AMD Fusion APU: Llano," *IEEE Micro*, vol.32, no.2, pp.28–37, March 2012.
- [7] E. Rotem, A. Naveh, A. Ananthakrishnan, E. Weissmann, and D. Rajwan, "Power-management architecture of the Intel microarchitecture code-named Sandy Bridge," *IEEE Micro*, vol.32, no.2, pp.20–27, March 2012.
- [8] K. Itoh, "Adaptive circuits for the 0.5-V nanoscale CMOS era," *IEEE International Solid-State Circuits Conference Digest of Technical Papers*, pp.14–20, Feb. 2009.
- [9] A.J. Bhavnagarwala, X. Tang, and J.D. Meindl, "The impact of intrinsic device fluctuations on CMOS SRAM cell stability," *IEEE J. Solid-State Circuits*, vol.36, no.4, pp.658–665, April 2001.
- [10] S. Borkar, T. Karnik, and V. De, "Design and reliability challenges in nanometer technologies," *Proc. Design Automation Conference*, p.75, June 2004.
- [11] S. Mukhopadhyay, H. Mahmoodi, and K. Roy, "Modeling of failure probability and statistical design of SRAM array for yield enhancement in nanoscaled CMOS," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol.24, no.12, pp.1859–1880, Dec. 2005.
- [12] P.P. Shirvani and E.J. McCluskey, "PADded cache: A new fault-tolerance technique for cache memories," *Proc. VLSI Test Symposium*, pp.440–445, April 1999.
- [13] A. Agarwal, B.C. Paul, H. Mahmoodi, A. Datta, and K. Roy, "A process-tolerant cache architecture for improved yield in nanoscale technologies," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol.13, no.1, pp.27–38, Jan. 2005.
- [14] M.A. Makhzan, A. Khajeh, A. Eltawil, and F. Kurdahi, "Limits on voltage scaling for caches utilizing fault tolerant techniques," *Proc. International Conference on Computer Design*, pp.488–495, Oct. 2007.
- [15] J. Kim, N. Hardavellas, K. Mai, B. Falsafi, and J. Hoe, "Multi-bit error tolerant caches using two-dimensional error coding," *Proc. International Symposium on Microarchitecture*, pp.197–209, Dec. 2007.
- [16] C. Wilkerson, H. Gao, A.R. Alameldeen, Z. Chishti, M. Khellah, and S.-L. Lu, "Trading off cache capacity for reliability to enable low voltage operation," *Proc. International Symposium on Computer Architecture*, pp.203–214, June 2008.
- [17] H. Fujiwara, S. Okumura, Y. Iguchi, H. Noguchi, H. Kawaguchi, and M. Yoshimoto, "A 7T/14T dependable SRAM and its array structure to avoid half selection," *Proc. International Conference on VLSI Design*, pp.295–300, Jan. 2009.
- [18] E. Seevinck, F.J. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," *IEEE J. Solid-State Circuits*, vol.SC-22, no.5, pp.748–754, Oct. 1987.
- [19] S. Yoshimoto, T. Amashita, S. Okumura, K. Yamaguchi, M. Yoshimoto, and H. Kawaguchi, "Bit error and soft error hardenable 7T/14T SRAM with 150-nm FD-SOI process," *Proc. International Reliability Physics Symposium*, pp.SE.3.1–SE.3.6, April 2011.
- [20] H. Fujiwara, S. Okumura, Y. Iguchi, H. Noguchi, Y. Morita, H. Kawaguchi, and M. Yoshimoto, "Quality of a Bit (QoB): A new concept in dependable SRAM," *Proc. International Symposium on Quality Electronic Design*, pp.98–102, March 2008.
- [21] S. Rusu, S. Tam, H. Muljono, D. Ayers, J. Chang, B. Cherkauer, J. Stinson, J. Benoit, R. Varada, J. Leung, R.D. Limaye, and S. Vora, "A 65-nm dual-core multithreaded Xeon® processor with 16-MB L3 cache," *IEEE J. Solid-State Circuits*, vol.42, no.1, pp.17–25, Jan. 2007.
- [22] Texas Instruments, "3A processor supply with I²C compatible interface and remote sense," [Online]. Available: <http://www.ti.com/lit/ds/slvsa9b/slvsa9b.pdf>
- [23] S.E. Schuster, "Multiple word/bit line redundancy for semiconductor memories," *IEEE J. Solid-State Circuits*, vol.SC-13, no.5, pp.698–703, Oct. 1978.
- [24] P. Kongetira, K. Aingaran, and K. Olukotun, "Niagara: A 32-way multithreaded Sparc processor," *IEEE Micro*, vol.25 no.2, pp.21–29, March 2005.
- [25] N. Quach, "High availability and reliability in the Itanium processor," *IEEE Micro*, vol.20, no.5, pp.61–69, Sept. 2000.
- [26] H.-H.S. Lee, G.S. Tyson, and M.K. Farrens, "Eager writeback — A technique for improving bandwidth utilization," *Proc. International Symposium on Microarchitecture*, pp.11–21, Dec. 2000.
- [27] S. Okumura, S. Yoshimoto, K. Yamaguchi, Y. Nakata, H. Kawaguchi, and M. Yoshimoto, "7T SRAM enabling low-energy simultaneous block copy," *Proc. IEEE Custom Integrated Circuits Conf.*, pp.1–4, Sept. 2010.
- [28] V. Srinivasan, D. Brooks, M. Gschwind, P. Bose, V. Zyuban, P.N. Strenski, and P.G. Emma, "Optimizing pipelines for power and performance," *Proc. International Symposium on Microarchitecture*, pp.333–344, Nov. 2002.
- [29] N. Muralimanohar, R. Balasubramanian, and N.P. Jouppi, "CACTI 6.0: A tool to model large caches," Technical Report, HPL-2009-85, Hewlett Packard Laboratories, April 2009.
- [30] N. Binkert, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M.D. Hill, D.A. Wood, B. Beckmann, G. Black, S.K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D.R. Hower, and T. Krishna, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol.39, no.2, pp.1–7, Aug. 2011.
- [31] Standard Performance Evaluation Corporation, "The SPEC CPU 2006 Benchmark Suite," <http://www.specbench.org>



IEEE.

Jinwook Jung received the B.E. degree in computer and systems engineering from Kobe University, Hyogo, Japan, in 2011, with the support of the Korea-Japan Joint Government Scholarship Program for the Students in Science and Engineering Departments. He is currently working toward the M.E. degree in system informatics at the same university. His current research interests include low-power variation-tolerant circuit techniques and architecture designs for reliability. He is a student member of



Yohei Nakata received the B.E. and M.E. degrees in computer and systems engineering from Kobe University, Hyogo, Japan in 2008 and 2010, respectively, where he is currently pursuing the Ph.D. degree in engineering. His current research interests include multi-core processor architecture, low-power processor and dependable processor designs. He is a student member of the IEEE and IPSJ.



Shunsuke Okumura received his B.E. and M.E. degrees in Computer and Systems Engineering in 2008 and 2010, respectively from Kobe University, Hyogo, Japan, where he is currently working in the doctoral course. His current research is high-performance, low-power SRAM designs, dependable SRAM designs, and error correcting codes implementation. He is a student member of IPSJ and IEEE.



Hiroshi Kawaguchi received B.Eng. and M.Eng. degrees in electronic engineering from Chiba University, Chiba, Japan, in 1991 and 1993, respectively, and earned a Ph.D. degree in electronic engineering from The University of Tokyo, Tokyo, Japan, in 2006. He joined Konami Corporation, Kobe, Japan, in 1993, where he developed arcade entertainment systems. He moved to The Institute of Industrial Science, The University of Tokyo, as a Technical Associate in 1996, and was appointed as a

Research Associate in 2003. In 2005, he moved to Kobe University, Kobe, Japan. Since 2007, he has been an Associate Professor with The Department of Information Science at that university. He is also a Collaborative Researcher with The Institute of Industrial Science, The University of Tokyo. His current research interests include low-voltage SRAM, RF circuits, and ubiquitous sensor networks. Dr. Kawaguchi was a recipient of the IEEE ISSCC 2004 Takuo Sugano Outstanding Paper Award and the IEEE Kansai Section 2006 Gold Award. He has served as a Design and Implementation of Signal Processing Systems (DISPS) Technical Committee Member for IEEE Signal Processing Society, as a Program Committee Member for IEEE Custom Integrated Circuits Conference (CICC) and IEEE Symposium on Low-Power and High-Speed Chips (COOL Chips), and as an Associate Editor of IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences and IPSJ Transactions on System LSI Design Methodology (TSLDM). He is a member of the IEEE, ACM, and IPSJ.



Masahiko Yoshimoto joined the LSI Laboratory, Mitsubishi Electric Corporation, Itami, Japan, in 1977. From 1978 to 1983 he had been engaged in the design of NMOS and CMOS static RAM. Since 1984 he had been involved in the research and development of multimedia ULSI systems. He earned a Ph.D. degree in Electrical Engineering from Nagoya University, Nagoya, Japan in 1998. Since 2000, he had been a professor of Dept. of Electrical & Electronic System Engineering in Kanazawa University, Japan.

Since 2004, he has been a professor of Dept. of Computer and Systems Engineering in Kobe University, Japan. His current activity is focused on the research and development of an ultra-low power multimedia and ubiquitous media VLSI systems and a dependable SRAM circuit. He holds on 70 registered patents. He has served on the program committee of the IEEE International Solid State Circuit Conference from 1991 to 1993. Also he served as Guest Editor for special issues on Low-Power System LSI, IP and Related Technologies of IEICE Transactions in 2004. He was a chair of IEEE SSCS (Solid State Circuits Society) Kansai Chapter from 2009 to 2010. He is also a chair of The IEICE Electronics Society Technical Committee on Integrated Circuits and Devices from 2011–2012. He received the R&D100 awards from the R&D magazine for the development of the DISP and the development of the real time MPEG2 video encoder chipset in 1990 and 1996, respectively. He also received 21st TELECOM System Technology Award in 2006.