

Low-Capacitance and Charge-Shared Match Lines for Low-Energy High-Performance TCAMs

Nitin Mohan, *Member, IEEE*, and Manoj Sachdev, *Senior Member, IEEE*

Abstract—Emerging high-speed lookup-intensive applications are demanding ternary content addressable memories (TCAMs) with large word-sizes, which suffer from lower search speeds due to larger match line capacitance. Therefore, low-energy high-performance design techniques are needed, which improve the search speeds of TCAMs without increasing their power consumption. In this paper, we present a cell-level comparison logic and a charge-shared match line scheme. Both schemes reduce search time and energy in TCAMs. Measurement results of the above schemes, implemented in 0.18- μm CMOS technology, show a search time reduction of 42% and 11%, and a search-energy reduction of 25% and 9%, respectively.

Index Terms—Content addressable memory (CAM), low energy, high performance, charge recycling, ternary content addressable memory (TCAM).

I. INTRODUCTION

TERNARY content addressable memories (TCAMs) are attractive for high-speed lookup-intensive applications such as packet classification and forwarding in network routers [1]. The increasing line rates are demanding fast TCAMs [2]. On the other hand, the growing deployment of IPv6 is increasing the word-size of TCAMs [3]. A wide TCAM is inherently slower, and it requires circuit techniques to improve its search speed. Typically, TCAMs are power-hungry due to their high-speed parallel searches. The increasing speed requirements are further escalating their power consumption. Arguably the power consumption can be sustained at its original level if the speed-enhancement techniques also reduce the search-energy.

In this paper, we present a cell-level comparison logic that offers smaller capacitance on the match lines (MLs). We also present a charge-shared ML scheme, which recycles and transfers one ML segment's charge to another ML segment. Both schemes reduce the search time and energy in TCAMs.

II. LOW-CAPACITANCE MATCH LINE

A significant portion of the TCAM power is consumed in switching highly capacitive MLs. Fig. 1(a) shows a conven-

tional 16T TCAM cell. The ML capacitance depends on the masking conditions of the connected cells. For example, if a cell is not masked, it adds a capacitance of $4C_D$ to the corresponding ML, where C_D is the drain capacitance of a comparison logic transistor that includes the bottom-plate and side-wall junction capacitances [Fig. 1(b)]. Similarly, globally and locally masked cells add capacitances of $2C_D$ and $4C_D$, respectively [Fig. 1(c), (d)].

We present a cell-level comparison logic that offers a smaller ML capacitance (Fig. 2). The proposed comparison logic requires an additional line (SelGbl) to keep node G at ground under the global masking condition ($SL1 = SL2 = "0"$). SelGbl is generated by NORing SL1 and SL2, and it is shared by all the cells in the same column. Since $SL1 = SL2 = "1"$ is an invalid state, the possibility of shorting the inverter outputs is eliminated. Similar comparison logic (without transistor M2) has been used in binary CAMs [4]. However, it has not been reported in TCAMs possibly due to floating node G in a globally masked cell. The proposed comparison logic employs transistor M2 for driving node G to ground in a globally masked cell.

If none of the bits are globally masked and interconnect capacitance is ignored, the proposed comparison logic reduces the ML capacitance by 75% ($4C_D$ to C_D). Similarly, if all the bits are globally masked, the ML capacitance is reduced by 50% ($2C_D$ to C_D). Therefore, the capacitance reduction varies between 75% and 50% depending on the number of globally masked bits. However, this reduction in ML capacitance comes at the expense of additional lines (SelGbl) and associated energy consumption. Fortunately, the rate of updating the global mask registers is negligibly less than the table lookup frequency in most TCAM applications [5]–[8]. Thus, the power consumed in switching SelGbl is negligibly less than the power consumed in switching MLs.

Typically, each ML employs an ML sense amplifier (MLSA) to differentiate between “match” and “mismatch” conditions by sensing ML pull-down currents, which can be given by the following equations:

$$I_{ON_ML} = N_{MM}I_{ON_CELL} + (N - N_{MM})I_{OFF_CELL} \quad (1)$$

$$I_{OFF_ML} = NI_{OFF_CELL} \quad (2)$$

where N is the word-size, N_{MM} is the number of mismatched bits, I_{ON_CELL} is the ML current contribution of each mismatched cell, and I_{OFF_CELL} is the ML current contribution of each matched cell. MLSA is designed to distinguish between “match” and “mismatch” conditions even in the worst case when $N_{MM} = 1$. The proposed comparison logic has only one nMOS transistor in the ML pull-down path (Fig. 2). Hence,

Manuscript received October 24, 2005; revised March 27, 2007. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) postgraduate scholarship (PGS B) and Micronet R&D. Test chip fabrication was provided by Canadian Microelectronics Corporation (CMC).

N. Mohan is with Advanced Micro Devices (AMD), Inc., Boston Design Center, Boxborough, MA 01719, USA (e-mail: nitin.mohan@amd.com; nitinm@ieee.org).

M. Sachdev is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: msachdev@ece.uwaterloo.ca).

Digital Object Identifier 10.1109/JSSC.2007.903089

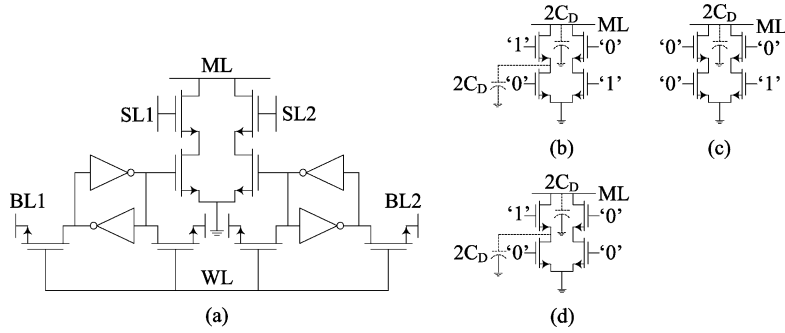


Fig. 1. (a) A 16T TCAM cell and its contribution to the ML capacitance under the following conditions: (b) no masking, (c) global masking, (d) local masking.

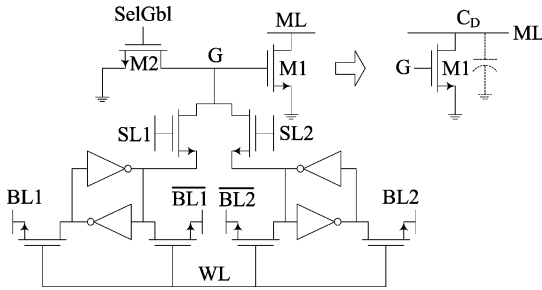


Fig. 2. A TCAM cell based on the proposed comparison logic, and its contribution to the ML capacitance under all the masking conditions.

its ML ON current (I_{ON_ML}) is greater than that of the conventional comparison logic, which has two series-connected transistors instead (Fig. 1). Simulation and measurement results also support the above deduction even though the voltage swing of node G is limited to $(V_{DD} - V_{th})$. Furthermore, the proposed comparison logic has only one ML leakage path per cell. Thus, its ML OFF current (I_{OFF_ML}) is less than that of the conventional comparison logic, which has two ML leakage paths per cell. Although one of the ML leakage paths of a masked cell can have a lower leakage due to body effect, the I_{OFF_ML} of the conventional scheme is still larger than that of the proposed scheme due to absence of the body effect in other leakage paths [Fig. 1(b), (c), (d)]. A higher I_{ON_ML}/I_{OFF_ML} ratio has a number of advantages. First of all, it makes the ML sensing less sensitive to process variations and operating conditions. Secondly, it can be exploited to reduce the search time without causing a false “match”. For example, a larger I_{ON_ML} can handle a larger pull-down or pull-up current in ML sensing. Finally, a higher I_{ON_ML}/I_{OFF_ML} ratio allows the implementation of wide TCAMs because I_{OFF_ML} is proportional to N .

We analyzed the proposed and conventional comparison logic circuits by implementing them in two 145-bit-wide TCAM words as shown in Fig. 3(a). A charge-redistribution MLSA is used for ML sensing whose timing diagram is shown in Fig. 3(b) [9]. All the control signals are common to both MLSAs. Initially, the MLs are discharged to ground using PRE. The search operation is initiated by the rising edge of EN, and the falling edges of FastPre and PRE. The ML voltage swing is restricted by the nMOS transistors (M1 and M2) whose gates are connected to a reference voltage (V_{REF}). The FastPre pulse

precharges the MLs to a voltage near $(V_{REF} - V_{th})$. The evaluation begins with the rising edge of the FastPre signal. Under the “match” condition, the ML does not have a pull-down path, and its node SP remains at V_{DD} . Under the “mismatch” condition, the node SP is pulled down to 0. A small current source (I_{REF}) at the node SP compensates for ML leakages. In our design, I_{REF} has been set to one-fifth of I_{ON} .

The width of the FastPre pulse (T_{FP}) is the most critical parameter in the charge-redistribution MLSA [9]. If T_{FP} is too small, MLs will be precharged to a voltage much lower than $(V_{REF} - V_{th})$. Under the “match” condition, this incomplete precharge can cause a false glitch at the MLSA output by charge-sharing ML_{New} and SP (Fig. 3). This false glitch increases energy consumption and affects the operation of the next stage. However, a wider T_{FP} pulse also increases the energy consumption due to direct current paths in the mismatched words. Larger values of T_{FP} also increase the search time. Therefore, the T_{FP} window is chosen just wide enough to avoid a false glitch under the “match” condition. The false glitch problem is less severe for the proposed low-capacitance ML due to two main reasons. First, a lower capacitance implies a faster precharging of MLs to $(V_{REF} - V_{th})$, which also reduces search time and energy. Second, the voltage drop at node SP due to charge-sharing between nodes SP and ML is less severe if ML capacitance is smaller.

III. CHARGE-SHARED MATCH LINES

Most low-power TCAMs divide large MLs into smaller segments and sense them sequentially [10]–[12]. For example, a 144-bit-wide ML can be divided into two segments of 36 and 108 bits [11]. Each segment has a separate MLSA. First the smaller segment (ML1) is sensed. The larger segment (ML2) is sensed only if ML1 matches the corresponding portion of the search key. Therefore, this scheme saves power only in the best case, which occurs when first segments of most words do not match the same portion of the search key. The optimum size of ML1 is determined from the data statistics. If a TCAM has segments optimized for one application, it does not give the best-case power in other applications. The average power consumption varies between the best case and the worst case depending on the application.

We propose a charge-shared ML scheme that reduces the worst-case energy consumption. Fig. 4 illustrates the scheme and its timing diagram using a current-race MLSA [13]. The

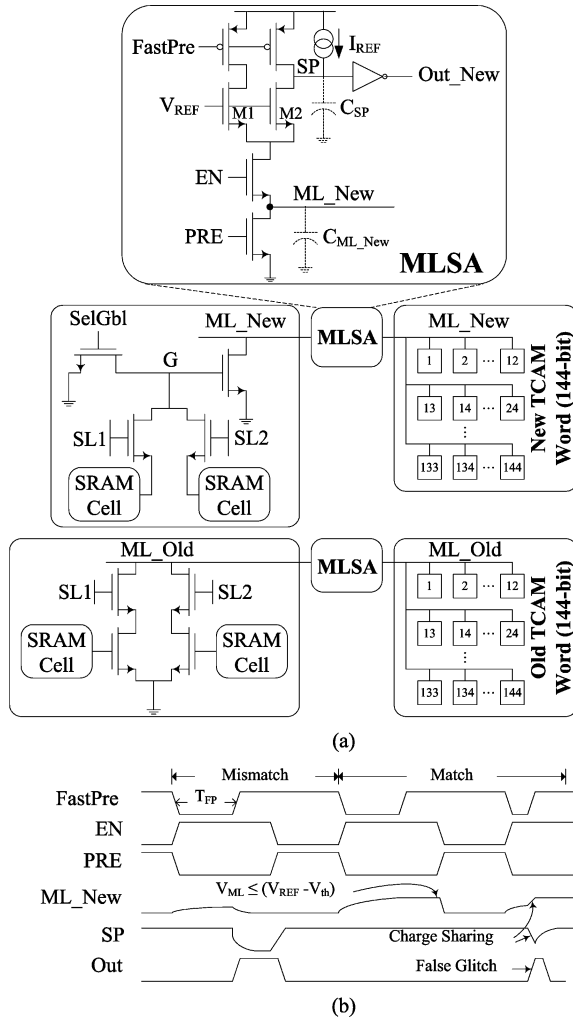


Fig. 3. (a) TCAM words employing the proposed and conventional comparison logic circuits with the charge-redistribution MLSA, and (b) its timing diagram.

current-race MLSA requires a dummy word to generate the control signals. The dummy ML is also divided into two segments (DML1 and DML2). All the cells of the dummy segments are locally masked, so both the dummy MLSAs generate a “match” in every search cycle. Since the ML capacitance varies with global masking, DML1 and DML2 should also track these variations. This is ensured by sharing the common search lines (SLs) with the dummy word. A rising edge of MLEN1 begins the search operation by enabling all the MLSAs in the first segment (MLSA1, DMLSA1, etc.). A rising edge of DMLSO1 indicates the completion of the search operation in the first segment. A delayed version of DMLSO1 signal ($\overline{MLOFF1}$) is used to turn off the MLSAs in the first segment. The delay (T_{CS}) ensures that all the matched words are detected before MLSAs are turned off. If the first segment of an ML matches with the corresponding portion of the search key, its MLSO1 turns on the corresponding MLSA2.

In the end of every search cycle, the conventional schemes discard the residual ML1-charge to the ground. The proposed scheme recycles the ML1-charge to reduce the search time and the worst-case energy consumption. If the first segment of a word results in a “match”, its ML1 is charge-shared with its

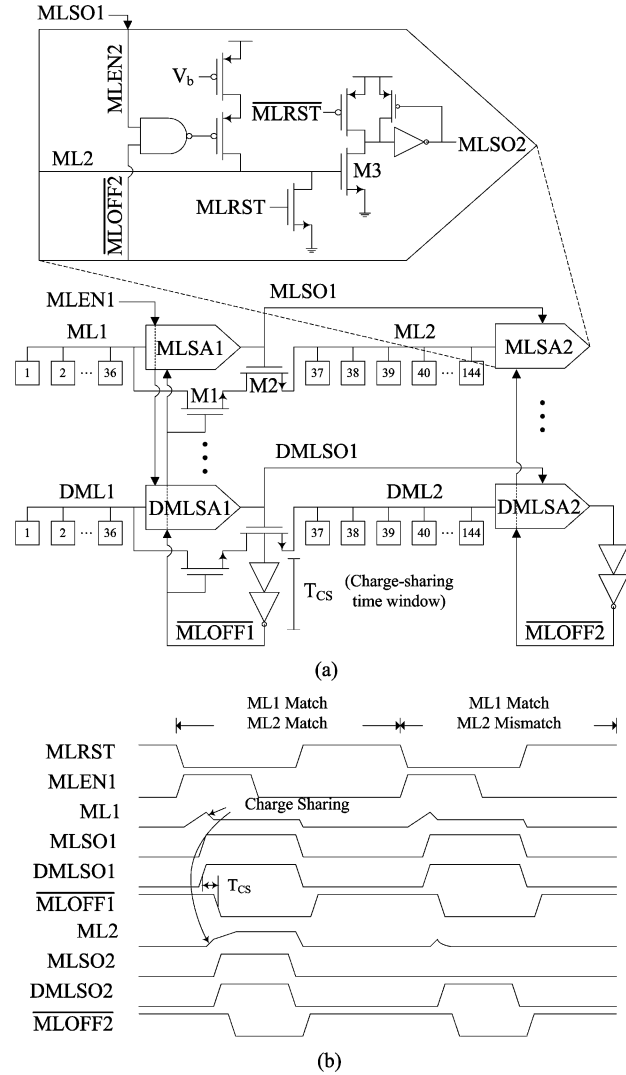


Fig. 4. (a) Circuit schematic of the proposed charge-shared ML scheme using a current-race ML sense amplifier, and (b) its timing diagram.

ML2 using transistors M1 and M2 (Fig. 3). The charge-sharing between DML1 and DML2 expedites the arrival of DMLSO2, which turns off MLSA2s. Since MLSA2s are enabled for a smaller duration, this scheme reduces both search time and worst-case energy. The charge-sharing between ML1 and ML2 begins at the rising edge of MLSO1 and ends at the falling edge of $\overline{MLOFF1}$ [Fig. 3(b)]. The time needed to charge-share ML1 and ML2 depends on the size of transistors M1 and M2. Larger transistors equalize ML1 and ML2 faster. However, oversized transistors also increase the ML capacitance. Therefore, their sizes should be optimized by simulations.

IV. TEST CHIP MEASUREMENT RESULTS

We implemented the low-capacitance ML and the charge-shared ML schemes (illustrated in Figs. 3 and 4, respectively) on two test chips in 0.18- μm CMOS technology. The micrographs of the two test chips are shown in Fig. 5. The first test chip contains a 144-bit word of the conventional TCAM cells and a 144-bit word of the proposed TCAM cells. Both types of cells consume the same number of transistors and hence, almost the




Fig. 5. Micrographs of the first and second test chips showing low-capacitance and charge-shared ML schemes, respectively.

same area. Each 144-bit word is arranged in an array of 12×12 cells and all the cells are hard-wired for “match” condition due to area constraints. One conventional cell and one proposed cell are connected to the respective words in parallel. The “match” and “mismatch” conditions for the two words are obtained by changing the status of these two cells. Thus, each word is effectively 145 bits wide. Both words have separate MLSAs and power-supply pins to measure the energy consumption. A reference circuit is also included to generate bias voltages for the MLSAs. All the control signals are common to both MLSAs.

The second test chip contains two 144-bit TCAM words and their dummies. One word and its dummy employ the standard current-race MLSA [13]. The other word and its dummy employ the current-race MLSA with the proposed charge-shared MLs (Fig. 3). Typically, a full-size TCAM block (256×144) contains only one dummy word. Thus, the signals $\overline{MLOFF1}$ and $\overline{MLOFF2}$ are shared by 256 words (Fig. 4). In order to imitate this capacitive loading, we included dummy loads in the second test chip (Fig. 5).

Fig. 6 shows measurement results of the first test chip. Here, T_{FP} and energy of the conventional and proposed schemes are shown for different values of TCAM cell supply voltage (V_{DD_Cell}) while the supply voltage of SLs and MLSAs remains at 1.8 V. Energy is measured for the “mismatch” condition since most words fail to match the search key in typical TCAM applications. A reduction in V_{DD_Cell} reduces I_{ON} and I_{REF} .




Fig. 6. FastPre pulse duration and energy measurement results of conventional and low-capacitance ML schemes for different values of V_{DD_Cell} .

However, the search operation is performed successfully as long as I_{REF} is large enough to compensate for ML leakages. A small V_{DD_Cell} also reduces the static power which is becoming a serious issue in sub-100-nm technologies [14]. Measurement results confirm that the low-capacitance ML scheme gives consistent energy (25%) and time (42%) savings for a large range of V_{DD_Cell} . Since C_{SP} is much smaller than C_{ML_New} , the SP voltage drops almost immediately after the rising edge of the FastPre pulse (Fig. 3). Hence, the reduction in the ML sensing time is almost equal to the reduction in T_{FP} . For smaller values of V_{DD_Cell} , T_{FP} increases due to a reduction in I_{REF} . Energy is less affected by the variations in V_{DD_Cell} because an increase in T_{FP} is compensated by a reduction in I_{ON} . For very low values of V_{DD_Cell} , the reduction in I_{ON} becomes more prominent and the energy consumption decreases.

Fig. 7 shows the search time and energy of the conventional and the charge-shared ML schemes measured for a range of I_{BIAS} (Fig. 4). Increasing I_{BIAS} reduces the search time and increases the search-energy. The charge-shared ML scheme gives a consistent improvement over the conventional scheme.

V. DISCUSSION

In the chip implementation of the conventional comparison logic (Fig. 5), the drains of upper two transistors were merged (Fig. 1). Thus, each unmasked cell added a capacitance of $3C_D$ to the corresponding ML (instead of $4C_D$ as estimated in Section II). We calculated the ML capacitances of the two 145-bit words, which are comprised of nMOS transistors with $W/L = 0.6\mu\text{m}/0.18\mu\text{m}$ and $C_D = 0.606$ fF:

$$C_{ML_Old} = 145 \times 3C_D + 141 \text{ fF} = 404.61 \text{ fF}$$

$$C_{ML_New} = 145 \times C_D + 148 \text{ fF} = 235.87 \text{ fF}$$

where second terms are the extracted interconnect capacitance (including the bottom-plate, fringe, and coupling capacitance) for the two MLs. Thus, the new comparison logic reduces the ML capacitance by 42%.




Fig. 7. Search time and energy measurement results of conventional and charge-shared ML schemes for different values of I_{BIAS} .

In order to verify the above calculations, we performed an indirect measurement of the ML capacitance. As explained in Section II, T_{FP} is typically chosen (3–6 ns) just large enough to avoid the false glitch under the match condition. In order to charge both C_{ML_Old} and C_{ML_New} approximately to the same voltage ($V_{REF} - V_{th}$), we set both MLs in match condition and chose a much larger value of T_{FP} ($T_{FP} = 15$ ns, period = 20 ns). Since a matched ML has no conducting path to ground, the whole current drawn from the power-supply is consumed in charging the ML capacitance. The average currents drawn by C_{ML_Old} and C_{ML_New} from the power-supply have been measured to $I_{Old} = 21.7 \mu A$ and $I_{New} = 13.9 \mu A$. Therefore, the ML capacitance is reduced by 36%, which is less than the expected value (42%). This implies that the value of T_{FP} (15 ns) is not large enough, and C_{ML_New} is charged to a slightly higher voltage than C_{ML_Old} . Both ML capacitances can be charged to a voltage much closer to ($V_{REF} - V_{th}$) by further increasing T_{FP} . However, large values of T_{FP} reduce the average power-supply currents, and this method loses its accuracy due to two main reasons. First, the average power-supply currents become comparable to ML leakages. Second, it becomes difficult to measure a very small current accurately. We also observed that a variation in V_{DD_Cell} does not change the measured power-supply currents, which reinforces the fact that there is no V_{DD_Cell} -dependent conducting path from ML to ground under the match condition.

The proposed comparison logic has greater I_{ON_ML} than the conventional comparison logic as explained in Section II. We indirectly measured the approximate I_{ON_ML} of the conventional and proposed comparison logic circuits. We set both MLs in mismatch condition, and chose $T_{FP} = 15$ ns and period = 20 ns. In this case, the average current drawn from the power-supply is proportional to I_{ON_ML} once the ML voltage reaches steady state. Fig. 8 shows the measured I_{ON_ML} of the two comparison logic circuits. For $V_{DD_Cell} = 1.8$ V, I_{ON_ML} of the proposed comparison logic is 14% higher than that of the conventional comparison logic. For smaller values of V_{DD_Cell} , the




Fig. 8. Measurement results for I_{ON} of the conventional and proposed comparison logic circuits.

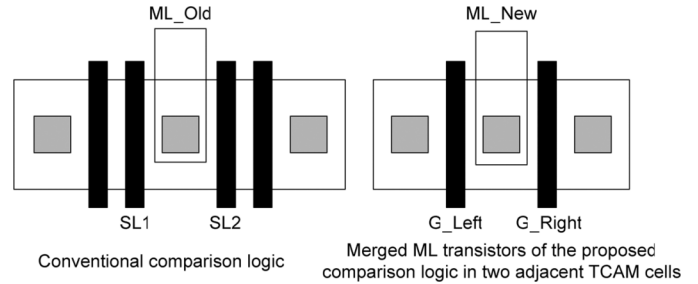


Fig. 9. Suggested layouts of conventional and proposed comparison logic circuits.

difference between the two ON currents reduces because the time taken in reaching steady state becomes comparable to T_{FP} , and the measurements become less accurate.

The low-capacitance ML scheme can be further optimized using efficient layout techniques. For example, ML transistors (M1) of the two adjacent cells can share the same drain contact (Fig. 9). Such a layout results in a smaller ML capacitance. In the present chip, we laid out the 144-bit cell in an array of 12×12 (Fig. 5). The extracted interconnect capacitance for ML is found to be 145 fF. When we laid out each ML as one row, the interconnect capacitance is reduced to 52 fF. Fig. 9 also shows an improved layout of the conventional comparison logic where the contacts are removed from the internal nodes in order to reduce their capacitance. Using the above layout techniques and minimum size transistors ($W = 0.42 \mu m$) in 0.18- μm CMOS technology, capacitances of 144-bit ML_Old and ML_New can be calculated as

$$C_{ML_Old} = 144 \times (0.825 \text{ fF}) + 52 \text{ fF} = 170.8 \text{ fF}$$

$$C_{ML_New} = 144 \times (0.309 \text{ fF}) + 52 \text{ fF} = 96.5 \text{ fF}$$

Thus, the proposed comparison logic with the modified layout achieves a 44% reduction in ML capacitance.

The energy reduction in charge-shared ML scheme varies with the ratio of ML1 and ML2 capacitances. If ML1 and ML2 capacitances are C_1 and C_2 , and the ML voltage swing is V_{ML} ,

the total charge consumed in the worst-case ML sensing for a conventional ML:

$$Q_{\text{OLD}} = (C_1 + C_2) V_{\text{ML}}. \quad (3)$$

Similarly, the total charge consumed in the worst-case ML sensing for a charge-shared ML:

$$Q_{\text{NEW}} = C_1 V_{\text{ML}} + C_2 (V_{\text{ML}} - V_{\text{CS}}) \quad (4)$$

where V_{CS} is the common-voltage of ML1 and ML2 after charge-sharing. V_{CS} can be calculated by applying charge-conservation before and after the charge-sharing:

$$V_{\text{CS}} = \frac{C_1 V_{\text{ML}}}{(C_1 + C_2)}. \quad (5)$$

Substituting V_{CS} from (5) and Q_{OLD} from (3) to (4)

$$Q_{\text{NEW}} = Q_{\text{OLD}} - \frac{C_1 C_2 V_{\text{ML}}}{(C_1 + C_2)}. \quad (6)$$

Therefore, the charged-shared ML achieves a worst-case energy (or charge) reduction of

$$\begin{aligned} E_{\text{Red}} &= \frac{(Q_{\text{OLD}} - Q_{\text{NEW}})}{Q_{\text{OLD}}} = \frac{C_1 C_2}{(C_1 + C_2)^2} \\ &= \frac{\left(\frac{C_1}{C_2}\right)}{\left(1 + \frac{C_1}{C_2}\right)^2}. \end{aligned} \quad (7)$$

Fig. 10 shows a plot of E_{Red} for different values of C_1/C_2 . It reaches a maximum of 25% for $C_1 = C_2$. Therefore, the charge-shared ML scheme is more suitable for TCAMs that have ML1 comparable to ML2. Substituting $C_1/C_2 = 36/108 = 1/3$ in (7), $E_{\text{Red}} = 19\%$. The measured reduction in energy (9%) is less than the theoretical value calculated from (7). There are two possible reasons for this difference. First, the charge-sharing time window (T_{CS}), which is fixed and equal to the inverter-chain delay, is not wide enough to fully equalize ML1 and ML2 (Fig. 4). Second, in the present implementation, a current source (I_{BIAS}) charges ML1 during the charge-sharing time window (Fig. 4). Thus, ML1 voltage is always slightly higher than ML2 voltage. The charge-sharing time window can be optimized by using a digitally controlled delay between DMLSO1 and MLOFF1 [15]. The second issue can be eliminated by using the rising edge of MLSO1 to turn off the corresponding I_{BIAS} during T_{CS} (Fig. 4).

In order to compare our results with the existing designs, we surveyed the published literature. The only published TCAM design with current-race MLSA and chip measurement results is found in [13]. This 144-bit TCAM, also implemented in 0.18- μm CMOS technology, achieves a search time of 3 ns for $I_{\text{BIAS}} = 260 \mu\text{A}$ [13]. Our charge-shared ML scheme achieves a search time of 4.7 ns for $I_{\text{BIAS}} = 120 \mu\text{A}$ (Fig. 7). Extrapolating the above results, our scheme shows 27% improvement in speed for the same I_{BIAS} .




Fig. 10. Worst-case energy reduction in charge-shared MLs for different values of the ML1 and ML2 capacitance ratio.

VI. CONCLUSION

We presented a low-capacitance ML scheme and a charge-shared ML scheme for TCAMs. Both schemes reduce search time and energy, and they can be combined with the existing design techniques to achieve a significant reduction in search time and energy. They can also reduce power if the search speed is not changed. The former scheme shows 42% and 25% reduction in search time and energy respectively. The latter scheme shows 11% and 9% reduction in search time and energy respectively. We analyzed the measurement results and proposed possible improvements. The low-capacitance ML scheme can be further improved by the efficient layout techniques described in Section V. Similarly, the charge-shared ML scheme can achieve better results for MLs with larger first segment, digitally controlled charge-sharing time window, and slightly modified MLSAs as discussed in Section V.

ACKNOWLEDGMENT

The authors would like to thank Canadian Microelectronics Corporation (CMC) for fabricating the test chips. They are grateful to W. Fung and D. Wright for stimulating discussions and feedback on this work, and D. Rennie and A. Pavlov for helping with the design and measurements of the test chips.

REFERENCES

- [1] H. J. Chao, "Next generation routers," in *Proc. IEEE*, Sep. 2002, vol. 90, no. 9, pp. 1518–1558.
- [2] M. Kobayashi, T. Murase, and A. Kuriyama, "A longest prefix match search engine for multi-gigabit IP processing," in *Proc. IEEE Int. Conf. Communications*, 2000, vol. 3, pp. 1360–1364.
- [3] N.-F. Huang, K.-B. Chen, and W.-E. Chen, "Fast and scalable multi-TCAM classification engine for wide policy table lookup," in *Proc. 19th Int. Conf. Advanced Information Networking and Applications (AINA'05)*, 2005, vol. 1, pp. 792–797.
- [4] H. Miyatake, M. Tanaka, and Y. Mori, "A design for high-speed low-power CMOS fully parallel content-addressable memory macros," *IEEE J. Solid-State Circuits*, vol. 36, no. 6, pp. 956–968, Jun. 2001.

- [5] D. Shah and P. Gupta, "Fast updating algorithms for TCAMs," *IEEE Micro*, vol. 21, no. 1, pp. 36–47, Jan./Feb. 2001.
- [6] Z. Wang, H. Che, M. Kumar, and S. K. Das, "CoPTUA: Consistent policy table update algorithm for TCAM without locking," *IEEE Trans. Computers*, vol. 53, no. 12, pp. 1602–1614, Dec. 2004.
- [7] M. Miller and B. Tezcan, Investigating design criteria for searching databases Integrated Device Technology Inc., White Paper, Feb. 2005 [Online]. Available: http://www.idt.com/content/solutions_InvestigatingDesignCriteriaForSearchingDatabases_wp.pdf
- [8] K. Lakshminarayanan, A. Rangarajan, and S. Venkatachary, "Algorithms for advanced packet classification with ternary CAMs," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 4, pp. 193–204, Oct. 2005.
- [9] P. Vlasenko and D. Perry, "Matchline sensing for content addressable memories," U.S. Patent 6,717,876, Apr. 6, 2004.
- [10] C. A. Zukowski and S.-Y. Wang, "Use of selective precharge for low-power content-addressable memories," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, 1997, pp. 1788–1791.
- [11] J.-K. Kim, P. Vlasenko, D. Perry, and P. B. Gillingham, "Low power content addressable memory architecture," U.S. Patent 6,584,003, Jun. 24, 2003.
- [12] K. Pagiamtzis and A. Sheikholeslami, "A low-power content-addressable memory (CAM) using pipelined hierarchical search scheme," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1512–1519, Sep. 2004.
- [13] I. Arsovski, T. Chandler, and A. Sheikholeslami, "A ternary content-addressable memory (TCAM) based on 4T static storage and including a current-race sensing scheme," *IEEE J. Solid-State Circuits*, vol. 38, no. 1, pp. 155–158, Jan. 2003.
- [14] N. Mohan and M. Sachdev, "A static power reduction technique for ternary content addressable memories," in *Proc. IEEE Canadian Conf. Electrical and Computer Engineering (CCECE)*, May 2004, pp. 711–714.
- [15] M. Maymandi-Nejad and M. Sachdev, "A digitally programmable delay element: Design and analysis," *IEEE Trans. VLSI Syst.*, vol. 11, no. 5, pp. 871–878, Oct. 2003.




Nitin Mohan (S'01–M'07) received the B.Tech. degree (with honors) in electronics engineering from the Institute of Technology, Banaras Hindu University (IT-BHU), Varanasi, India, in 1999, and the M.A.Sc. and Ph.D. degrees in electrical and computer engineering from the University of Waterloo, Waterloo, ON, Canada, in 2001 and 2006, respectively.

From 1999 to 2000, he was with Wipro Technologies working on the design and verification of FPGAs. During 2001–2002, he was with Sirific

Wireless Corporation designing CMOS integrated circuits. He spent the summer of 2005 at DALSA Corporation designing analog/mixed signal circuits. He is currently with Advanced Micro Devices Inc. in Boston, MA, where he designs high-performance low-power circuits for next-generation microprocessors. He has authored over a dozen journal/conference papers and a pending U.S. patent.

Dr. Mohan has received the Natural Sciences and Engineering Research Council of Canada Postgraduate Scholarship, Ontario Graduate Scholarship, President's Graduate Scholarship, and Doctoral Thesis Completion Award.



Manoj Sachdev (SM'97) received the B.E. degree (with honors) in electronics and communication engineering from the University of Roorkee, India, and the Ph.D. degree from Brunel University, U.K.

He was with Semiconductor Complex Limited, Chandigarh, India, from 1984 to 1989, where he designed CMOS integrated circuits. From 1989 to 1992, he worked in the ASIC division of SGS-Thomson, Agrate, Milan, Italy. In 1992, he joined Philips Research Laboratories, Eindhoven, The Netherlands, where he researched various

aspects of VLSI testing and design for manufacturability. He is currently a Professor in the Department of Electrical and Computer Engineering, University of Waterloo, ON, Canada. His research interests include low-power and high-performance digital circuit design, mixed-signal circuit design, and test and manufacturing issues of integrated circuits. He has written three books and two book chapters and has contributed to over 125 technical articles in conferences and journals. He holds more than 15 granted and several pending U.S. patents in VLSI circuit design and test.

Dr. Sachdev has received several awards, including the 1997 European Design and Test Conference Best Paper Award, the 1998 International Test Conference Honorable Mention Award, and 2004 VLSI Test Symposium Best Panel Award.