


Article

Multi- V_{dd} Design for Content Addressable Memories (CAM): A Power-Delay Optimization Analysis

Siddhartha Joshi ^{1,*} , Dawei Li ¹, Seda Ogrenci-Memik ¹, Grzegorz Deptuch ², James Hoff ², Sergio Jindariani ², Tiehui Liu ², Jamieson Olsen ² and Nhan Tran ²

¹ Electrical Engineering & Computer Science, Northwestern University, Evanston, IL 60201, USA; daweli2016@u.northwestern.edu (D.L.); seda@eecs.northwestern.edu (S.O.-M.)

² Fermi National Accelerator Laboratory, Batavia, IL 60208, USA; deptuch@fnal.gov (G.D.); jimhoff@fnal.gov (J.H.); sergio@fnal.gov (S.J.); thliu@fnal.gov (T.L.); jamieson@fnal.gov (J.O.); ntran@fnal.gov (N.T.)

* Correspondence: sidjos@u.northwestern.edu; Tel.: +1-630-815-3608

Received: 15 June 2018; Accepted: 27 July 2018; Published: 30 July 2018



Abstract: In this paper, we characterize the interplay between power consumption and performance of a matchline-based Content Addressable Memory and then propose the use of a multi- V_{dd} design to save power and increase post-fabrication tunability. Exploration of the power consumption behavior of a CAM chip shows the drastically different behavior among the components and suggests the use of different and independent power supplies. The complete design, simulation and testing of a multi- V_{dd} CAM chip along with an exploration of the multi- V_{dd} design space are presented. Our analysis has been applied to simulated models on two different technology nodes (130 nm and 45 nm), followed by experiments on a 246-kb test chip fabricated in 130 nm Global Foundries Low Power CMOS technology. The proposed design, operating at an optimal operating point in a triple- V_{dd} configuration, increases the power-delay operation range by 2.4 times and consumes 25.3% less dynamic power when compared to a conventional single- V_{dd} design operating over the same voltage range with equivalent noise margin. Our multi- V_{dd} design also helps save 51.3% standby power. Measurement results from the test chip combined with the simulation analysis at the two nodes validate our thesis.

Keywords: content addressable memory (CAM); TCAM; multi- V_{dd} ; multi supply; associative memory; tunable operation; standby power; searchline power; matchline power

1. Introduction

Content Addressable Memories (CAM) operate by comparing data in parallel, which makes the search operation extremely fast; however, it also makes the CAMs power hungry. High power consumption is the primary design issue with CAM design, and there have been various attempts towards improving the power efficiency of CAMs. We discuss related techniques in this avenue in Section 3. In this paper, we analyze and model the power consumption behavior of content addressable memories and then propose a multi- V_{dd} design for content addressable memories to reduce power consumption and achieve a highly tunable design. We evaluate our design, present simulation analysis on two nodes (130 nm and 45 nm) and finally present experimental results from our fabricated test chip at 130 nm Global Foundries Low Power CMOS technology.

A post-fabrication tunable design, with variable search speeds operating over a large power-delay space, helps us adapt performance based on the workload, accommodate for process variations, and reduce design margin overheads. Our design does not require the use of any level converters thereby avoiding their power and area overhead. The chip design incurred a two-pin overhead for

the full chip along with design overhead on the test board to create and route the additional external supplies. However, these could also be traded-off with the power and area overhead of an internal on-chip voltage generator. The power saved and optimum operating points for the multi- V_{dd} design are workload dependent. For our workloads, which consisted of particle physics data, we could get a dynamic power saving of 25.3% and a standby power saving of 51.3%. Other power-saving techniques can be applied on top of our multi- V_{dd} design to yield even better results. For most designs and workloads, we find that reducing searchline swing by reducing the driver voltage, along with decreasing the matchline charging time, by increasing the voltage for the current driver to the matchline, together, to be the most effective way to optimize power for a given performance point.

The paper is organized as follow: First, we give a background on CAMs and the architecture of our CAM chip. Here, we also introduce our target application, which lies in high energy particle physics. Then, we summarize the related work in this field in Section 3. In Section 4, we explain our multi- V_{dd} CAM design and its operation, along with the power saving techniques we employed. We explore the power consumption behavior of matchline-based content addressable memories in Section 5. Then, in Section 6, we show the behavior of the multi- V_{dd} CAM chip through simulation analysis over a range of input voltages and explore its power-delay space, which indicates opportunities for optimization. In Section 7, we present testing results from the test chip to validate our analysis, and we present our conclusions in Section 8.

2. Background

A CAM is a memory device that performs two functions: (1) It stores data; and (2) It compares the stored data with input data, in parallel across the whole chip. Most CAM designs are able to search their entire stored data in a single clock cycle. Figure 1 shows the block diagram of our CAM chip. The main components of the CAM chip are: data and clock drivers, row and column address decoding blocks and the CAM core [1,2]. The CAM core consists of a 2D array of CAM cells, with 3D IC implementations in development [3–7]. CAM cells have three main types: NAND, NOR, and Ternary [2].




Figure 1. Block Diagram of the CAM Chip.

The primary applications of CAMs lie in Internet Protocol packet classification and forwarding [1,8]. They are also used in translation look-aside buffers, Huffman coding/decoding [2], and also for classifying experimental data in sciences, such as high energy particle physics. CAMs have also been found to be suited for accelerating data-intensive parallel workloads [9] including graphics [10–12] and search engines [13]. The target application for our CAM chip lies in high energy particle physics to filter data obtained from particle collisions in particle colliders. The filtering and detection of collision data needs to be done in real-time, and CAMs are the best hardware resource to do just this [3,6,14–18].

Data is loaded in the CAM cells by inserting the row and column addresses along with the data to be stored and a write signal. Data is most commonly stored in SRAMs, however, other technologies like RRAM and STTRAM are being explored [1,9,19]. This operation is usually not as timing or power critical as the search operation itself, as writes are infrequent in CAMs. So, the writes can be done at a lower clock frequency. Once the data is loaded in the CAM cells, we move to the search operation. In the search operation, in every clock cycle, data is sent to each of the CAM cells, where it is compared. Data is compared in parallel in each of the CAM cells; and inside each cell, all the bits are compared in parallel, too. This high level of parallelism helps make the comparison operation extremely fast. Post-processing of the outputs of CAMs, which is either a match or mismatch, is highly application-specific. Common post-processing blocks are priority encoders, whose output is based on a given priority order; or threshold logic comparators, whose output is based on the number of matches/ mismatches in a group of CAM cells. In our previous work [20], we presented our design for a multi- V_{dd} CAM and initial testing results with an emphasis on our tracking trigger target application. In this paper, we expand on that work, explore our design in much greater detail, and generalize the modeling of the power behavior.

3. Related Work

Approaches to reduce the power consumption in CAMs have targeted primarily two main components: matchline power and searchline power [1,2,21–32]. In most applications, a much larger number of mismatches occur than matches [32], and hence many techniques have been tailored to reduce the mismatch power consumption. The selective pre-charge technique [26] proposes the use of NAND type cells in the higher order bits of the matchline which stop current propagation at the mismatched bit when input data mismatches with the stored data in a CAM cell. Although this technique saves power, it also increases the impedance along the matchline and, therefore, slows the comparison operation. Hence, usually NOR type cells are preferred for workloads requiring higher speed. For NOR type cells, a mismatch-dependent power saving scheme utilizes a precharge controller which predicts the mismatched state on the matchlines [21] and avoids precharging them to their full high level in case of a mismatch which saves matchline power. Low swing searchlines [31] have been proposed to save the searchline power. Another technique utilizes low swing searchlines [22] on global searchlines with higher capacitive loads, and full swing on local searchlines which have lower capacitive loads. This technique can be combined with a pipelined search scheme [22] which activates fewer portions of the CAM initially, and then proceeds by activating only those portions which match. To address the precharging power consumption, a pre-charge free design was proposed [23]. A design, which reduces the leakage power in CAM's proposed two novel TCAM cells [24]. A charge-shared matchline [30] reduces the matchline power consumption by reducing the capacitance on the matchline. In the current-race scheme [25], the matchline is pre-discharged, as opposed to traditional precharging. Then, when new data is sent to the CAM cells for comparison, the matchline is charged using a current source. Since most workloads incur a higher number of mismatches than matches, this scheme saves power by not discharging the pre-charged matchline to the ground in most clock cycles. This scheme also has the added benefit of removing the searchline reset [2,25], thereby saving searchline power too. These aforementioned techniques can be applied along with a multi- V_{dd} scheme with an additional design effort and area overhead as and where it is deemed necessary. We utilize the selective precharge and current race scheme in our design, as we will elaborate in Section 4.

TCAMs have also been employed to help reduce power in other systems such as in the FPU inside GPUs [10], by replacing some of their operations. These TCAMs store the result of common operations, and then consequently the re-execution of the core is instead replaced by just a search and read inside the TCAMs, which saves power. TCAMs can also be used in applications which can tolerate approximate results like in multimedia applications running on GPUs, which leverage the acceptability of a lower quality of service through voltage over-scaling [10–12] to get a much lower power consumption. These techniques benefit especially from the use of non-volatile memories which

can provide higher densities, and lower power consumption under certain conditions. However, most applications, including ours, cannot tolerate approximate results, and the integration of non-volatile memory is not feasible for most applications.

The multi- V_{dd} technique is used in CMOS VLSI designs [33–36], including SRAMs [37–40], FPGAs [41–46], and CAMs [28,29,47] to reduce power and improve reliability. It has been shown to have an outsized effect on VLSI designs when compared to transistor sizing and the use of dual threshold transistors [48]. Voltage scaling is highly effective as power reduces quadratically; however, it also impacts the delay [49,50]. The idea for using multi- V_{dd} in systems usually encompasses using the extra timing slack or running paths off the critical path at a lower supply [34]. The multi- V_{dd} approach is useful even when wire capacitance is dominant over gate capacitance, which is not true of most other power saving techniques. Multi- V_{dd} designs have the added benefit of allowing much larger post-fabrication tuning, which not only helps in reducing power consumption but can also increase the yield [51]. Multi- V_{dd} has been proposed for CAM designs; however, most studies only focus on the static power saving [28] or only on a particular design [29]. One study suggests the use of higher V_{dd} for the priority encoder [47], however, this restricts it to CAM devices where the priority encoder delay is significant and part of the critical path.

We present a general study where we separate the post-processing of CAM match/mismatch outputs from the CAM delay and operation. Our study also involves testing of a fabricated chip, which is paramount in reaching convincing conclusions on the multi- V_{dd} feasibility. We have used some of the existing power saving techniques that are complementary to the study of multi- V_{dd} operation space. However, there are other compatible techniques that can be used along with our multi- V_{dd} CAM design to further improve results depending on the workload. We present a general study to correctly evaluate the benefits of a multi- V_{dd} design in CAMs starting from the basic principles of power consumption behavior of the components.

4. Multi- V_{dd} CAM Design

To evaluate multi- V_{dd} design in CAMs, we use a traditional design incorporating both binary and ternary unit cells. We employ a matchline-based architecture with a selective precharge scheme [26] along with a current race scheme [25]. Our total input word size is 60 bits, divided into 15-bit segments (Figure 2). The outputs of 4 CAM cells are AND-ed or OR-ed in a threshold control logic cell to form an effective 60-bit CAM cell, which gives the final output. Figure 3 shows the layout of the test chip used in our evaluation. It is 5.46 mm \times 5.46 mm, and each 15-bit CAM cell macro is 25 μ m \times 25 μ m at the 130 nm node. Figure 4 shows a photograph of our fabricated test chip. Table 1 summarizes the three power supplies we use in our design.

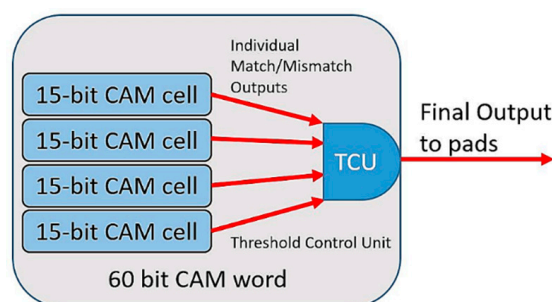


Figure 2. CAM Word block diagram.




Figure 3. Chip layout.


Figure 4. Photograph of the multi- V_{dd} CAM chip.

Table 1. Power supplies used and supplied components.

Supply	Components Supplied
VCHARGE	Matchline
CVDD	SRAM, sense-amplifier (SR-FF)
DVDD (Driver-Vdd)	Data (Senseline) and Clock (CLK) Drivers

Our CAM cell design is a mixed-type CAM with both Binary and Ternary Cells. Ternary cells require a larger area, and hence were used only where required. We employed a combination of NAND, NOR, and Ternary type cells, in order to satisfy power-delay requirements and also to keep our analysis general. Figure 5 shows the schematic of our CAM cell. Each CAM cell is 15 bits where bits 14–11 are NAND type cells which form the selective precharge; bits 9, 8, 7, 6, 4, 3, 2, and 1 are NOR type cells; and finally bits 10, 5, and 0 are NOR-type-Ternary cells [2] for wildcard inputs. The input and clock drivers are laid out along the periphery with the CAM core in the center. Altogether our CAM core forms a $2D$ array of 128×128 15-bit CAM cells.

In our design, we leverage the selective precharge technique [26] and the current race scheme [25]; both of which we introduced in the previous section. We use four NAND type cells in the higher order bits of the matchline to stop further propagation of the matchline current in case of a mismatch thereby saving matchline power. The current-race scheme [25], pre-discharges the matchline, as opposed to traditional precharging. Then, when new data is sent to the CAM cells for comparison, the matchline is charged using the VCHARGE supply. In our workload, we expect a much larger number of mismatches than matches, and hence, this scheme saves power by not discharging the pre-charged matchline to the ground in most clock cycles.

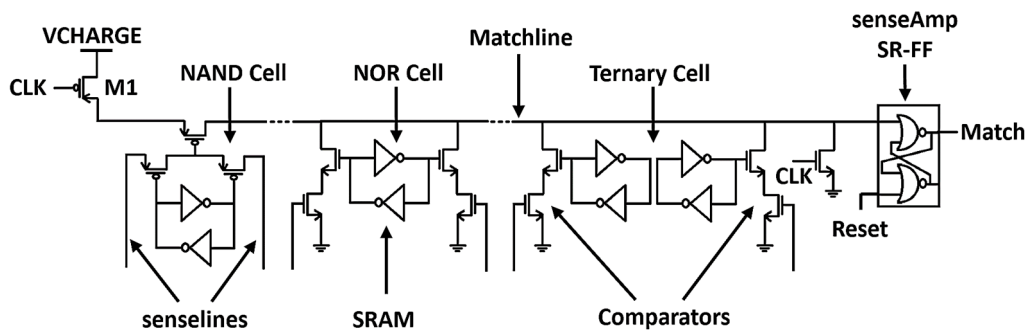


Figure 5. Multi- V_{dd} CAM Cell.

CAM operation begins by loading data into the SRAM cells of the CAMs after which the search cycle commences. The search clock-cycle can be divided into two parts. First, when the clock signal is high, the previous charge on the matchline is discharged to ground. Also, during this time, new data is sent to the CAM cells through the searchlines. Then, when the clock signal drops to 0 the search period begins and the matchline starts charging from the VCHARGE node. If there is no mismatch between the input data and the stored data, the matchline charges and once it reaches the designed threshold voltage of the SR-FF sense amplifier, the state of the CAM switches to a match. If there is a mismatch between the input data and the stored data, depending on the location(s) of the mismatched bit, the matchline is connected to the ground by one of the NOR or Ternary cells or the charging of the matchline is blocked by one of the NAND cells, higher up the bits in the matchline. This stops the matchline from charging up to the threshold voltage of the senseAmp and so the CAM stays in the mismatched state. These two periods, the pre-discharge and the compare/charge periods, constitute one complete clock cycle in the CAM operation. The state of the CAM is stored until a reset signal is sent. At the end of the workload input data stream, the output data of the CAM is read out, and the state of the CAM is reset to mismatched. We store the output in the SR-FF to be read-out/processed separately in a different phase. By storing the output, we can separate the delay and power behavior of the post-processing circuitry (threshold logic cell and output drivers), which can be highly application-specific. For our design, we used a simple threshold comparator, however, as explained, its delay and power do not influence the operation of the main CAM core.

We employed a multi- V_{dd} design where we separated the voltage domains based on the different power consumption behaviors of the functional blocks of a CAM chip (Table 1). We detail the power consumption behaviors of the functional blocks in Section 5. The functional blocks have been divided into three voltage domains, each controlled by an independent voltage source. The data input and clock drivers are supplied by the Driver-VDD (DVDD) supply, the matchlines inside the CAM cells are supplied by the VCHARGE supply, and finally, the SRAM storage cells and the matchline sense-amplifiers in the CAM cells are supplied by the CVDD supply. The availability of separate supplies not only allowed us to save power and optimize the operation of the CAM chip based on the workload, but it also allowed us to study the power consumption behavior of CAMs in detail.

The primary issues with multi- V_{dd} designs are the area overhead due to level converters and static power loss at the boundary of voltage domains [52]. In our design, for optimal operation, care is taken such that there is no path where a component operating at a lower supply domain will drive a component operating at a higher supply voltage [34]. This meant that we did not need any level converters in our design, which helped us avoid the area and power overhead. This was done by having the different voltage domains connect inside the CAM cells, connected by the matchline through the SR-FF senseAmp. The matchline signifies the match/mismatch state of the CAM cell, and it is not a digital signal. Its state, defined by its voltage level, is detected by the SR-FF, which can be designed to operate over a larger voltage range than the traditional digital voltage abstraction. Another boundary between the supply domains occurs inside the comparators of the CAM cells. Similarly,

this is not an issue, and a voltage converter is not needed, as all that is required to evaluate the result of the comparison is that the logic-high signal of the senselines crosses the NMOS threshold voltage. Two serially connected NMOS comparator transistors' gates could be connected to different voltage domains, and this chain of transistors would still perform the comparator operation without the gate voltages being the same, as long as they are above the threshold voltage and not part of the critical delay (See Figure 5). The absence of level converters though does mean that we will need to operate over a limited total voltage supply range or a limited range of difference between the supply voltages. This range can be increased if level converters are incorporated into the design. However, as we will show in Section 6, a very small range of difference between the supplies in our design, of only about 0.2 V, was sufficient to optimize the operation of the CAM chip and achieve more than a two-fold improvement in the power-delay operation space.

Other optimizations in our design are that the charging of the matchline through the VCHARGE node can be blocked by the transistor M1 (shown in Figure 5) when the search operation is not being performed, which helps us reduce the static power loss through VCHARGE. To improve the noise margin and to ensure that we do not get false matches, the transistors in the comparators are sized so that the worst case single bit mismatch of the input with the stored word keeps the matchline voltage low enough so as not to switch the senseAmp.

Providing the different power supplies externally requires additional pins. Based on estimated power requirements, we budgeted two pins for VCHARGE, two pins for CVDD and four pins for DVDD. This leads to a two-pin overhead for the multi- V_{dd} design, as otherwise, a single- V_{dd} design could be supplied by six pins. The overhead of supplying the different voltages externally with additional pins can be traded-off with an internal voltage generator and its additional design complexity, area and power overhead.

5. Power Consumption Behavior of Matchline Based Content Addressable Memories

It is important to understand the power consumption behavior of CAMs to aid the design of CAMs and to evaluate the power saving mechanisms that have been proposed. The CAM matchline and senseline drivers are the two major sources of power consumption within a CAM chip. We use a modified compact modeling scheme [53] to model the power consumption behavior in the CAM cell.

The different functional blocks of the chip consume power in different fashions in the CAM chip. The power consumption of the data drivers, supplied by DVDD in our design, only depends on the Hamming distance of the input bits with respect to the previous inputs. DVDD power is independent of whether the input data matched or mismatched, and hence, independent of the power behavior of the CAM cells themselves. The CAM cell power consumption behavior depends on the match/mismatch of the data and also on the location of the mismatched bit. The power consumed by the matchline inside the CAM cell, which is supplied by VCHARGE in our design, is highly dependent on the match/mismatch of the input data with the stored data. Further, in the selective precharge scheme, in the mismatch case, the power consumed also depended on whether the mismatch happened in the selective precharge bits or in one of the NOR or ternary cells. As explained in Section 3, this is because a mismatch in the selective precharge bits stops the current from flowing further in the matchline. The power consumed by CVDD depends on whether or not the state of the CAM switches as CVDD supplies the SR-FF as well. Otherwise, the CVDD power consumption trends are the same as any typical SRAM array with additional power consumption in the SR-FF senseAmp.

Based on the observations described above, the behavior of the power supplies can be modeled using Equations (1)–(3):

$$P_{DVDD} = \sum_{D=\text{All Drivers}} (H_D * P_D^{\text{Bit}}) \quad (1)$$

$$P_{CVDD} = P_{\text{SR-FF}}^{\text{Match}} * N_{\text{Match}} + P_{\text{SRAM}} \quad (2)$$

$$P_{VCHARGE} = P_{\text{NAND}}^{\text{Miss}} * N_{\text{NAND}} + P_{\text{NOR}}^{\text{Miss}} * N_{\text{NOR}} + P_{\text{CAM}}^{\text{Match}} * N_{\text{Match}} \quad (3)$$

In Equation (1), P_{DVDD} is the total driver power over the set of all drivers as a function of H_D , the Hamming distance of the inputs of a given driver D at the current cycle and P_D^{Bit} , which is the power consumption per input bit switch. In Equation (2), P_{CVDD} is the total CVDD power, P_{SR-FF}^{Match} is the power consumed by SR-FF from a cell match, N_{Match} is the number of CAM cells that match, and P_{SRAM} is the power consumed in the SRAMs of the CAM cells during the search operation, which is usually negligible compared to the other parameters in the search operation. However, it is the major component in standby mode. In Equation (3), $P_{VCHARGE}$ is the matchline supply power composed of three components: The first component is P_{NAND}^{Miss} , the power per CAM cell mismatch triggered by a NAND cell. The second term contributing to the matchline supply power is the power per CAM cell mismatch, P_{NOR}^{Miss} , triggered by a NOR cell. Both these terms are multiplied by the number of CAM cells that experience each corresponding type of miss. The final term indicates dissipation due to a CAM cell's match state, P_{CAM}^{Match} , multiplied by N_{Match} , which is the number of matched CAM cells. This simple model classifies the power consumption behavior of the three supplies as outlined in Table 1.

The values of P_D^{Bit} , P_{SR-FF}^{Match} , P_{SRAM} , P_{NAND}^{Miss} , P_{NOR}^{Miss} and P_{CAM}^{Match} will vary based on the implementation. The coefficients denoted by $N_{\#}$ are entirely workload-dependent. The aggregate behavior as depicted in Equations (1)–(3), however, will remain similar for traditional matchline-based CAMs.

When analyzing standby power, we observe that all the supplies burn power depending on the fraction of the chip they supply power to. In conventional CAM designs, to go into standby, the supply for most of the chip will need to be kept high, so that the data in the SRAMs is not lost. However, in our design, we can completely turn off the Driver-VDD (DVDD) and VCHARGE in standby mode, which helps us save power. This is one of the reasons we chose to connect the SR-FF in the CAM cells to the same supply as the SRAMs so that the state of the CAMs can also be saved in standby.

6. Multi-V_{dd} CAM Behavior

In this section, we describe the behavior of our proposed multi-V_{dd} CAM design in detail. To simulate as well as test our chip, we use worst-case, and pseudo-realistic data banks obtained using simulations of a particle collider expected in future runs [17]. In this section, we present the behavior of our design with simulation analysis.

6.1. Unit CAM Cell

We simulated our CAM cell design in 130 nm Global Foundries PDK using Cadence tools. The CAM cell, as summarized in Table 1, is supplied by the VCHARGE and CVDD lines. Matchline risetime for a single CAM cell in our design is defined as the time it takes to charge the matchline from 10% to 90% of the charging supply. Figure 6 shows the simulation results of matchline risetime vs. the different supply voltages, where each trace corresponds to varying a single supply, while the others are kept constant at 1.35 V. From the figure, we can see that VCHARGE has the maximum effect on the matchline risetime. Increasing VCHARGE by 0.1 V effectively reduces the matchline risetime in half. The other supply in the CAM cell, CVDD, affects the matchline risetime to a much less extent. Also, shown in the figure is the behavior of the CAM cell when only a single supply is used. With only a single supply available we do not observe the same amount of gains in performance (reduction in matchline risetime delay), nor do we observe the same range of operation.




Figure 6. CAM cell matchline risetime vs. VCHARGE/CVDD.

The propagation delay of a CAM cell is the duration of time from when the inputs are available in the CAM cell to be compared, to when the output is available at the SR-FF. This delay consists of the pre-discharge delay, when the previous charge on the matchline is discharged, the matchline risetime up to the threshold of the SR-FF, and finally the delay inside the SR-FF. The first part of the delay where the previous charge is discharged depends on the size of the discharging transistor and also on how much time is to be allocated to the input data to settle inside the comparators. For our design, we found 1 ns to be sufficient to discharge the matchline and completely eliminate false positives. So, for our designs, we utilize a clock cycle, which stays high for 1 ns and low for the rest. The matchline risetime follows the behavior we observed in Figure 6 and finally, the delay through the senseAmp depends on the senseAmp design. For our design, we found the matchline risetime up to the threshold voltage to be the most sensitive to the supply voltages. This is due to two factors: The charging rate of the matchline is controlled by VCHARGE, and the threshold voltage of the SR-FF is weakly dependent on the CVDD supply. Together, these factors imply that for the fastest operation, VCHARGE should be maintained as high as possible, along with a low CVDD.

A scatter plot of the two power supplies inside the CAM versus their corresponding matchline risetime is depicted in Figure 7. Using more power from VCHARGE directly corresponds to better performance, whereas such a benefit is not apparent when budgeting more power from CVDD.




Figure 7. Scatter plot showing the effect of power on matchline risetime.

6.2. Full CAM Chip

In the full chip, the important parameter to consider is the critical propagation delay of the CAM chip, which is the path from the input pads, through the senseline interconnects, to the CAM cells, ending at the match/mismatch output of the CAM cells. Inside the CAM cell, the delay is composed of the pre-discharge time, the matchline risetime, and the delay through the SR-FF senseAmp. In our design, the critical path ending in the SR-FF helps us separate the post-processing of the match data from the performance of the CAM structures.

We simulated our triple- V_{dd} CAM design in the 130 nm Global Foundries PDK using Cadence tools. Figure 8 shows the variation in the total propagation delay versus the supply voltages. Each trace corresponds to varying a particular supply while the others are kept constant at 1.35 V. Also shown is

the behavior when only a single supply is available in the design. We can observe that VCHARGE has the most drastic effect on the performance of the CAM. The behavior shows the same dependency as that of a single CAM cell (Figure 6), and hence, shows that the matchline risetime inside the CAM cell is the critical delay whose behavior influences the whole chip behavior the most. We again see a smaller dependency between CVDD and the propagation delay. Again, this is primarily due to the relation between the threshold voltage of the SR-FF and CVDD. The propagation delay does not depend on DVDD to any appreciable degree, and hence, its effect on the delay can be ignored. This factor is dependent on the design of the drivers. Reducing searchline swing and trading off the corresponding delay penalty by increasing the supply for matchline charging can obtain power savings.




Figure 8. Propagation delay vs. value of supply voltages.

Figure 9 shows the propagation delay versus VCHARGE for different values of CVDD. Two observations are important to note from this plot: (1) The delay of the full chip depends most on the matchline risetime behavior in the CAM cells; and (2) having access to multiple supply lines allows us to operate the chip over a much larger delay range. Operating over the same voltage range, in a single supply design, the range of propagation delay is from 4.65 ns to 11.29 ns. On the other hand, when we have access to multiple supplies, we can vary the propagation delay from 3.27 ns to 26.2 ns, which is an increase in the delay range by 2.4 times.




Figure 9. Propagation delay vs. VCHARGE/CVDD.

6.3. Noise-Margin Analysis

The range of voltages we use is determined by the acceptable level of noise margin for the design. Through our analysis, we found and confirmed that the most critical case, which can result in an erroneous result from the CAM is when we have a single bit mismatch in one of the NOR bits of the CAM cell [25]. This could lead to a matchline voltage that is high enough to cause a false positive. Care must also be taken that the level of voltage is high enough to allow a low enough propagation delay for a given range of operating frequency requirement. In a current race scheme to guarantee correct operation, we need to make sure the matchline voltage remains low enough in the worst case single bit mismatch in one of the NOR cells. To assist in doing this, we size the discharge transistors in the comparators of the NOR cells to be large enough to keep the matchline low. This also increases the

area and thus, it is application dependent. Another technique is to have a higher threshold voltage for the senseAmp. However, this increases the matchline risetime required to reach that threshold and in turn, also increases the matchline power consumption.

We define the noise margin for our design as the difference in the matchline voltage levels between a perfect match and a single-bit NOR cell mismatch. By appropriately sizing the discharge transistors in the NOR cells, we can keep this well under control. For our design, we performed Monte Carlo simulations for the worst-case single-bit mismatch scenario with both process and mismatch variations. Our results showed that the voltage on the matchline increased to between 6.59 mV to 39.6 mV, with a mean of 20.66 mV, and a standard deviation of 4.79 mV. This is well below our designed matchline match voltage of 600 mV. This leads to an ample noise-margin of about 580 mV, which can be reduced judiciously for a more compact design. Figure 10 shows a Monte Carlo simulation of the distribution of the matchline voltage, for the worst case single-bit mismatch, for our design at 1.35 V. We performed this simulation for the corner cases and chose our range of voltages as those where a minimum 250 mV of noise margin was available in the worst case. A smaller noise margin can be chosen as acceptable by other applications at a higher risk of false positives.

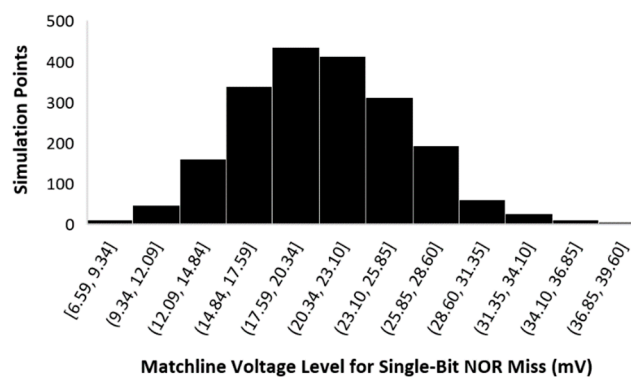


Figure 10. Distribution of matchline voltage for worst case single-bit mismatch in 2000-point Monte Carlo simulation.

6.4. Power-Delay Optimization

Finally, let us examine the power-delay operation space of our multi- V_{dd} CAM design. Figure 11 shows the total propagation delay on the x-axis and the corresponding power consumption on the y-axis for a single pattern of the chip. Each data point on the scatter plot is a combination of the three supply voltages. In this plot, operation points are also depicted for when only a single-supply is available in the design. Immediately we can see that a multi- V_{dd} design operates over a much larger power-delay space when compared to a single- V_{dd} design, which is constrained. This larger space provides us with an opportunity for optimizing the setting of the voltages to obtain better performance for the same power consumption or lower power consumption for the same performance. At every point on the single- V_{dd} power delay line, we can find an equivalent multi- V_{dd} operating point, which gives the same propagation delay while consuming less power. In Figure 11, these are points which have the same X-intercept as the single- V_{dd} operating point, but a lower Y-intercept. The lowest point we find on the multi- V_{dd} space satisfying such a condition is the optimum operating point for the respective performance/delay requirement.




Figure 11. Propagation delay vs. power dissipated.

For example, consider a point on the single- V_{dd} line, and another on the multi- V_{dd} scatter plot, both with a propagation delay of 6.4 ns. The single- V_{dd} point corresponds to setting the supply at 1.35 V, and the multi- V_{dd} point with least total power consumption corresponds to setting VCHARGE at 1.3 V, CVDD at 1.2 V and DVDD at 1.2 V. Even though they have the same delay, the multi- V_{dd} design point burns 20.04% less power. Overall, we achieve a power reduction of 25.3% in the best case and 10.61% on average using empirically obtained optimum values of supply voltages when compared to a case of similar performance with only a single supply design, operating over the same supply range. At worst, the multi- V_{dd} design consumes the same power as the single- V_{dd} design.

We extended our analysis to a more advanced 45 nm node and simulated our design using the 45 nm NCSU PDK. Figure 12 shows simulation results for a single CAM cell at the 45 nm node. Comparing Figure 12 with Figures 6, 7 and 11 shows that the overall behavior remains the same. This implies that opportunities for optimization with multi- V_{dd} supplies exist even at more advanced nodes. The only difference from the 130 nm node is that the optimal operation points will correspond to a different combination of voltages, which again can be determined empirically using a plot like Figure 11 for any particular application.

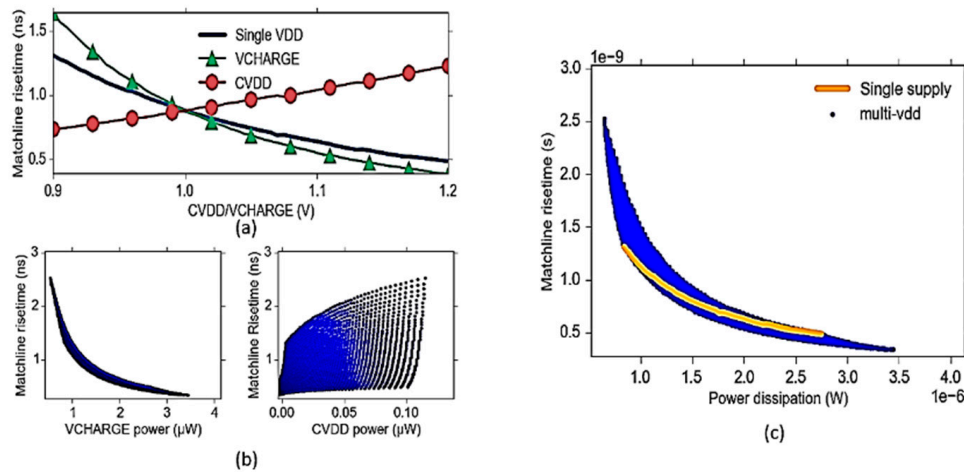


Figure 12. Multi- V_{dd} CAM at 45 nm node. (a) Matchline risetime variation; (b) scatter plot showing the effect of power on matchline risetime; (c) matchline risetime vs. power dissipated.

6.5. Standby Power Consumption

In the standby mode, for a typical chip even with clock gating, in order to conserve the data in the SRAMs and to preserve the state of the CAM chip, a large fraction of the chip would still have to be kept powered. However, in our particular design, because we connect the SRAMs and the SR-FF inside the CAM cells to a separate CVDD supply, we can turn off DVDD and VCHARGE completely, which helps us save power. CVDD can be reduced to the data retention voltage [54] thereby saving even more power. Considering a multi- V_{dd} design from the beginning also enables us to leverage

multiple SRAM optimizations, which would not be possible otherwise, or would require separate internal voltage generation [24,25]. Our measurements with our test chip revealed that by turning off DVDD and VCHARGE in the standby mode, we can save ~51.3% of the standby power, as we detail in the next section.

7. Experimental Results

To verify our analysis, we fabricated our design in 130 nm Global Foundries 8 metal layer technology. Figure 4 shows a photograph of the test chip. The chip is part of the VIPRAM project at Fermilab [6]. It is a prototype PRAM chip for future upgrades to the Level-1 trigger system of the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) in CERN [17]. The chip die size is 5.46 mm by 5.46 mm. A single 15-bit CAM cell macro takes 25 μm by 25 μm . The test chip has a CAM array of 128 rows by 128 columns for a total memory size of 246 kilobits. The chip was wire bonded on a 144-pin quad flat package. The typical power consumption of the chip is approximately 190 mW when operating at 75 MHz with a supply of 1.35 V. Testing was performed on a custom FPGA board [15] connected to a Linux machine [55]. The setup permitted changing of voltages for all the supplies independently using the external voltage regulator. Figure 13 shows a picture of our testing board. Figures 14–17 show some of the essential experimental results from the chip.

Figure 14 shows the active power dissipation distribution while the chip is performing search operations. From Figure 14 we see that the maximum amount of power is burned in driving the data across the chip (DVDD), followed by charging the matchline (VCHARGE), and finally the SRAMs and the SR-FF (CVDD). This implies decreasing power dissipation from DVDD, i.e., the searchline driver power should have the most drastic effect on reducing total power consumption. We scanned the chip to find the maximum achievable clock frequency for different voltage combinations. The inverse of these frequencies gave us an estimation of the delays inside the CAM chip. Thereby, we obtained a voltage-delay operation region for the CAM chip where the different data points correspond to different settings of the input voltage levels. Again, we observed no dependence of internal delay on DVDD. Figure 15 shows these measured results, showing the relation between the voltage levels of VCHARGE and CVDD/DVDD plotted against the critical propagation delay. These measurements can be compared with the simulation data from Figure 9. We observe the same trend in both. Power measurements from the chip agreed with our simulations and showed a wide range of power consumption with dissipation increasing from every supply with increasing voltage. Choosing two points from Figure 11 we see that in a single supply setup, with all voltages at the same level of 1.33 V, we measured a propagation delay of 12.7 ns; whereas in a multi- V_{dd} setup, we were able to obtain a delay of 12.7 ns with voltage settings of 1.3 V for VCHARGE, 1.25 V for CVDD and 1.25 V for DVDD while achieving a 15% reduction in power.




Figure 13. Test CAM chip mounted on the FPGA testing board.




Figure 14. Active Power consumption distribution among the three supplies.

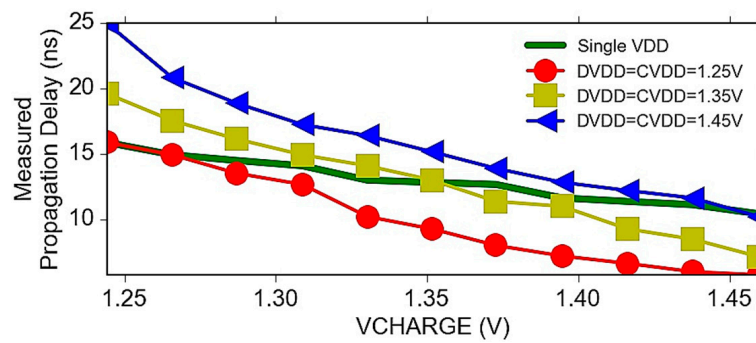


Figure 15. Testing results showing the measured delay.




Figure 16. Standby power consumption behavior w.r.t. supply voltage.




Figure 17. Standby Power consumption distribution among the three supplies.

We performed standby current measurements with our experimental setup. We summarize our experimental results in Figure 16, which shows the standby power variation in our design when all three supplies are powered and kept at the same voltage. At 1.35 V, we observed a current of 0.88 mA from DVDD, 0.92 mA from CVDD and 0.0084 mA from VCHARGE. Figure 17 shows the distribution of this standby power. This amounts to a total of ~2.5 mW standby power. This can be reduced to just 1.242 mW without affecting operation by simply turning off the DVDD and VCHARGE supplies, which is a saving of 51.3%. These savings can be further improved by scaling down the CVDD supply to the data retention voltage (DRV) of the SRAMs [54]. Such high standby power savings are very encouraging, especially as our design does not incorporate any specialized design techniques to reduce standby power, besides using a multi- V_{dd} supply. Specialized standby power saving techniques [24] can be further applied along with the multi- V_{dd} technique to achieve even better results.

8. Conclusions

We presented a thorough power characterization of matchline-based content addressable memories. Through this analysis, we proposed a customized multi- V_{dd} scheme in CAMs. From the power model, simulation analysis and testing results we see that the use of the multi- V_{dd} scheme in CAMs helps reduce power consumption and at the same time makes the chip's performance highly tunable. We showed the existence of an optimum operation point for a particular delay requirement in the power-delay space of multi- V_{dd} CAM devices, which provides the same performance as a single- V_{dd} device at much lower power consumption. We found significant standby power savings when the storage cells of the CAMs and the output register were connected to a separate supply. Finally, we validated our analysis and design by presenting measurement results from a test chip employing our multi- V_{dd} design in 130 nm 8 metal layer Global Foundries technology.

Author Contributions: Conceptualization and Investigation, S.J.; Funding acquisition, S.O.-M. and T.L.; Software, S.J., J.O., S.J., D.L. and N.T.; Supervision, S.O.-M., G.D., J.H. and T.L.; Writing—original draft, S.J.; Writing—review & editing, S.O.-M., J.H. and T.L.

Funding: This work was partially supported by the NSF Grant CCF-1422489 and the Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the US Department of Energy.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Karam, R.; Puri, R.; Ghosh, S.; Bhunia, S. Emerging Trends in Design and Applications of Memory-Based Computing and Content-Addressable Memories. *Proc. IEEE* **2015**, *103*, 1311–1330. [[CrossRef](#)]
2. Pagiamtzis, K.; Sheikholeslami, A. Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey. *IEEE J. Solid-State Circuits* **2006**, *41*, 712–727. [[CrossRef](#)]
3. Hoff, J.R.; Deptuch, G.W.; Joshi, S.; Liu, T.; Olsen, J.; Shenai, A. VIPRAM_L1CMS: A 2-Tier 3D Architecture for Pattern Recognition for Track Finding. In Proceedings of the 2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD), Strasbourg, France, 29 October–6 November 2016; pp. 1–7.
4. Tony, S.M.; Wu, F.; Li, H.; Huang, P.; Rahimi, A.; Rabaey, J.M.; Wong, H.-S.P.; Shulaker, M.M. Brain-Inspired Computing Exploiting Carbon Nanotube FETs and Resistive RAM: Hyperdimensional Computing Case Study. In Proceedings of the 2018 IEEE International Solid-State Circuits Conference, San Francisco, CA, USA, 11–15 February 2018; pp. 492–494.
5. Li, H.; Wu, T.F.; Rahimi, A.; Li, K.-S.; Rusch, M.; Lin, C.-H.; Hsu, J.-L.; Sabry, M.M.; Eryilmaz, S.B.; Sohn, J.; et al. Hyperdimensional computing with 3D VRRAM in-memory kernels: Device-architecture co-design for energy-efficient, error-resilient language recognition. In Proceedings of the 2016 IEEE International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 3–7 December 2016; pp. 16.1.1–16.1.4.
6. Liu, T.; Hoff, J.; Deptuch, G.; Yarema, R. A New Concept of Vertically Integrated Pattern Recognition Associative Memory. *Phys. Procedia* **2012**, *37*, 1973–1982. [[CrossRef](#)]

7. Hu, Y.J.; Li, J.F.; Huang, Y.J. 3-D content addressable memory architectures. In Proceedings of the 2009 IEEE International Workshop on Memory Technology, Design, and Testing, Hsinchu, Taiwan, 31 August–2 September 2009; pp. 59–64.
8. Mathan, K.; Ravichandran, T. Data Intelligent Low Power High Performance TCAM for IP-Address Lookup Table. *Circuits Syst.* **2016**, *7*, 3734–3745. [[CrossRef](#)]
9. Guo, Q.; Guo, X.; Bai, Y.; İpek, E. A Resistive TCAM Accelerator for Data-Intensive Computing Categories and Subject Descriptors. In Proceedings of the 2011 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), Porto Alegre, Brazil, 3–7 December 2011; pp. 339–350.
10. Rahimi, A.; Ghofrani, A.; Cheng, K.; Benini, L.; Gupta, R.K. Approximate Associative Memristive Memory for Energy-Efficient GPUs. In Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition (DATE), Grenoble, France, 9–13 March 2015; pp. 1497–1502.
11. Ghofrani, A.; Rahimi, A.; Lastras-Montano, M.A.; Benini, L.; Gupta, R.K.; Cheng, K.T. Associative Memristive Memory for Approximate Computing in GPUs. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **2016**, *6*, 222–234. [[CrossRef](#)]
12. Imani, M.; Rahimi, A.; Rosing, T.S. Resistive Configurable Associative Memory for Approximate Computing. In Proceedings of the 2016 Design, Automation & Test in Europe Conference & Exhibition (DATE), Dresden, Germany, 14–18 March 2016; pp. 1327–1332.
13. Wade, J.P.; Sodini, C.G. A Ternary Content Addressable Search Engine. *IEEE J. Solid-State Circuits* **1989**, *24*, 1003–1013. [[CrossRef](#)]
14. Deptuch, G.; Hoff, J.; Jndariani, S.; Liu, T.; Olsen, J.; Tran, N.; Joshi, S.; Li, D.; Ogrenici-Memik, S. Performance Study of the First 2D Prototype of Vertically Integrated Pattern Recognition Associative Memory (VIPRAM). *arXiv*, **2017**, arXiv:1709.08303.
15. Liu, T.; Deptuch, G.; Hoff, J.; Jndariani, S.; Joshi, S.; Olsen, J.; Tran, N.; Trimpl, M. Design and testing of the first 2D Prototype Vertically Integrated Pattern Recognition Associative Memory. *J. Instrum.* **2015**, *10*, 1–8. [[CrossRef](#)]
16. Annovi, A.; Bertolucci, F.; Biesuz, N.; Calabro, D.; Calderini, G.; Citraro, S.; Crescioli, F.; Dimas, D.; Dell’Orso, M.; Donati, S.; et al. Highly parallelized pattern matching execution for the ATLAS experiment. In Proceedings of the 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), San Diego, CA, USA, 31 October–7 November 2015; pp. 15–17.
17. CMS Collaboration. Technical Proposal for the Phase-II Upgrade of the CMS Detector. CERN-LHCC-2015-010, LHCC-P-008. 2015. Available online: <https://cds.cern.ch/record/2020886> (accessed on 28 July 2018).
18. Annovi, A.; Beretta, M.M.; Calderini, G.; Crescioli, F.; Frontini, L.; Liberali, V.; Shojaii, S.R.; Stabile, A. AM06: The Associative Memory chip for the Fast Tracker in the upgraded ATLAS detector. *J. Instrum.* **2017**, *12*. [[CrossRef](#)]
19. Lehtonen, E.; Poikonen, J.H.; Laiho, M.; Kanerva, P. Large-scale memristive associative memories. *IEEE Trans. Very Large Scale Integr. Syst.* **2014**, *22*, 562–574. [[CrossRef](#)]
20. Joshi, S.; Li, D.; Ogrenici-Memik, S.; Deptuch, G.; Hoff, J.; Jndariani, S.; Liu, T.; Olsen, J.; Tran, N. A Content Addressable Memory with Multi-Vdd Scheme for Low Power Pattern Recognition. In Proceedings of the 60th IEEE International Midwest Symposium on Circuits and Systems, Boston, MA, USA, 6–9 August 2017.
21. Zackriya, M.; Kittur, H.M. Content Addressable Memory—Early predict and terminate precharge of Match-Line Content Addressable Memory. *IEEE Trans. Very Large Scale Integr. Syst.* **2016**, *25*, 385–387. [[CrossRef](#)]
22. Pagiamtzis, K.; Sheikholeslami, A. A low-power content-addressable memory (CAM) using pipelined hierarchical search scheme. *IEEE J. Solid-State Circuits* **2004**, *39*, 1512–1519. [[CrossRef](#)]
23. Mahendra, T.V.; Mishra, S.; Dandapat, A. Self-Controlled High-Performance Precharge-Free Content-Addressable Memory. *IEEE Trans. Very Large Scale Integr. Syst.* **2017**, *25*, 2388–2392. [[CrossRef](#)]
24. Mohan, N.; Sachdev, M. Low-Leakage Storage Cells for Ternary Content. *IEEE Trans. Very Large Scale Integr. Syst.* **2009**, *17*, 604–612. [[CrossRef](#)]
25. Arsovski, I.; Chandler, T.; Sheikholeslami, A. A ternary content-addressable memory (TCAM) based on 4T static storage and including a current-race sensing scheme. *IEEE J. Solid-State Circuits* **2003**, *38*, 155–158.

26. Zukowski, C.; Wang, S.-Y. Use of selective precharge for low-power content-addressable memories. In Proceedings of the 1997 IEEE International Symposium on Circuits and Systems. Circuits and Systems in the Information Age ISCAS'97, Hong Kong, China, 12 June 1997; pp. 1788–1791.
27. Do, A.; Yin, C.; Velayudhan, K.; Lee, Z.; Yeo, K.; Kim, T. 0.77 fJ/bit/search Content Addressable Memory Using Small Match Line Swing and Automated Background Checking Scheme for Variation Tolerance. *IEEE J. Solid-State Circuits* **2014**, *49*, 1487–1498. [[CrossRef](#)]
28. Mohan, N.; Sachdev, M. A Static Power Reduction Technique for Ternary Content Addressable Memories. In Proceedings of the Canadian Conference on Electrical and Computer Engineering 2004, Niagara Falls, ON, Canada, 2–5 May 2004; pp. 711–714.
29. Agarwal, A.; Hsu, S.K.; Kaul, H.; Anders, M.A.; Krishnamurthy, R.K. A Dual-Supply GHz 13 fJ/bit/search 64×128 CAM in 65 nm CMOS. In Proceedings of the 2006 32nd European Solid-State Circuits Conference, Montreux, Switzerland, 19–21 September 2006; pp. 4–7.
30. Mohan, N.; Sachdev, M. Low-capacitance and charge-shared match lines for low-energy high-performance TCAMs. *IEEE J. Solid-State Circuits* **2007**, *42*, 2054–2060. [[CrossRef](#)]
31. Yang, B.D.; Lee, Y.K.; Sung, S.W.; Min, J.J.; Oh, J.M.; Kang, H.J. A low power content addressable memory using low swing search lines. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2011**, *58*, 2849–2858. [[CrossRef](#)]
32. Arsovski, I.; Sheikholeslami, A. A mismatch-dependent power allocation technique for match-line sensing in content-addressable memories. *IEEE J. Solid-State Circuits* **2003**, *38*, 1958–1966. [[CrossRef](#)]
33. Igarashi, M.; Usami, K.; Nogami, K.; Minami, F.; Kawasaki, Y.; Aoki, T.; Takano, M.; Sonoda, S.; Ichida, M.; Hatanaka, N. A low-power design method using multiple supply voltages. In Proceedings of the 1997 International Symposium on Low Power Electronics and Design, Monterey, CA, USA, 18–20 August 1997; pp. 36–41.
34. Usami, K.; Horowitz, M. Clustered voltage scaling technique for low-power design. In Proceedings of the 1995 International Symposium on Low Power Design, Dana Point, CA, USA, 23–26 April 1995; pp. 3–8.
35. Mondal, H.K.; Gade, S.H.; Kishore, R.; Deb, S. Adaptive Multi-Voltage Scaling in Wireless NoC for High Performance Low Power Applications. In Proceedings of the 2016 Design, Automation & Test in Europe Conference & Exhibition (DATE), Dresden, Germany, 14–18 March 2016; pp. 1315–1320.
36. Friedman, E.G.; Kursun, V. *Multi-Voltage CMOS Circuit Design*; John Wiley & Sons: Hoboken, NJ, USA, 2006.
37. Shibata, N.; Watanabe, M.; Okiyama, H. A high-speed low-power multi-VDD CMOS/SIMOX SRAM with LV-TTL level input/output pins-write/read assist techniques for 1-V operated memory cells. *IEEE J. Solid-State Circuits* **2010**, *45*, 1856–1869. [[CrossRef](#)]
38. Kulkarni, J.; Khellah, M.; Tschanz, J.; Geuskens, B.; Jain, R.; Kim, S.; De, V. Dual-VCC 8T-bitcell SRAM Array in 22 nm tri-gate CMOS for energy-efficient operation across wide dynamic voltage range. In Proceedings of the 2013 Symposium on VLSI Circuits, Kyoto, Japan, 12–14 June 2013; pp. 352–353.
39. Koo, K.-H.; Wei, L.; Keane, J.; Bhattacharya, U.; Karl, E.A.; Zhang, K. A 0.094 μm^2 High Density and Aging Resilient 8T SRAM with 14 nm FinFET Technology Featuring 560 mV VMIN with Read and Write Assist. In Proceedings of the 2015 Symposium on VLSI, Kyoto, Japan, 17–19 June 2015.
40. Takeda, K.; Hagihara, Y.; Aimoto, Y.; Nomura, M.; Nakazawa, Y.; Ishii, T.; Kobatake, H. A Read-Static-Noise-Margin-Free SRAM Cell for low-VDD high speed applications. *IEEE J. Solid-State Circuits* **2006**, *41*, 113–121. [[CrossRef](#)]
41. Mondal, S.; Memik, S.O. A low power FPGA routing architecture. In Proceedings of the 2005 IEEE International Symposium on Circuits and Systems, Kobe, Japan, 23–26 May 2005; pp. 1222–1225.
42. Mukherjee, R.; Memik, S. Evaluation of dual VDD fabrics for low power FPGAs. In Proceedings of the 2005 Design Automation Conference, Anaheim, CA, USA, 13–17 June 2005; pp. 1240–1243.
43. Mukherjee, R.; Liu, S.; Merik, S.O.; Mondal, S. A high-level clustering algorithm targeting dual Vdd FPGAs. *ACM Trans. Des. Autom. Electron. Syst.* **2008**, *13*, 57. [[CrossRef](#)]
44. Li, F.; Lin, Y.; He, L. FPGA power reduction using configurable dual-vdd. In Proceedings of the 2004 Design Automation Conference, San Diego, CA, USA, 7–11 June 2004.
45. Mondal, S.; Memik, S.O. Power optimization techniques for SRAM-based FPGAS. In Proceedings of the 2006 International Conference on Field Programmable Logic and Applications, Madrid, Spain, 28–30 August 2006; pp. 959–960.
46. Mukherjee, R.; Memik, S.O. Power-Driven Design Partitioning. In Proceedings of the International Conference on Field Programmable Logic and Applications, Leuven, Belgium, 30 August–1 September 2004.

47. Do, A.T.; Chen, S.; Kong, Z.H.; Yeo, K.S. A low-power CAM with efficient power and delay trade-off. In Proceedings of the 2011 IEEE International Symposium of Circuits and Systems (ISCAS), Rio de Janeiro, Brazil, 15–18 May 2011; pp. 2573–2576.
48. Augsburger, S.; Nikolić, B. Combining Dual-Supply, Dual-Threshold and Transistor Sizing for Power Reduction. In Proceedings of the IEEE International Conference on Computer Design: VLSI in Computers and Processors, Freiberg, Germany, 18 September 2002.
49. Chandrakasan, A.P.; Sheng, S.; Brodersen, R.W. Low-power CMOS digital design. *IEEE J. Solid-State Circuits* **1992**, *27*, 473–484. [[CrossRef](#)]
50. Gonzalez, R.; Gordon, B.M.; Horowitz, M.A. Supply and Threshold Voltage Scaling for Low Power CMOS. *IEEE J. Solid-State Circuits* **1997**, *32*, 1210–1216. [[CrossRef](#)]
51. Bijansky, S.; Lee, S.K. TuneLogic: Post-Silicon Tuning of Dual-Vdd Designs. In Proceedings of the 2009 10th International Symposium on Quality Electronic Design, San Jose, CA, USA, 16–18 March 2009.
52. Usami, K.; Igarashi, M. Low-power design methodology and applications utilizing dual supply voltages. In Proceedings of the Design Automation Conference, Yokohama, Japan, 25–28 January 2000; pp. 123–128.
53. Li, D.; Joshi, S.; Ogreni-Memik, S.; Hoff, J.; Jindariani, S.; Liu, T.; Olsen, J.; Tran, N. A methodology for power characterization of associative memories. In Proceedings of the 2015 33rd IEEE International Conference on Computer Design (ICCD), New York, NY, USA, 18–21 October 2015; pp. 491–498.
54. Qin, H.; Cao, Y.; Markovic, D.; Vladimirescu, A.; Rabaey, J. SRAM leakage suppression by minimizing standby supply voltage. In Proceedings of the 2003 International Symposium on Signals, Circuits and Systems, San Jose, CA, USA, 22–24 March 2004; pp. 2–7.
55. Frazier, R.; Iles, G.; Newbold, D.; Rose, A. Software and firmware for controlling CMS trigger and readout hardware via gigabit Ethernet. *Phys. Procedia* **2012**, *37*, 1892–1899. [[CrossRef](#)]



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).