A Family of Cells to Reduce the Soft-Error-Rate in Ternary-CAM*

Navid Azizi and Farid N. Najm
Department of ECE, University of Toronto
Toronto, Ontario, Canada
(nazizi,najm)@eecg.utoronto.ca

ABSTRACT

Modern integrated circuits require careful attention to the soft-error rate (SER) resulting from bit upsets, which are normally caused by alpha particle or neutron hits. These events, also referred to as single-event upsets (SEUs), will become more problematic in future technologies. This paper presents a ternary content-addressable memory (CAM) design with high immunity to SEU. Conventionally, error-correcting codes (ECC) have been used in SRAMs to address this issue, but these techniques are not immediately applicable to CAMs because they depend on processing the full contents of the memory word outside the array, which is not possible in a normal CAM access. We propose a family of TCAM cells that reduce the SER at the cost of some area increase. An SER reduction of up to 40% can be obtained with a 18% increase of area; another design reduces the SER by 16% with only a 5% increase in area.

Categories and Subject Descriptors B.3.1 [Memory Structures]: Semiconductor Memories;

General Terms: Design

Keywords: Content-Addressable Memory, Soft-Error Rate.

1. INTRODUCTION

CMOS scaling has been driven by the desire for higher transistor densities and faster devices. Successive technology generations have shrunk the transistor dimensions and reduced the operating voltages of Integrated Circuits (ICs), which in turn have increased the sensitivity of ICs to the surrounding electromagnetic radiation [1]. A radiation event, such as an alpha particle hitting a semiconductor device, causes charge to be collected by the source and drain diodes of nearby transistors [1]; if the radiation event hits a memory element, the collected charge may be large enough that the contents of the storage element are flipped thus causing a Single-Event-Upset (SEU) [1]. These errors are designated *soft-errors* since they do not cause a permanent failure of the circuit [2] and the frequency of soft-errors for a device is designated the Soft-Error-Rate (SER).

Static Random Access Memory (SRAM) cells are usually used to implement memories that require fast access times and low power dissipation. Historically, SRAM cells were robust against soft-errors because of the inherent feedback of the cell, but as the SRAM cell has been made smaller and with the exponential increase in the amount of SRAM on chip, the SER in memories has increased with each process generation [1].

Content-Addressable Memories (CAM) are SRAM memories enhanced with comparison transistors that enable searching a word across all memory contents in a single clock cycle [3]. A CAM returns the location of the input word, effectively performing a table lookup operation, which speeds up a variety of lookup-intensive applications, but the most pervasive use of CAM today is in routers for the purposes of packet forwarding and classification [4].

In an SRAM memory, the SER is normally reduced to an acceptable level by using Error-Correcting Codes (ECC). ECC techniques are not immediately applicable to CAMs because they typically depend on processing the full contents of the memory word outside the array, which is not possible in a normal CAM access, which only returns match or miss results. Most techniques for avoiding errors in CAMs use Dynamic Random Access Memory (DRAM) cells, which have high soft-error immunity due to their large capacitor at the storage node, instead of SRAM cells [5]. One technique implements a DRAM block alongside an SRAM-based CAM [6]; the DRAM block, which includes ECC circuitry, is used to continuously write correct data into the CAM. Thus, in the worst case, any soft error in the CAM is overwritten in the amount of time it takes to refresh all entries in the CAM. One problem with this approach is that SEUs that happen in a CAM word and from which a match or miss result is obtained before correct data is overwritten from the DRAM array can lead to incorrect operation. Furthermore the additional DRAM block has a high (∼30%) area overhead. Using DRAM cells also results in increased design complexity and fabrication costs. We have developed two novel families of SRAM-based CAM cells that reduce the SER in CAMs. The techniques augment Ternary Content-Addressable Memory (TCAM) cells with extra transistors to make them more immune to SEUs. The different cells

*This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

DAC 2006, July 24–28, 2006, San Francisco, California, USA.

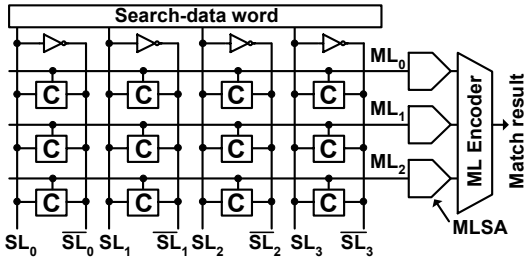

Figure 1: Organization of a Content Addressable Memory

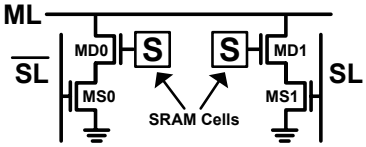

Figure 2: Ternary CAM cell

have different SER/area characteristics and allow for area to be traded off for a reduction in the SER.

The remainder of this paper proceeds as follows. Section 2 provides background on soft-errors, their relevant metrics and general SER reduction techniques as well as reviewing CAM basics. Section 3 describes the proposed new family of ternary CAM cells that reduces the SER. Section 4 presents the simulation results for the new cells, and finally, Section 5 concludes the paper.

2. BACKGROUND

2.1 Content Addressable Memories

Content-Addressable Memories (CAM) are memories which allow searching of an input word across all the memory contents in a single clock cycle. A CAM returns the locations of the stored data that match the input data, effectively performing a table lookup operation [3]. The organization of a CAM is shown in Fig. 1. A CAM search operation consists of three phases: data broadcast, word comparison and ML encoding [3].

The data broadcast phase is accomplished by driving the input search-data word onto the complementary Search-lines (SLs) labeled SL_i, and $\overline{SL_i}$ in Fig. 1. In the word comparison phase, the broadcast search-data word is compared to each stored word in parallel and the results appear on the Match-lines (MLs) labeled ML_i. During this comparison operation, each CAM cell compares its SL bit to its stored bit. If at least one cell in a word has a mismatch (or miss) between its SL bit and its stored bit, there will be a path from the ML to ground. On the other hand, if all cells in a stored word match the bits of the search word, there is no path from the ML to ground. To complete the word comparison phase, the Match-line Sense-amplifiers (MLSAs), which are connected to each ML, detect the state of their match-line (match or miss) and output a logic high for a match and a logic low for a miss. Finally, during the ML encoding phases, the ML encoder maps the MLSA outputs to a binary-encoded match result.

TCAM cells, as shown in Fig. 2, are composed of two SRAM cells which are augmented with additional compare circuitry, composed of two pull-down paths [3]. A TCAM cell, in addition to being able to store a logic-0 or a logic-1, allows for the storage of don't cares (X), which act as wild cards and allow pattern matching between the search and stored data [3]. A stored X will match if the search data is either a logic-0 or a logic-1. Table 1 shows the encoding of the SRAM bits to store a logic-0, logic-1 and don't care value in the TCAM cell. The match circuitry within each TCAM cell works as follows: if the TCAM cell is storing a logic-0, transistor MD0 will be turned off and transistor MD1 will be turned on; if then the search data that is presented to the TCAM cell is also a logic-0, SL will be low, turning transistor MS1 off, and $\overline{SL}$ will be high turning transistor MS0 on. In this case, where the stored and search data are equal, both of the pull-down paths will have one of their transistors being off, and thus the ML is disconnected from ground signaling a *match*. If, on the other hand, the search data that is presented to the CAM cell was a logic-1, one of the pull-down paths would have both transistors on, thus creating a path from the ML to ground signaling a *miss*. When storing a don't care value transistors MD0 and MD1 are both off at the same time, allowing the TCAM cell to match both a logic-0 and a logic-1 regardless of the search data, since there is no path to ground possible.

Table 1: SRAM encoding in a TCAM

SRAM values	TCAM state
00	don't care (X)
01	logic-0
10	logic-1
11	invalid

2.2 Soft Errors

Soft errors, which are also called SEUs, are errors in ICs due to external radiation rather that a design or manufacturing defect [7]. There are three key mechanisms that cause soft errors in ICs [1]: alpha particles emitted by decaying radioactive impurities in the packaging and interconnect materials, atmospheric neutrons with energies below 1 MeV and atmospheric neutrons with energies above 1MeV [8].

At the onset of an ionizing radiation event, such as an alpha particle or neutron collision, a track of electron-hole pairs is generated in the path of the ion's passage [1]. If the path of electron-hole pairs is near the reverse-biased junctions at the drain and source of transistors, where there is an electric field, the electric field collects the carriers thus creating a current and voltage glitch at that node [1]. If the radiation occurs near the drain of an n-channel Metal-Oxide Semiconductor (NMOS) transistor, electrons will be collected by the electric field onto the node thus causing a downward glitch of the node [2]. On the other hand, if the radiation occurs near the drain of a p-channel Metal-Oxide Semiconductor (PMOS) transistor, the holes will be collected at the node causing an upward voltage glitch [2].

The amount of charge that is *collected* by a transistor drain after a radiation event depends on a many factors including the substrate structure, the device doping and biasing of circuit nodes among others [1]. The collection slope, Q_S, of a device is a measure of the charge collection efficiency of a device measured in fC, and it is heavily dependent on the process [7]. As technology has scaled down, Q_S has become smaller and Q_S for a PMOS is usually less than that of an NMOS [8].

The collection efficiency of a device, however, is not the only factor that determines if a soft-error occurs; the sensi-

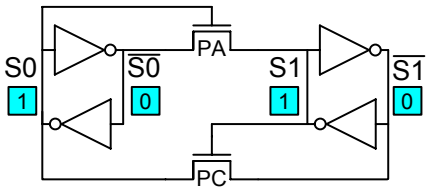

Figure 3: Simple Pass-Transistor Feedback-Enhanced TCAM

tivity of the device to the excess charge must also be considered. Q_{crit} is the critical amount of charge that has to be deposited at a circuit node for a soft-error to occur [7]. Q_{crit} depends on many factors including the nodal capacitance, the operating voltage, and the circuit topology [1].

Circuit simulators can be used to efficiently calculate the critical charge of a circuit [2] and, together with the use of empirical SER models, can be used to estimate the SER of circuits and ICs as a whole. One empirical model has the form [8]:

$$SER \propto F \times A \times e^{-\frac{Q_{crit}}{Q_S}} \quad (1)$$

where F is the neutron flux in particle/(cm^{2}s) and A is the area of the circuit that is sensitive to particle strikes in cm^2.

This empirical model can be used to determine the SER at a single node. The model has to be used twice, once for when the node is at logic-0 and once when the node is at logic-1 since Q_{crit}, Q_S and A will be different when considering the two states. When a logic-1 to logic-0 transition is being considered the Q_S of the NMOS is used, and the area of the NMOS drains that are attached to the node are considered, since an upset can only happen when electrons are injected onto the NMOS drain. Conversely when a logic-0 to logic-1 transition is considered the Q_S of the PMOS is used, and only the area of the PMOS drains are used since an upset can only happen when holes are injected onto the PMOS drain [2]. Since the collection efficiency of PMOS transistors is lower than that of an NMOS transistors and the hole mobility is less than the electron mobility, most soft errors are triggered by a radiation event on an NMOS transistor causing a logic-1 to logic-0 transition [9].

To determine the SER of the complete IC, the SER of every node is summed up to obtain the total SER [10]. The total SER is computed through a sum of the SER of each node instead of the average of the SER of each node because the area of each node is already factored into the nodal SER.

3. FEEDBACK-ENHANCED TCAM CELLS

We propose two new families of feedback-enhanced TCAM cells, built on the following strategy: select the state of the TCAM cell that is invalid (11), which is never used in a normally functioning CAM and make it unstable; the techniques used to make the invalid state unstable, will at the same time, reinforce the ability of the other states to hold their value and make them more stable. One of the families is built using pass-transistors, and the other family is built using cross-coupled transistors.

3.1 Pass Transistor CAM cells

In the pass-transistor CAM family, the decrease in stability of the invalid state is accomplished by connecting pass-transistors between the two SRAM cells in the TCAM cell.

Fig. 3 shows the circuit topology of a simple pass-transistor feedback-enhanced TCAM when storing the invalid state. For clarity, the pass-transistors that allow reading and writing into the cell, and the match circuitry, are not shown; only the two back-to-back inverters that comprise the storage of the cell are shown. Two additional NMOS transistors, PA and PC, have been connected between the storage nodes. In this invalid configuration, the addition of the two transistor connects nodes S1 and $\overline{\text{S0}}$ together and nodes S0 and $\overline{\text{S1}}$ together, which are at different logic values, and which cause the state of the TCAM to change from the invalid state to the don't care state.

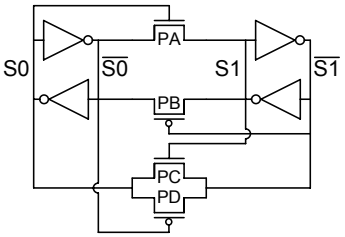

Figure 4: Full Pass-Transistor Feedback-Enhanced TCAM

When the TCAM shown in Fig. 3 is storing a logic-1 the nodes S1 and $\overline{\text{S0}}$ are connected through PA and are more immune to soft-errors because each node reinforces the other and allows for an extra path for the injected charge to be discharged; when the cell is storing a logic-0 nodes S0 and $\overline{\text{S1}}$ reinforce each other. On the other hand when the TCAM is storing a don't care value, both pass-transistors are off and there is no additional immunity to soft-errors.

This cell provides increased immunity against a logic-0 to logic-1 transition by providing extra discharge paths through NMOS pass-transistors. It also provides an increased immunity for a logic-1 to logic-0 transition by reducing the feedback that exists in the cross-couple inverter pair. For example, when a cell is storing a logic-1, a radiation event occurring on node S0 will reduce the voltage on that node, which in a typical TCAM, would have caused an increase in the voltage on node $\overline{\text{S0}}$ and caused the cell to flip; but in the new TCAM cell since node $\overline{\text{S0}}$ is also being driven by node S1, the voltage at node $\overline{\text{S0}}$ will not increase as much as before thus limiting the positive feedback that may flip the cell.

A further variation on the idea which more actively protects against logic-1 to logic-0 transitions is to include two PMOS pass-transistors that also connect the storage nodes as shown in Fig. 4. This "full pass-transistor feedback-enhanced TCAM cell" has extra immunity since two pass-transistors are ON when the cell is holding a logic-0 or logic-1.

3.2 Cross-Coupled CAM cells

A second family of feedback-enhanced TCAM cells has been developed that uses additional pull-down and pull-up transistors arranged in a cross-coupled fashion.

This family is also built on the strategy of decreasing the stability of the TCAM cell when it is in the invalid state, in order to increase the stability when the cell is holding the other states. Fig. 5 shows the circuit topology of a simple cross-coupled feedback-enhanced TCAM when storing the invalid state. Two additional NMOS transistors, DA and DC, have been arranged in a cross-coupled fashion between nodes S0 and S1. In this invalid configuration, the addition

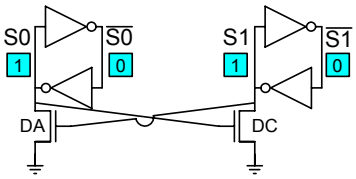

Figure 5: Simple Cross-Coupled Feedback-Enhanced TCAM

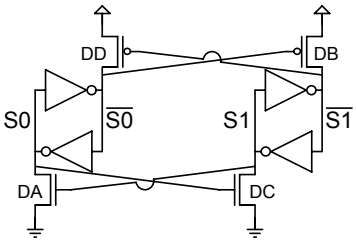

Figure 6: Full Cross-Coupled Feedback-Enhanced TCAM cell

of the two transistors causes both those nodes to be driven low, causing the TCAM cell to switch from the invalid state to the don't care state, thus making the invalid state unstable. Furthermore, the addition of these transistors increases the stability when the cell is holding a logic-0 or a logic-1; for example when the cell is holding a logic-0 transistor DA is on and it reinforces the low voltage at node S0.

The "full cross-coupled feedback-enhanced TCAM cell" is shown in Fig. 6; in this circuit two PMOS transistors have been added in a cross-coupled fashion between nodes $\overline{S0}$ and $\overline{S1}$ which further reduces the stability of the invalid state, and allows for extra discharge paths when the cell is holding one of the three other states thus increasing their stability. When the cell is holding a logic-0 the voltage on node S1 reinforces the low voltage at node S0 by turning on transistor DA, and the voltage at node $\overline{S1}$ reinforces the high voltage at node $\overline{S0}$ by turning on transistor DD; the converse happens when the cell is holding a logic-1. There is, however, no obvious increased protection when the cell is in the don't care state as all the extra transistors are off.

4. RESULTS

To consider the benefits of the two families of feedback-enhanced TCAM cells, sixteen different cells from each family were designed, which contained all the permutations of including or not including the four transistors from the full feedback-enhanced TCAM cells.

4.1 Assessment Methodology

All simulation results reported in this report are based on HSPICE, using Berekely Predictive Technology Models (BPTM) [11] for a 70nm technology.

As explained in Section 2.2 the SER of a circuit is the sum of the SER of its nodes, and thus the SER of a TCAM cell is the sum of the SER for each of the four storage nodes. Since a TCAM can hold one of three states, a logical-1, a logical-0 and a don't care value, the SER has to be calculated when the cell is storing each of three values, and averaged with the probability that the cell is holding that value. Without having more information about the distribution of the values stored in TCAM cells, it will be assumed that a TCAM cell is equally likely to hold each of the three values (33.33%).

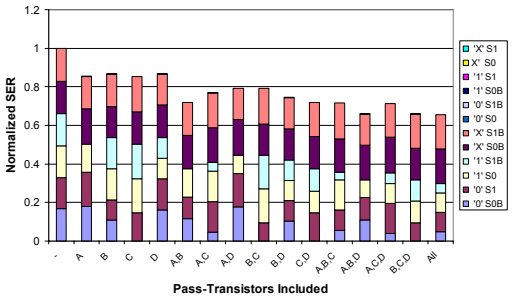

Figure 7: Reduction in the SER of the different pass-transistor feedback-enhanced TCAM cells broken up into different components

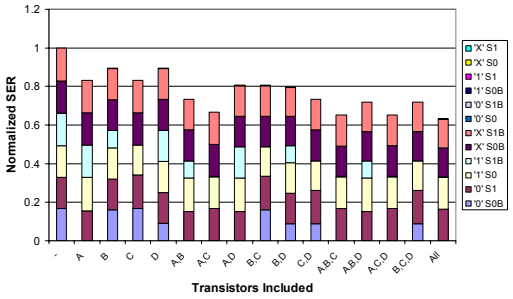

Figure 8: Reduction in the SER of the different Cross-Coupled feedback-enhanced TCAM cells broken up into different components

To differentiate between the SER of a node and the SER of a node when the TCAM is holding a particular state, the term *node-state* will be used. The node-state 0/S1 will refer to node S1 when the TCAM cell is holding a logic-0.

To determine the reduction of the SER of a single node-state, we use (1). The susceptible area, A, for the SRAM cell within the TCAM was measured off a layout [12], and for the new feedback-enhanced TCAM cells the increase in the value of A was projected from the same layout. Q_{Sp} and Q_{Sn} for a 70nm process was extrapolated from the trends found in [8] and the critical charge, Q_{crit} for each node, was found through an HSPICE simulation [7] by injecting charge with a piece-wise linear current waveform in the shape shown in [13]. For nodes which are at logic-0, a positive charge was injected due to radiation events on PMOS drains, and for nodes which are at logic-1 a negative charge was injected due to radiation events on NMOS drains. The charge was increased until the value that was stored in the TCAM was corrupted and that charge was defined to be Q_{crit}.

4.2 Soft-Error-Rate

The reduction in the SER for the pass-transistor family is shown in Fig. 7. It can be seen that the best permutation is when all four pass transistors are used, which reduces the SER by approximately 34.5%. The permutation that uses transistors PA, PB, and PD reduces the SER by almost the same amount, 34%, even though it uses one less transistor.

For the cross-coupled family of cells it can be seen that the best permutation is when all four transistors are used which reduces the SER by approximately 36.7% as seen in Fig. 8. The permutation that uses DA, and DC reduces the SER by slightly less, 33.1%, even though it uses two less transistors. The PMOS transistors are too weak to provide much benefit in this case.

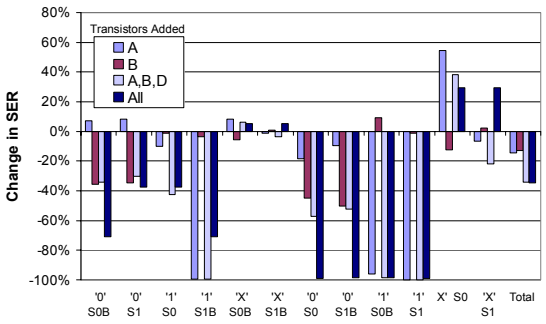

Figure 9: Breakdown in the reduction in the SER of various pass-transistor cells

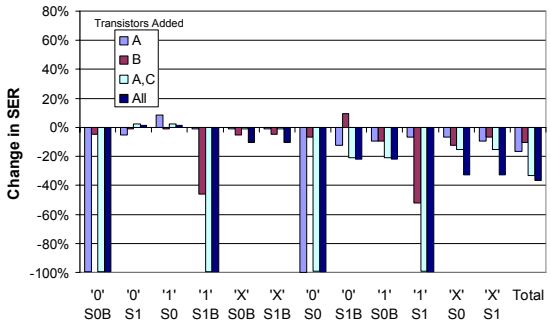

Figure 10: Breakdown in the reduction in the SER of various cross-coupled cells

Fig. 9 provides more detail as to why these reductions occur for the pass-transistor family. The figure shows the reduction in SER for each node-state for four different cells from the pass-transistor family. The first six groups of bars in the figure represent radiation events on NMOS transistors which cause logic-1 to logic-0 transitions, and the last six bars represent radiations events on PMOS transistors which cause logic-0 to logic-1 transitions. When only one NMOS transistor, transistor PA, is added it can be seen that when the TCAM cell is holding a logic-1 it has a high immunity to soft-errors; virtually eliminating any soft-errors that occur due to radiation events on node-states $1/\overline{S1}$, $1/\overline{S0}$ and 1/S1. In almost all other cases there is only a small decrease in the SER or even a small increase, which is due to the increase in the susceptible drain area with the inclusion of the additional transistor. There is, however, a large increase in the SER of the node-state X/S0; in this case the added pass-transistor is nominally off and separates two nodes which are at different logic values; a radiation event on node S0, which is connected to the gate of the pass-transistor, partially turns ON the pass-transistor making it easier to flip the cell. The increase in the SER of the node-state X/S0 is not really an issue because the SER of all the logic-0 to logic-1 transitions, due to PMOS transistors, are orders of magnitude lower that the SER due to NMOS transistors [9] as can be seen in Fig. 7 where the SER of the logic-0 to logic-1 transitions, which are at the top of the stacked bars, are virtually inconsequential to the cell SER.

When only transistor PB is used, while no node becomes invulnerable to soft-errors, the rate of most of the logic-1 to logic-0 transitions is lowered leading to a better reduction in the cell SER. When two PMOS and one NMOS transistor are used, it can be seen that when the TCAM is holding a logic-0 or logic-1 the SER of all nodes is considerably reduced. The TCAM cell is also virtually invulnerable at node-state $1/\overline{S1}$ due to the asymmetry in the connections within the cell. When all four pass-transistors are added, the SER at each node becomes more moderated.

For the cross-coupled family similar, trends can be seen in Fig. 10. When an NMOS pair is used four node-states become almost invulnerable and when all four transistors are added, the SER further reduces slightly at each node-state. A difference in the results compared to the pass-transistor family of TCAM cells is that the SER of the don't care states are reduced, instead of becoming higher since the added transistors do not have voltage differentials between their drain and source, as is the case in the pass-transistor family.

4.3 Performance

The read times of both families of enhanced-feedback TCAM cells are virtually unchanged compared to that of a conventional TCAM cell. The worst-case write times, however, show an increase, which at worst is 87ps (less than 6%). The increase in write times is due to the extra feedback in the TCAM cell which causes an increased contention when writing the cell. The increase in write times, while not ideal, is not of primal concern since the critical performance metric in a CAM is the match time.

4.4 Area

The additional transistors in the feedback-enhanced families of TCAM cells will increase the cell area. Each SRAM cell takes up 20% of the TCAM cell area [14]. Furthermore, in [15] an additional transistor was added to an SRAM cell and the area increased by 16.6%. These two numbers can be used to determine a projected area increase by adding up to four transistors in feedback-enhanced families of cells. For example, when adding one extra transistor the cell area would increase by 3.3%, but we choose a slightly more pessimistic value of 5%. When two transistors of the same type are added, the projected area increase is 10% but when two transistors of different types are used the projected area increase is less (8%) because the two transistors will expand the layout of the cell in the same direction thus limiting the cell area increase. When adding four additional transistors, the projected area increase is 16%.

The reduction in the SER is not due to the increase in cell area, but instead it is due to the new cell topologies introduced. Increasing the width of the transistors in the TCAM cell by-itself, while increasing Q_{crit}, would also increase the susceptible drain area and would actually increase the SER. In our simulations increasing the width of the pull-down transistors by 50% increases the SER by around 2%.

4.5 Static Noise Margin

While the SER of the new family of TCAM cells is reduced by up to around 35% resulting in increased stability, the worst-case Static Noise Margin (SNM) of the pass-transistor family of cells, which is a measure of tolerance to DC disturbances during a read, is decreased by up to 6%, and the SNM of symmetric cells in the cross-coupled family is decreased by up to 5%.

While this sounds unintuitive, the reason for it is that the SER is a measure of tolerance to transient noise sources, while the SNM is a measure of tolerance to DC noise sources.

When measuring the SNM, a DC value is swept at one storage node from 0V to V_{DD}, providing a butterfly curve and the largest square that can fit into the lobes of the curve is the SNM [16]. When the voltage is being swept to intermediate values, this DC measurement will cause all other nodes to settle at intermediate values, which would never occur in the transient case. Due to the extra connections within the new TCAM family, these intermediate values cause the shape of the butterfly curve to be distorted compared to the conventional curves, thus reducing the SNM.

To recover the SNM and make it equal to that of a conventional cell the pull-down NMOS transistors in the core of the two SRAM cells can be made slightly wider. By making the pull-down transistors wider by 8.5% the SNM of the useful variations of the pass-transistor and cross-coupled family of TCAM cells can be made virtually equal to that of a conventional TCAM cell. There is, however, an increase in the area of the TCAM cell by less than 1%. Furthermore the SER of these cells is virtually unchanged compared to the cells whose pull-down transistors are not made wider.

4.6 Larger PMOS transistors

In designing the above two families of SER tolerant CAM cells, the added transistors were of minimum size, but non-minimum size transistors can also be used. While increasing the width of transistors usually increases the susceptible area for soft-errors to occur, the addition of PMOS transistors has almost no effect on the SER since the logic-0 to logic-1 transition rate is much lower than the logic-1 to logic-0 transition [9]. With this in mind, the PMOS transistors in both the pass-transistor and cross-coupled family of TCAM cells were doubled in size to investigate the effect of increased PMOS width.

In the pass-transistor family, the increase in the PMOS width helps to further reduce the SER of all designs that have PMOS transistors in them by around 2% to 6.5%. An interesting observation is that the design that has transistors PA, PB, and PD is now the best design, reducing the SER by more than 40%. For the cross-coupled family, the larger PMOS transistors have a smaller effect on the best designs, where the complete cross-coupled design SER reduces by only a further 1.7% to a 38.4% reduction.

4.7 Summary

Fig. 11 shows the reduction in the SER for all designs plotted against the projected increase in the cell area for those designs. If an increase in cell area of around 17% is tolerable, then the SER can be reduced by 40.6% by using the PMOS transistors and an NMOS transistor from the pass-transistor family. If on the other hand, only a 10% cell area increase is tolerable, using the design from the cross-coupled family which only has two NMOS transistors is the best, and it reduces the SER by 33%.

5. CONCLUSION

Two novel families of SER tolerant feedback-enhanced TCAM cells have been presented. The different cells have different SER/area characteristics and allow for area to be traded-off for a reduction in the SER. The best design reduces the SER by 40% with a 18% increase of area; another design reduces the SER by 16% with only a 5% increase in area.

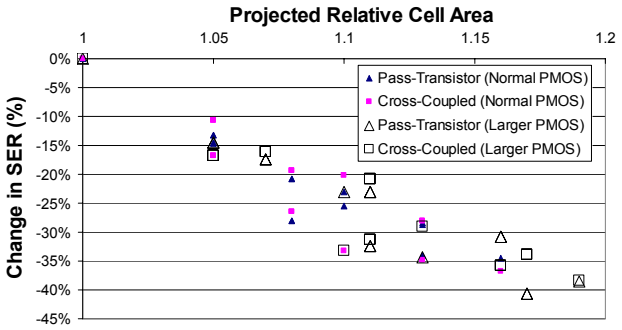

Figure 11: Summary of tradeoff between cell-area and SER

6. REFERENCES

[1] R. Baumann. Soft errors in advanced computer systems. *Design & Test of Computers*, 22(3):258–266, May-June 2005.

[2] T. Heijmen. Analytical semi-empirical model for SER sensitivity estimation of deep-submicron CMOS circuits. *Inter. On-Line Testing Symp.*, pages 3–8, July 2005.

[3] Kostas Pagiamtzis and Ali Sheikholeslami. Content-addressable memory (CAM) circuits and architectures: A tutorial and survey. *IEEE Journal of Solid-State Circuits*, 41(3):712–727, March 2006.

[4] H.J. Chao. Next generation routers. *Proceedings of the IEEE*, 90(9):1518–1558, Sept. 2002.

[5] H. Noda, K. Inoue, M. Kuroiwa, et al. A cost-efficient high-performance dynamic TCAM with pipelined hierarchical search and shift redudancy architecture. *Journal of Solid State Circuits*, 40(1):245–253, Jan. 2005.

[6] H. Noda, K. Dosaka, F. Morishita, K. Arimoto, et al. A soft-error immune maintenance-free TCAM architecture with associated embedded DRAM. *Custom Integrated Circuits Conf.*, pages 451–454, Sept. 2005.

[7] P. Shivakumar, M. Kistler, S.W. Keckler, D. Burger, and L. Alvisi. Modeling the effect of technology trends on the soft error rate of combinational logic. *Inter. Conf. on Dependable Systems and Networks*, pages 389–398, 2002.

[8] P. Hazucha and C. Svensson. Impact of CMOS technology scaling on the atmospheric neutron soft error rate. *Trans. on Nuclear Science*, 47(6):2586–2594, Dec. 2000.

[9] T. Juhnke and H. Klar. Calculation of the soft error rate of submicron CMOS logic circuits. *Journal of Solid-State Circuits*, 30(7):830–834, July 1995.

[10] N. Seifer, D. Moyer, N. Leland, and R. Hokinson. Historical trend in alpha-particle induced soft error rates of the Alpha™microprocessor. *Inter. Reliability Physics Symp.*, pages 259–265, April 2001.

[11] http://www-device.eecs.berkeley.edu/~ptm/.

[12] R. Venkatraman et al. The design, analysis, and development of highly manufacturable 6-T SRAM bitcells for SoC applications. *Transactions on Electron Devices*, 52(2):218–226, Feb. 2005.

[13] L.B. Freeman. Critical charge calculations for a bipolar SRAM array. *IBM Journal of Research and Development*, 40(7):119–129, Jan. 1996.

[14] P.F. Lin and J.B. Kuo. A 0.8-V 128-kb four-wayset-associative two-level CMOS cache memory using two-stage wordline/bitline-oriented tag-compare (WLOTC/BLOTC) scheme. *Journal of Solid-State Circuits*, 37(10):1307–1317, Oct. 2002.

[15] N. Azizi and F.N. Najm. An asymmetric SRAM cell to lower gate leakage. *Inter. Symp. on Quality Electronic Design*, pages 534–539, 2004.

[16] E. Seevinck Sr., F. J. List, and J. Lohstroh. Static-noise margin analysis of MOS SRAM cells. *Journal of Solid-State Circuits*, 22:748–754, Oct. 1987.