LOW POWER DUAL MATCHLINE TERNARY CONTENT ADDRESSABLE MEMORY

Nitin Mohan and Manoj Sachdev

Electrical and Computer Engineering, University of Waterloo, Ontario, Canada N2L 3G1

ABSTRACT

Content addressable memories (CAMs) are very attractive for high-speed table lookups in modern network systems. This paper presents a low-power dual match line (ML) ternary CAM (TCAM) to address the power consumption issue of CAMs. The highly capacitive ML is divided into two segments to reduce the active capacitance and hence the power. We analyze possible cases of mismatches and demonstrate a significant reduction in power (up to 43%) for a small penalty in search speed (4%).

1. INTRODUCTION

A phenomenal increase in the number of Internet users and real-time network applications requires the networks to support very high-speed data transfer. The table lookup tasks, which were traditionally done in software, are being replaced by hardware solutions. New requirements such as flow analysis, policy based routing and quality of service (QoS) are increasing the quantity and variety of the table lookups. For example at OC-192 line rates, up to 30 million packets have to be processed every second [1]. The packet processing requires multiple lookups for each packet. This increases the table lookup requirement to 100 million searches per second [1]. As the line rates increase from OC-192 to OC-768 and higher, software based solutions will not be able to meet the lookup requirements. Therefore, content addressable memory (CAM) will become the only viable solution for lookups at such speeds. A CAM searches an input keyword within a table of stored data and returns the address of the matched location. In the presence of multiple matches, a priority encoder returns the highest priority address. The search is performed in parallel, and hence it shows a fixed latency. However, CAMs are more expensive and power hungry due to per cell comparison logic and parallel search operation.

A number of techniques on circuit design, and architecture style have been proposed in the past to reduce the power consumption in CAMs, but most of them are focused on binary CAMs. Recently, an increasing demand for more efficient and feature rich ternary CAM (TCAM) has driven the research and development of low-power TCAMs. A TCAM is similar to a binary CAM with the additional feature of storing and searching three states: '0', '1' and 'X' (don't care). The 'X' option can be used to search the data in a given range of values. The TCAMs are particularly suitable for classless inter-domain routing (CIDR) longest prefix matching and packet classification in single search operation [1].

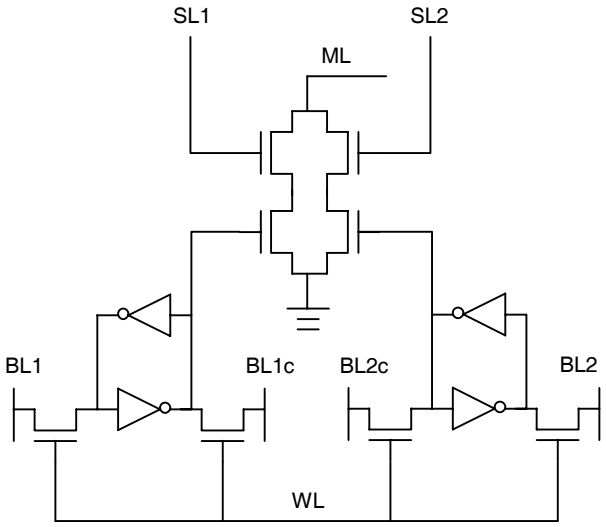

Figure 1. Circuit schematic of 16T TCAM cell

2. TCAM REVIEW

A TCAM cell consists of two RAM cells and a pass transistor based comparison logic (Figure 1). A TCAM chip is implemented as an array of TCAM cells with bit lines (BLs) and search lines (SLs) running vertically and word lines (WLs) and match lines (MLs) running horizontally. All the cells in the same column share the corresponding BLs and SLs. Similarly, all the cells in the same row share the corresponding WL and ML. Data is read and written using BLs and WLs similar to the conventional RAMs. The parallel data search operation is performed by precharging the MLs to power supply voltage (V_{DD}) and placing the input keyword at the SLs. If there is a mismatch in any cell connected to an ML, it will discharge to ground through the comparison logic of that cell. Typically, only a few entries in the lookup table

match with the input keyword, and most of the MLs are discharged to ground. This leads to high power consumption due to the large capacitance of MLs. The second major source of power consumption in CAMs is the switching activity of highly capacitive SLs. During the ML precharge phase, the SLs must be discharged to ground to avoid static current through the comparison logic of the mismatched cells.

Several techniques have been proposed to reduce the power consumption in CAMs. A selective precharge CAM reduces the power by precharging the MLs only if there is match in the selected bits of the words [2]. The number of selected bits is optimized for minimum average energy. In an application specific CAM for data compression, the power consumption is reduced by shutting down the power for redundant comparisons [3]. The redundant comparisons are directly derived from the data compression algorithm. In another implementation, the voltage swing of ML is reduced by charging it through an NMOS transistor and discharging it by PMOS comparison logic [4]. A precomputation-based CAM was recently proposed that reduces the power by extracting a smaller size parameter from each word [5]. The parameter of the input keyword is first compared with all the pre-computed parameters. Then, only the words with matched parameters are compared with the input keyword. Hence, the power reduction is achieved at the expense of slower search speed. Moreover, increased complexity due to the additional logic makes it difficult to cascade multiple CAMs. Recently, a current-race sensing scheme is proposed that reduces the power by reducing SL switching activity and limiting the voltage swing of MLs [6]. The current-race sensing scheme is shown in Figure 2. Initially, all the MLs are discharged to ground, and then each ML is charged by a current source. A dummy word is also added to generate a reference signal (MLOFFc) that turns off all the current sources when the match detection is done. The dummy word has a forced match condition. If there is a mismatch in an ML, it has a conducting path to ground, and it charges slower than the dummy ML (DML). When the DML voltage reaches the threshold voltage of the sense amplifier (SA), its output (DMLSO) rises from '0' to '1'. The delayed inverted version of the DMLSO signal is used to turn off all the current sources. The delay element ensures that voltages of all the matched MLs reach the threshold voltages of their respective sense amplifiers even in the presence of process variations.

Most of the schemes described above focus on reducing the ML voltage swing and the switching activity. We propose a dual ML TCAM that saves power by reducing the active ML capacitance. In this paper, we present this scheme for an SRAM-based TCAM. However, it can be easily extended to the DRAM-based TCAMs since both types of TCAMs have similar comparison logic circuits [1].

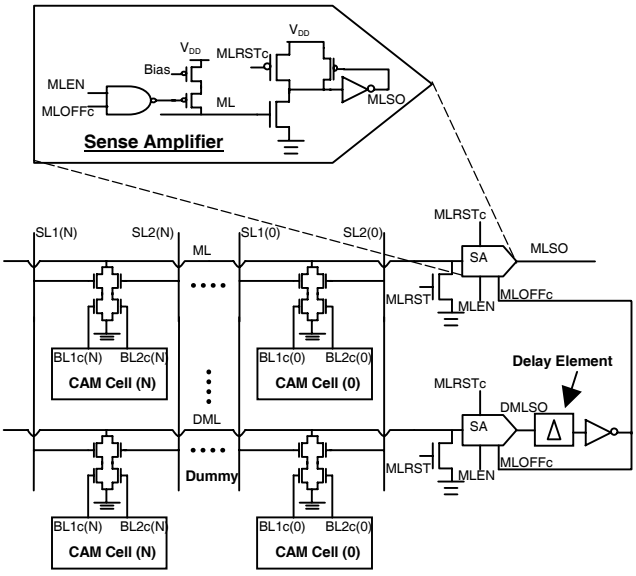

Figure 2. Current-race sensing scheme [6]

3. DUAL MATCHLINE TCAM

The dual ML TCAM is shown in Figure 3. Although we use the current-race ML sensing to illustrate the dual ML TCAM, our scheme can also be implemented with other types of ML sense amplifiers. The match lines ML1 and ML2 connecting to the left and the right side of the comparison logic respectively, have separate sense amplifiers. All the cells in a row share ML1 and ML2. If we neglect the interconnect capacitance of match lines ML1 and ML2, their capacitance will be half of the ML capacitance of the traditional scheme (Figure 1). ML1 and ML2 are initially discharged to ground and the input keyword is applied to SLs. Now the first sense amplifier (SA1) is enabled, and a current source charges ML1. If a mismatch is found in ML1, the second sense amplifier (SA2) remains off. Hence theoretically, the power consumption in this case is reduced by half. On the other hand if a match is found in ML1, SA2 is also enabled and there is no power savings in this case. Although the comparison in ML1 and ML2 takes place sequentially, the speed of this scheme remains unchanged because the capacitance of ML1 or ML2 is half of the ML capacitance. Therefore, they charge two times faster than ML for the same current. Moreover, the delay (Δ) of the dummy word output directly scales with the ML capacitance. Since smaller capacitance charges faster, it requires a smaller delay to compensate for the SA threshold variations. Therefore, the delays of the dummy words ($\Delta 1$ and $\Delta 2$) associated with ML1 and ML2 are also reduced by half.

In the above discussion we neglected the effect of interconnect capacitance. Since both the ML1 and ML2 run horizontally across the whole array, the interconnect capacitance of the dual ML TCAM is two times larger

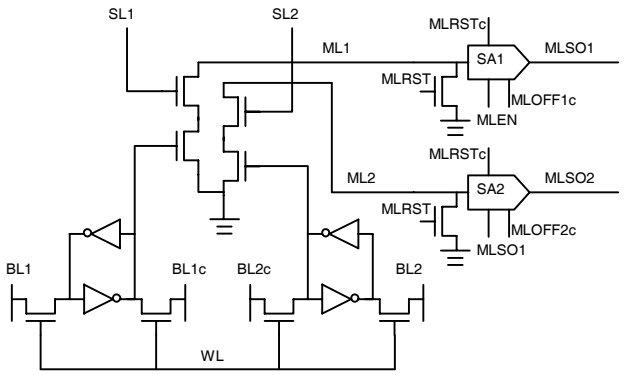

Figure 3: Dual ML TCAM scheme

than that of the traditional TCAM. Therefore, the dual ML TCAM pays some penalty in terms of performance. The layout of a typical 16T TCAM cell (Figure 1) in 0.18 µm CMOS technology has an approximate horizontal dimension of 10 µm. Considering the metal 4 routing of ML with minimum width, the interconnect capacitance of ML per cell is extracted to 0.19 fF. Typical word size (N) for CAMs in networking applications is 144 [1]. Hence, the total interconnect capacitance of each ML can be estimated to C_{INT} = 144 x 0.19 fF = 27.36 fF. The ML capacitance is composed of the drain capacitance of the NMOS transistors whose gates are connected to SLs. However, if some of the SLs are at V_{DD}, the corresponding NMOS transistors conduct, and the intermediate nodes also contribute to the total ML capacitance. The contribution of these intermediate nodes is absent for the global masking condition (SL1 = SL2 = '0'). Hence, the ML capacitance is smaller for this condition. The ML capacitance is also dependent on ML voltage, which varies from 0 to $V_{DD}/2$ for the current-race sensing scheme. In the absence of global masking, the average ML capacitance of the traditional scheme is calculated to 347 fF for the devices in the comparison logic with W/L = 0.6 µm/0.18 µm. If all the bits are globally masked, the average ML capacitance is calculated to 190 fF. Therefore, the energy reduction of the dual ML scheme varies from 44% to 46% depending on the number of bits with global masking condition. The above calculation assumes that most of the mismatches are detected by SA1 as we will demonstrate in the next section.

4. RESULTS AND DISCUSSION

We simulated the traditional and the dual ML TCAMs in Cadence using 0.18 µm TSMC CMOS technology design kit, and compared the power consumption and search time of the two schemes. In all the simulations, W/L of NMOS transistors in the comparison logic is 0.6 µm/0.18 µm and the delays of the dummy ML output for the traditional and dual ML TCAMs are 535 ps and 266 ps respectively. The typical configuration of a TCAM block is 256 x 144 [1]. Therefore in our simulations, we included a dummy load at the output of the dummy delay element (MLOFFc in Figure 2). The dummy load emulates the input capacitance of 256 NAND gates associated with the 256 SAs (Figure 2). We also connected half of the SLs to V_{DD} and rest of them to ground. This arrangement gives average value of the ML capacitance. The total energy consumption of TCAM is dominated by the mismatch energy since majority of the MLs in TCAM do not match. The simulation results with and without interconnect capacitance are given in Table 2 and Table 1 respectively.

Table 1: Simulation results without C_{INT}

	Traditional	Dual ML
Search time (T_S)	7.88 ns	7.88 ns
Mismatch Energy: ML1 (E_1)	747 fJ	394 fJ
Mismatch Energy: ML2 (E_2)	747 fJ	918 fJ

Table 2: Simulation results with C_{INT} = 27.36 fF

	Traditional	Dual ML
Search time (T_S)	8.14 ns	8.46 ns
Mismatch Energy: ML1 (E_1)	769 fJ	426 fJ
Mismatch Energy: ML2 (E_2)	769 fJ	973 fJ

Here E_1 is the energy consumption when a mismatch is detected by SA1. Similarly, E_2 is the energy consumption when a mismatch is detected by SA2. In the traditional TCAM, ML1 and ML2 are connected, and the energy of operation corresponds to a mismatch in ML.

As expected from the previous analysis, the search time of both schemes are the same when interconnect capacitance is ignored. The energy E_2 is higher than the mismatch energy of the traditional scheme. This is due to the transient current when a match is detected in ML1, and MLSO1 rises from '0' to '1'. The simulations results with interconnect capacitance show that the dual ML scheme results in a small trade-off (4%) in search speed. The mismatch energies also increase due to the increased capacitance of the MLs.

In the dual ML TCAM, a mismatch can be detected by SA1 or SA2. Their corresponding energies are E_1 and E_2 respectively. If the probability of mismatch detection by SA1 is P_{ML1}, the average energy can be given by,

$$E_{AVG} = P_{ML1} \times E_1 + (1 - P_{ML1}) \times E_2$$

A TCAM cell can result in two types of mismatches as shown in Table 3.

Table 3: Types of mismatches in a TCAM cell

Mismatch	SL1	SL2	BL1	BL2
Type I	0	1	1	0
Type II	1	0	0	1

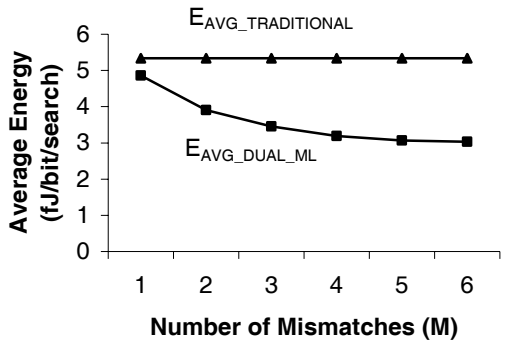

Figure 4: Reduction in search energy for Dual ML TCAM

Since both types of mismatches are equally probable, their probability of occurrence is 0.5. It can be shown from Figure 3 that Type I mismatch cannot be detected by SA1. However, all the cells in a row share the same ML1. Therefore, only one Type II mismatch is sufficient for SA1 to detect mismatch. If the number of mismatches in a given row is M, the probability of the event that all the mismatches are of Type I is $(0.5)^M$. Hence, the probability that at least one of the mismatch belongs to Type II i.e. the probability of detecting mismatch by SA1 is:

$$P_{ML1} = 1 - (0.5)^M$$

The following table highlights the energy and power reduction of the dual ML TCAM for different values of M:

Table 4: Energy (power) reduction of the dual ML TCAM

M	P_{ML1}	E_{AVG} (fJ/word/search)	Reduction in energy (power) (%)
1	0.50	700	9
2	0.75	563	27
3	0.87	497	35
4	0.94	459	40
5	0.97	442	43
6	0.98	437	43

Therefore, the energy (power) consumption is reduced by 43% if 5 or more bits of the stored data words (144 bits) do not match with the corresponding bits of the input keyword. Figure 4 illustrates the variation of average energy per bit per search with the number of mismatches. For 5 or more mismatches, the average energy becomes almost constant.

In principle, the dual ML TCAM is similar to the selective precharge scheme [2]. The selective precharge scheme achieves power reduction by assuming that most of the TCAM entries have one or more mismatches in a small number of selected bits. However, for a given set of selected bits, this assumption may not be valid in different applications. The number and position of the selected bits for minimum energy consumption may vary with the type of data-traffic and application. This problem is eliminated in the dual ML TCAM by distributing ML1 among all the bits. Hence, the dual TCAM scheme is more general in application. The attractive features of this scheme come at the expense of slightly larger area to accommodate an additional ML and an additional sense amplifier per word. The ML sense amplifiers (SA1 and SA2) in every row should be placed horizontally to match their vertical height with TCAM cells for more compact chip layout.

5. CONCLUSION

A dual ML TCAM scheme is presented. It is shown by theoretical analysis and simulation results that the dual ML TCAM results in significant power reduction (up to 43%) at the expense of small trade-off in speed (4%). The proposed scheme is orthogonal to most existing power saving techniques and can easily be combined with the other schemes.

6. ACKNOWLEDGMENT

This work is supported by NSERC PGS B. The authors would like to thank Wilson W. Fung and Derek W. Wright for several stimulating discussions on this work.

7. REFERENCES

[1] M. Wirth, "The next generation of content addressable memories," *Application Note*, MOSAID Technologies, http://www.mosaid.com/semiconductor/papers/NextGenerationofCAMs.pdf.

[2] C. Zukowski and S. Wang, "Use of selective precharge for low-power content-addressable memories," *Proc. IEEE ISCAS*, pp. 1788-1791, Jun. 1997.

[3] K.-J. Lin and C.-W. Wu, "A low-power CAM design for lz data compression," *IEEE Trans. On Computers*, vol. 49, no. 10, pp. 1139-1145, Oct. 2000.

[4] H. Miyatake, M. Tanaka and Y. Mori, "A design for high-speed low-power CMOS fully parallel content-addressable memory macros," *IEEE Journal of Solid-state Circuits*, vol. 36, no. 6, pp. 956-968, Jun. 2001.

[5] C.-S. Lin, J.-C. Chang and B.-D. Liu, "A low-power precomputation-based fully parallel content-addressable memory," *IEEE Journal of Solid-state Circuits*, vol. 38, no. 4, pp. 654-662, Apr. 2003.

[6] I. Arsovski, T. Chandler, A. Sheikholeslami, "A ternary content-addressable memory (TCAM) based on 4T static storage and including a current-race sensing scheme," *IEEE Journal of Solid-state Circuits*, vol. 38, no. 1, pp. 155-158, Jan. 2003.