

A Low Power TLB Structure for Embedded Systems

Jin-Hyuck Choi[†], Jung-Hoon Lee[†], Seh-Woong Jeong^{*}, Shin-Dug Kim[†], and Charles Weems[‡]

[†]Dept. of Computer Science, Yonsei University, Seoul 120-749, Korea

^{*}Media SOC Team, System LSI Business, Samsung Electronics Co., Ltd., Kihueung, Korea

[‡]Dept. of Computer Science, University of Massachusetts, Amherst, MA01003-4610, USA

Abstract--We present a new two-level TLB (translation look-aside buffer) architecture that integrates a 2-way banked filter TLB with a 2-way banked main TLB. The objective is to reduce power consumption in embedded processors by distributing the accesses to TLB entries across the banks in a balanced manner. First, an advanced filtering technique is devised to reduce access power by adopting a sub-bank structure. Second, a bank-associative structure is applied to each level of the TLB hierarchy. Simulation results show that the Energy*Delay product can be reduced by about 40.9% compared to a fully-associative TLB, 24.9% compared to a micro-TLB with 4+32 entries, and 12.18% compared to a micro-TLB with 16+32 entries.

Keywords--Bank associative structure, filter mechanism, low power design, translation look-aside buffer

I. INTRODUCTION

Recently power consumption has become an important factor in designing high performance embedded processors. Architectural approaches offer significant improvements in reducing power consumption that go beyond techniques at the gate or circuit level. Thus the goal of this work is to design a low power TLB structure for high performance embedded processors. The TLB is a small on-chip cache for recently used virtual to physical address translations. Although it is very small, the TLB is costly in terms of power consumption because it is a fully associative structure [1]. To reduce power consumption, a micro or filter TLB with a small number of entries has been used to reduce the number of entries accessed at once and to support faster access.

Current TLBs are typically implemented with static RAM cells for data and CAM (content addressable memory) cells for tags. For low power TLB design, memory cells are redesigned using techniques such as modifying the CAM cell [1] and reducing voltage [3]. However, the method in [1] suffers from higher hardware cost and complexity, and

performance degradation. As [3] shows, supply voltage is an important parameter controlling CMOS power consumption.

For a fully associative TLB structure, power consumption tends to increase abruptly as the number of entries increases [1]. The total number of TLB entries accessed at once should be less than 64 or 128. However, higher performance can be achieved if more TLB entries are provided. Thus, one method is to divide the entire TLB space into two banked TLBs so that the number of entries accessed at once can be reduced to less than 32 or 64 [4]. This structure consumes less power than a fully associative TLB because only a half of the CAM entries are looked up on each access to the TLB. But a major drawback is performance degradation due to the tendency to encounter more capacity misses in a banked system.

Thus, we have devised a two-level dual TLB structure that integrates a 2-way banked filter TLB with a 2-way banked main TLB to achieve low power consumption and high performance. Experimental results show that our TLB design can reduce the power consumption and the miss ratio significantly, compared with other approaches.

II. LOW POWER TLB DESIGN

Using a fast lookup buffer at address translation time is a very common technique among embedded processor manufacturers and its hit ratio is approximately 50%. If a miss occurs at this buffer, another cycle is incurred to access the main TLB. One approach to reducing internal power dissipation is to avoid a large number of accesses to the main TLB, by using an extended filtering technique, which is borrowed from the filter cache design [2]. However, simply adapting the filtering scheme to the TLB can be problematic. In particular, the TLB space can be reduced by the inclusion property and this in turn tends to increase the overall number of TLB misses, increasing power consumption. Thus, a method to improve the filter TLB hit ratio is necessary.

One method is to enlarge the filter TLB capacity per page entry, i.e., the filter TLB is designed to hold more entry slots by maintaining more PPN (physical page number) slots per VPN (virtual page number) entry; called an extended sub-bank structure. Consider the k sub-bank case. The same tag

can be used for k PPNs for a given VPN entry, and the number of tag bits for the k sub-banks is reduced by $\log_2 k$ bits. This enables the filter to hold more entries, thereby increasing the filter hit ratio. Also, power reduction is achieved by activating only one sub-bank entry at a time, using the $\log_2 k$ selection bits.

Another approach to reducing power consumption is a bank associative structure. The x -way banked TLB structure partitions the entire TLB into x banks by modulo x , and so the least significant $\log_2 x$ tag bits of any virtual address are used as the selection bits to activate a specific bank at a time. By using this scheme, we can reduce the power dissipation to approximately $1/x$ in accessing the TLB. We have adopted a 2-way banked structure to minimize the additional cost of the sense-amps and output drivers necessitated by increasing the bank associativity. Thus our proposed TLB is constructed in two parts; a 2-way banked filter TLB with an extended sub-bank structure and a 2-way banked main TLB, as shown in Fig. 1.

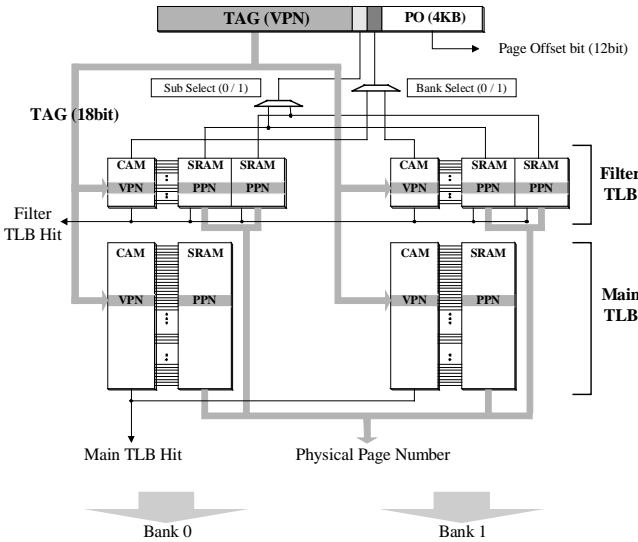


Fig. 1. Organization for a combination of filter and main TLBs.

The operational flow is based on a modified two-level TLB model. When the CPU generates a particular virtual address, the filter TLB is searched first. To select one of the two banks, we use the least significant bit of the tag field as a bank select bit, i.e., 0 or 1. Three possible cases can be considered as follows.

A. Hit in the selected bank of the filter TLB

If a page is found in the activated bank of the filter TLB, the actions are the same as for a conventional TLB hit. The requested physical address is sent to the cache and compared with its tag bits. But if a miss occurs at the filter TLB, the main TLB is activated during the next cycle for a possible match.

B. Hit in the selected bank of the main TLB

If a requested page is found in the main TLB, the requested physical page is sent to the cache and compared with its tag bits. At the same time, its corresponding entry is moved into the filter TLB and it is invalidated in the main TLB to prevent entry duplication. And if there is an evicted entry from the filter TLB, it is stored back to an invalidated entry of main TLB for temporal reuse.

C. Miss in both TLBs

When misses occur at both TLBs, while the memory management unit is handling the miss, a new page entry is loaded into the filter TLB directly. Only when a filter TLB entry is evicted according to the LRU (least recently used) policy, the evicted page entry is moved into the main TLB for temporal reuse. Using this scheme, the total number of effective VPNs is essentially one large TLB space consisting of the filter TLB space plus the main TLB space. In contrast, the conventional micro TLB is a hierarchical structure based on the inclusion property. Thus, we effectively increase TLB space by avoiding the inclusion property.

III. PERFORMANCE EVALUATION

There are four essential performance metrics, i.e., miss ratio, average memory access time (AMAT), power consumption, and Energy*Delay product which we use to evaluate and compare the proposed TLB system with other approaches. For our comparison, we focus on the two most commonly used TLB architectures, i.e., a fully associative TLB and a micro-TLB. Several experiments were performed to determine the optimum number of filter TLB entries and bank associativity. We concluded that the number of indexing entries should be 4 and that a 2-way banked structure should be used to minimize the additional hardware cost (e.g., additional comparator, sense-amps, output-driver, etc.).

Fig. 2 and Fig. 3 show the average miss ratios and the average TLB access times for the following TLB structures, assuming a 4KB page size. Fully-associative (FA) TLBs with 32 entries and 48 entries, denoted by FA32 and FA48, are used because they hold the same number of PPNs as our main TLB and the combined resources of our TLB structure, respectively. A micro (filter)-TLB of 4 entries with a main TLB of 32 entries, denoted by Micro TLB (4-32), and a micro (filter)-TLB of 16 entries with 32 entries in its main TLB, denoted by Micro TLB (16-32), are also used for comparison. The former is a typical configuration and the latter is more comparable in terms of the number of PPNs provided.

The time to access the filter TLB is assumed to be one cycle and if a miss occurs, then the main TLB is searched during the next cycle. The notation, “4-4-16-16 entries,” denotes the proposed TLB in a configuration of a 2-way

banked filter TLB with 4+4 entries and a 2-way main TLB with 16+16 entries. Every filter (micro) TLB in this experiment uses LRU replacement. The time needed for a miss handling interrupt routine is assumed to be 15cycles, which is based on values for common 32-bit embedded processors (e.g., Hitachi SH4 or ARM920T).

The power consumption of each TLB configuration is evaluated with the CACTI-II simulator [5], which is modified for TLB power evaluation. To evaluate power consumption in the proposed TLB, power dissipated at each bank is evaluated and accumulated separately. And then the overall power consumption is obtained by summing the power consumption for both TLBs and additional power consumption at the multiplexer. In the simulator, only one selected bank is activated to drive the tag on the bit lines of the CAM, and enable the SRAM word lines and sense amplifiers to read out the data. There is no additional delay due to the multiplexer because the bank select bit can be applied to the multiplexer, while the tag is compared at the CAM.

Benchmark analysis shows that the percentage of one-cycle hits at the filter TLB is almost 95% in the proposed TLB




Fig. 2. Miss ratios of proposed TLB and various TLB organizations.




Fig. 3. Average memory access time (AMAT) of the proposed TLB and various TLB organizations.

Table I shows the amount of power consumed for each event in a single TLB access. The significant difference in power consumption between a TLB with 128 entries and that with 64 entries, as shown in Table 1, comes from the fact that the power consumed at the match line and bit lines in the CAM tends to grow considerably when the number of entries increases beyond 64. Thus, we used 32 entries in our experiments to show the reduction in power consumption that we can realistically expect.

TABLE I
Power consumption per access for various TLB sizes.

Number of entries in fully associative TLB	Read hit (nJ)	Read/miss (nJ)	Write (nJ)
4	2.2630	0.6154	0.2397
16	2.9209	1.1021	0.5526
32	3.7575	1.7511	0.7097
48	4.6512	2.4502	0.8610
64	5.4200	3.0490	0.9908
128	8.7272	5.6447	1.8551

The average power consumption of a fully associative TLB is given by:

$$Avg.power = N_{hit} * P_{hit} + N_{miss} * P_{miss}, \quad (1)$$

where N_{hit} and N_{miss} are the ratios of hits and misses in the TLB respectively. Also P_{hit} and P_{miss} are the power consumption required to process a hit and a miss respectively. P_{miss} can be calculated as follows:

$$P_{miss} = P_{CAM} + P_{write} + P_{off}, \quad (2)$$

where P_{CAM} is the power dissipated at all the entries when the tag part of the TLB is accessed, and P_{write} is the power dissipated in the data area and tag area in order to update an entry in the case of a miss. P_{off} is the power dissipated at the cache and pads when a TLB miss occurs. P_{off} is calculated as follows:

$$P_{off} = P_{cache_acc} + M_{cache_miss} * (P_{cache_write} + P_{pad}), \quad (3)$$

where P_{cache_acc} is the power used to access a cache block, M_{cache_miss} is the cache miss ratio, P_{cache_write} is the power consumption when a cache write operation occurs in a cache miss, and P_{pad} is the power dissipated at the on-chip pad slot. A 32KB 2-way set associative data cache with 32-byte block size is assumed for the processor memory hierarchy and the supply voltage is 4.5 Volts. The values of M_{cache_miss} , P_{cache_acc} , P_{cache_write} , and P_{pad} are 5%, 21.291nJ, 10.145nJ, and 6.48nJ respectively.

Finally, the logic components required to construct an

additional TLB entry are a comparator, a valid signal driver, and a data output driver, whose power consumption values are assumed to be 0.03nJ, 0.01nJ, and 0.2nJ respectively. These values were obtained using modified DineroIV and CACTI-II simulators.

Fig 4 and Fig. 5 present the power consumption and Energy*Delay product for the different TLB structures. The delay value is the average TLB access time. Clearly, the proposed TLB system is the best structure in terms of power consumption and the performance improvement tends to be more significant than the other structures. It should be noted that our 4-4-16-16 configuration is smaller than the 16-32 Micro-TLB and the 48 entry associative TLB.




Fig. 4. Power consumptions of the proposed TLB and various TLB organizations.

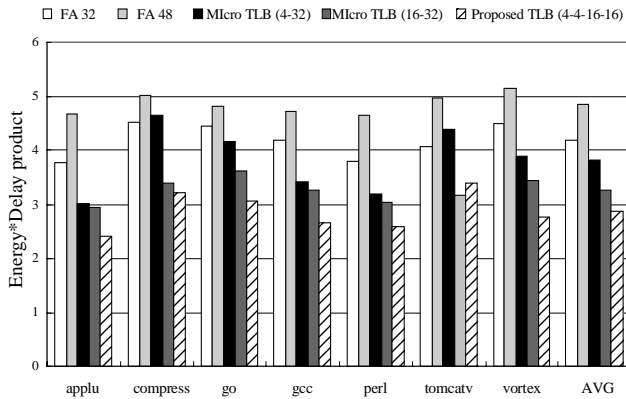


Fig. 5. Energy*Delay products of the proposed TLB and various TLB organizations.

Table II provides simulation results for performance evaluation. In order to compare the proposed TLB with the other TLB organizations, we used a metric called IR (Improvement Ratio), $IR = ((a-b)/b) * 100\%$.

TABLE II
Improvement ratio over various TLB structures.

	IR of miss ratio	IR of AMAT	IR of power	IR of E*D product
FA 32	59.34 (%)	- 8.41 (%)	36.94 (%)	31.51 (%)
FA 48	22.07 (%)	- 13.41 (%)	48.03 (%)	40.90 (%)
Micro TLB (4-32)	59.26 (%)	20.56 (%)	11.06 (%)	24.90 (%)
Micro TLB (16-32)	40.81 (%)	- 8.91 (%)	19.52 (%)	12.18 (%)

IV. CONCLUSION

We devised a TLB system for low power consumption in embedded processors. The proposed TLB is configured as a 2-way banked filter TLB and a 2-way banked main TLB. This is to minimize the power consumption by reducing the overall number of entries to be accessed at a time, and enhance the effective TLB space by avoiding the inclusion property. As shown in our experiments, the performance of our TLB design can reduce the power consumption by about 48.03%, compared with a fully associative TLB that has the same number of PPNs, and also it can reduce the miss ratio and Energy*Delay product by 59.26% and 24.9%, respectively, compared with the micro TLB with 4-32 entries, and 40.81% and 12.18%, compared to the micro TLB with 16-32 entries.

ACKNOWLEDGEMENT

This work was supported by the Samsung MMU project and partially by Korea Research Foundation under contract number EA0081.

REFERENCES

- [1] T. Juan, T. Lang, J. Navarro, "Reducing TLB Power Requirements," *Int'l Symp. on Low Power Electronics and design*, 1997.
- [2] J. Kin, M. Gupta, and W. H. Mangione-Smith, "The Filter Cache: An Energy Efficient Memory Structure," In *Int'l Symp. Microarchitecture*, 1997, pp. 184-193.
- [3] D. Liu and C. Svensson, "Trading Speed for Low Power by Choice of Supply and Threshold Voltages," *IEEE Journal of Solid State Circuits*, Vol. 28, No. 1, 1993. pp. 10-17.
- [4] S. Manne, A. Klauser, D. Grunwald, and F. Somenzi, "Low power TLB Design for High Performance Microprocessors," *Univ. of Colorado Technical Report*, 1997.
- [5] G. Reinman and N. Jouppi, "An Integrated Cache Timing and Power Model," *Compaq WRL Report*, 1999.