

Custom Memory Design for Logic-in-Memory: Drawbacks and Improvements over Conventional Memories

Original

Custom Memory Design for Logic-in-Memory: Drawbacks and Improvements over Conventional Memories / Ottati, Fabrizio; Turvani, Giovanna; Masera, Guido; Vacca, Marco. - In: ELECTRONICS. - ISSN 2079-9292. - ELETTRONICO. - 10:18(2021), p. 2291. [10.3390/electronics10182291]

Availability:

This version is available at: 11583/2926514 since: 2021-09-23T07:35:26Z

Publisher:

MDPI

Published

DOI:10.3390/electronics10182291

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Article

Custom Memory Design for Logic-in-Memory: Drawbacks and Improvements over Conventional Memories

Fabrizio Ottati , Giovanna Turvani , Guido Masera  and Marco Vacca 

VLSI Laboratory, Department of Electronics and Telecommunications (DET), Politecnico di Torino, Corso Duca Degli Abruzzi 24, 10129 Torino, Italy; giovanna.turvani@polito.it (G.T.); guido.masera@polito.it (G.M.); marco.vacca@polito.it (M.V.)

* Correspondence: fabrizio.ottati@polito.it

Abstract: The speed of modern digital systems is severely limited by memory latency (the “Memory Wall” problem). Data exchange between Logic and Memory is also responsible for a large part of the system energy consumption. Logic-in-Memory (LiM) represents an attractive solution to this problem. By performing part of the computations directly inside the memory the system speed can be improved while reducing its energy consumption. LiM solutions that offer the major boost in performance are based on the modification of the memory cell. However, what is the cost of such modifications? How do these impact the memory array performance? In this work, this question is addressed by analysing a LiM memory array implementing an algorithm for the maximum/minimum value computation. The memory array is designed at physical level using the FreePDK 45 nm CMOS process, with three memory cell variants, and its performance is compared to SRAM and CAM memories. Results highlight that read and write operations performance is worsened but in-memory operations result to be very efficient: a 55.26% reduction in the energy-delay product is measured for the AND operation with respect to the SRAM read one. Therefore, the LiM approach represents a very promising solution for low-density and high-performance memories.

Keywords: Logic-in-Memory (LiM); In-Memory Computing (IMC); Near-Memory Computing (NMC); Memory Wall



Citation: Ottati, F.; Turvani, G.; Masera, G.; Vacca, M. Custom Memory Design for Logic-in-Memory: Drawbacks and Improvements over Conventional Memories. *Electronics* **2021**, *10*, 2291. <https://doi.org/10.3390/electronics10182291>

Academic Editor: Kiat Seng Yeo

Received: 8 June 2021

Accepted: 13 September 2021

Published: 17 September 2021

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 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 (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Modern digital architectures are based on the Von Neumann principle: the system is divided into two main units, a central processing one and a memory. The CPU extracts the data from the memory, elaborates them and writes the results back. This structure represents the main performance bottleneck of modern computing systems: in fact, memories are not able to supply data to CPUs at a speed similar to the processing one, limiting the throughput of the whole system; moreover, high-speed data exchange between CPU and memory leads to large power consumption. This problem is commonly referred to as the “Memory Wall” problem or the “Von Neumann bottleneck”. A complex memory hierarchy is employed to partially compensate for this, but it does not completely solve it: the system results to be still limited by the impossibility to have a memory that is large and very fast at the same time.

For these reasons, companies and researchers are searching for a way to overcome the Memory Wall problem: Logic-in-Memory (LiM), also called In-Memory Computing (IMC) [1], is a computing paradigm that is being investigated for this purpose. In this model, part of the computation is executed inside the memory. This result is achieved modifying the memory architecture by adding logic circuits to it. Since part of the computation is performed directly inside the memory, the CPU is not limited by the memory latency when some operations have to be performed. In addition to this, the rate at which data are exchanged between CPU and memory is reduced, resulting in power consumption reduction.

Many approaches to Logic-In-Memory can be found in literature; however, two main approaches can be distinguished. The first one can be classified as Near-Memory Computing (NMC) [2–18], since the memory inner array is not modified and logic circuits are added at the periphery of this; the second one can be denoted as Logic-in-Memory (LiM) [19–28], since the memory cell is directly modified by adding logic circuits to it.

In an NMC architecture, logic and arithmetic circuits are added on the memory array periphery, in some cases exploiting 3D structures; therefore, the distance between computational and memory circuits is shortened, resulting in power saving and latency reduction for the data exchange between these. For instance: in [3], logic and arithmetic circuits are added on the bottom of an SRAM (Static Random Access Memory) array, where the data are transferred from different memory blocks, elaborated and, then, written back to the array; in [2], a DRAM (Dynamic Random Access Memory) is modified to perform logic bitwise operations on the bitlines, and the sense amplifiers are configured as programmable logic gates. Near-Memory Computing allows to maximise the memory density, with minimal modifications to the memory array itself, which is the most critical part of memory design; this results in a limited performance improvement with respect to computing systems based on conventional memories.

In a LiM architecture, the memory cells and periphery are modified by adding logic and arithmetic circuits to them, resulting in true in-memory processing, with the data being elaborated also inside each memory cell. For instance: in [28], a XOR logic gate is added to each memory cell to implement a Binary Neural Network (BNN) directly in memory; in [25], an SRAM is modified at the cell level to perform logic operations directly in the cell, which results are then combined by appositely designed sense amplifiers on the periphery of the array. This approach leads to a reduction in memory density since the cell footprint is increased; nevertheless, the resulting performance boost is huge, since all the data stored in memory can be elaborated at once from the inner array.

Many applications can benefit from the IMC approach, such as machine learning and deep learning algorithms [4,6,8–12,14,15,19,21–24], but also general purpose algorithms [2,5,7,13,16–18,20,25,26]. For instance: in [19], a 6T SRAM cell is modified by adding two transistors and a capacitor to it, in order to perform analog computing on the whole memory, which allows to implement approximated arithmetic operations for machine learning algorithms; in [18], logic layers consisting of latches and LUTs are interleaved with memory ones in an SRAM array, in order to perform different kinds of logic operations directly inside the array; in [26], the pass transistors of the 6T SRAM cell are modified to perform logic operations directly in the cell, which allows the memory to function as an SRAM, a CAM (Content Addressable Memory) or a LiM architecture. In general, every algorithm that works on high parallelism data and performs many element-wise operations in parallel (e.g., neural networks), is likely to receive a performance improvement when IMC solutions are employed.

Another interesting field of application is represented by Neuromorphic Computing [29,30] based on Beyond-CMOS technologies, such as memristive ones. This kind of device is well suited for IMC or LiM applications, thanks to their non-volatile characteristics and low cell area footprint. For instance, in [31] a VRRAM array is produced for a neuromorphic application, by implementing an in-memory XNOR operation for the synaptic weights.

The modification of the memory cell circuit by the addition of computational elements to it, is a risky solution: memories are circuits with a very high level of optimization; hence, even a minor modification can have a large impact on their behaviour and performance; moreover, this approach results in a reduction of the memory density. At the same time, a large boost in the overall system performance can be obtained, since all the stored data can be processed at once. As a consequence, the LiM approach represents an interesting option for low-density and high-performance memories, like caches. It is important to identify the impact that the modification of a memory cell circuit has on standard memory

operations (read and write) and on in-memory logic operations, evaluating objectively the advantages and disadvantages of the approach.

The goal of this work is to understand and quantify this impact. As a case study, an algorithm for the maximum/minimum computation [27] based on the bitwise logic AND operation is used. The array is designed and characterised at transistor and layout levels in Cadence Virtuoso, using FreePDK 45 nm CMOS process. Three different solutions for the memory cell circuit are investigated, that implement the same logic function, then, the array performance is compared to two conventional memories, a 6T SRAM and a NOR CAM, by considering the latency and energy consumption of each memory operation. The results highlight that modifying the memory certainly affects in a non negligible way the read and write operations performance, but this impact can be greatly reduced by proper design and optimisation of the memory cell circuit; nevertheless, in-memory logic operations result to be very efficient in terms of energy consumption. In fact, a 55.26% reduction in the energy-delay product of the AND operation, with respect to the SRAM read one, is observed. The results obtained suggest that LiM architectures represent a very good alternative for the implementation of algorithm accelerators which can be used as secondary memories, where the execution rate of read and write operations is lower than the in-memory logic operations one.

The paper outline is the following:

- in Section 2, the design of conventional memories (SRAM and CAM) implementations to act as performance references is discussed.
- in Section 3, the design of the LiM array and the three memory cells types is analyzed.
- in Section 4 the testbench for the characterisation of the memory arrays produced is presented.
- in Section 5, the simulation framework adopted is discussed.
- in Section 6, the obtained results are presented and analysed.
- in Section 7, some considerations about the results and the architecture are provided.

The main contributions of this paper are the following:

- a LiM array, implementing a specific algorithm [27] as a case study, is designed at physical level using the FreePDK 45 nm CMOS process and characterised through extensive SPICE simulations.
- three variants of the LiM cell are designed and characterised.
- the LiM array performance are compared to conventional memories ones; in particular, a SRAM and a CAM arrays are designed and simulated using the same parameters of the LiM array.
- to characterise the design for large memory sizes, a circuitual model that allows to strongly reduce the circuit netlist size is proposed and adopted to shorten as much as possible the simulation time of large arrays.
- to speed up the design of custom memory arrays such as LiM ones, a scripting approach is proposed and adopted.

2. Reference Architectures

In order to properly characterise the LiM architecture design, two standard memory arrays, SRAM and CAM, are produced in Cadence Virtuoso to be used as reference circuits: the SRAM array is chosen since it provides a lower ground for the memory cell circuit complexity that can be used as a reference by the other memory architectures; the CAM array, instead, is chosen since it is an example of Logic-in-Memory architecture (each memory cell performs an XNOR operation) widely used nowadays. The cell topologies chosen for these memory architectures are the 6T SRAM and the NOR CAM [32].

For the SRAM array, the standard 6T cell memory cell is chosen (Figure 1a): since the aim of this work is to produce a memory architecture capable of performing logic operations, the cell area dedicated to the memory function is minimised by picking the design with the smallest cell footprint possible for the SRAM core. For what concerns

the read sensing circuitry, a conventional voltage latch sense amplifier (SA) [33] is chosen, which circuit is depicted in Figure 1b. A commonly adopted SA circuit topology is selected to compare the read operation performance among the memories, in order to understand how much the added complexity affects the standard memory operation of the array. This circuit provides a high sensing speed and low idle power consumption, which are due to the non linearity of the bi-stable ring used as latch.




Figure 1. The SRAM cell and SA. (a) The 6T cell. (b) The SRAM sense amplifier (SA) [33]. (c) The SRAM cell layout.

For the CAM, a conventional NOR topology [32] (Figure 2a), is employed. For what concerns the CAM sensing circuitry, a current-saving scheme [34] is selected among the possible ones [32]. The correspondent matchline sense amplifier (MLSA) circuit is depicted in Figure 2b. In CAM memories, this circuit is employed to reduce, with respect to the standard sensing scheme, the energy consumption associated to a search operation, thanks to the fact that the matchline (ML) is charged in case of match instead of being discharged when a mismatch occurs. In fact, it is well known that during a search operation in a NOR CAM array, the mismatch result is the most frequent one (only one or few words in the memory array match the searched one). By associating a matchline voltage commutation to the match result instead of the mismatch one, a large reduction in the energy consumption

associated to the search operation is obtained, since only few lines experience a variation of their electric potential.

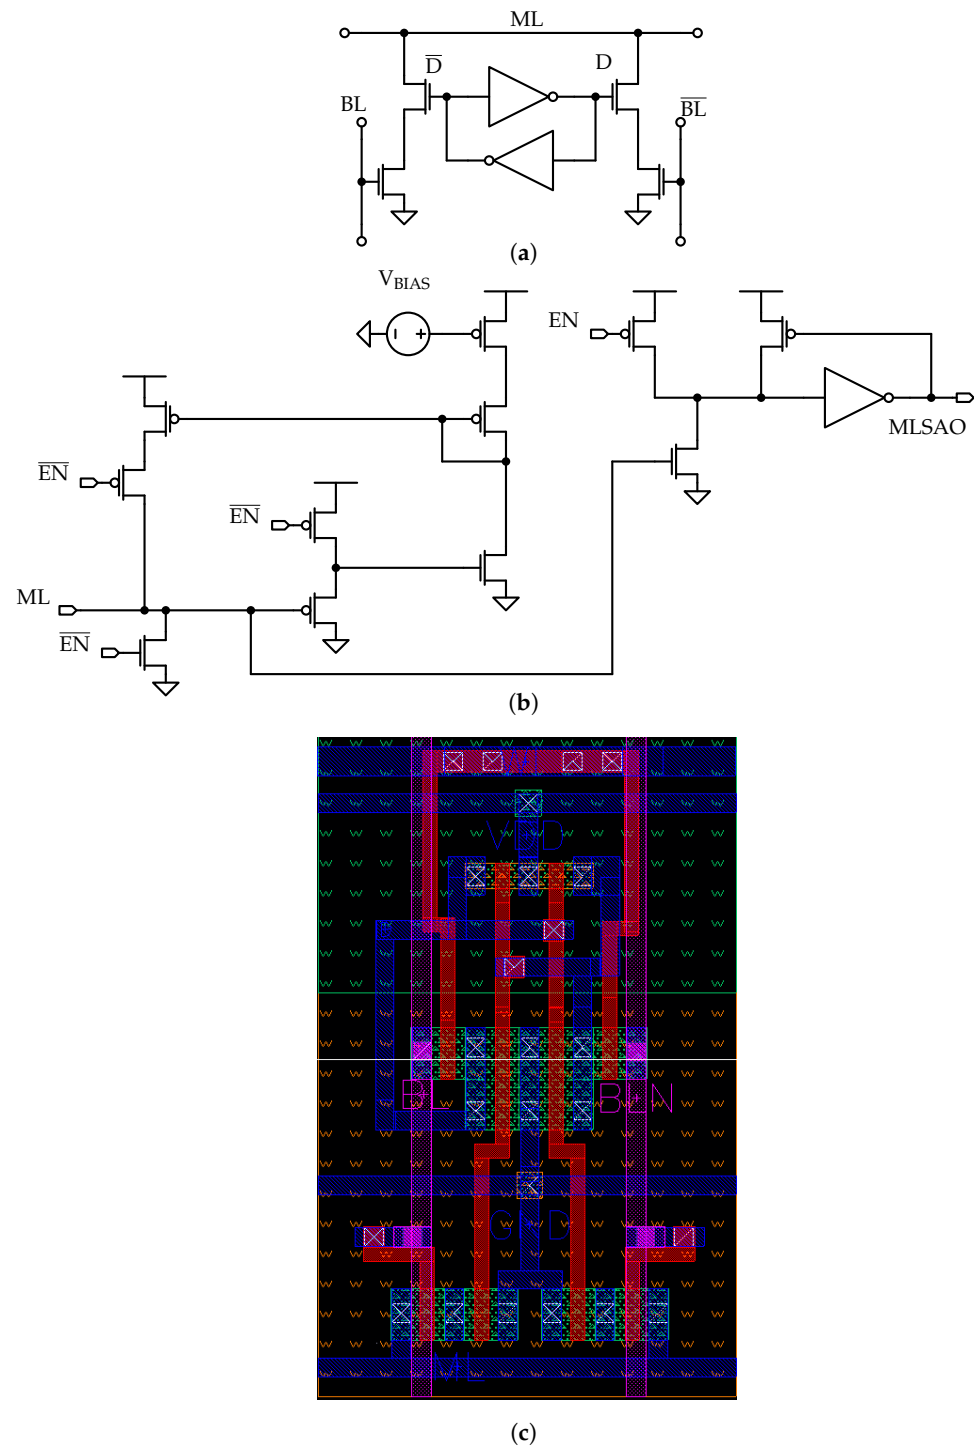


Figure 2. The CAM cell and MLSA. (a) Simplified schematic of the 10T NOR CAM cell [32]. The access transistors of the SRAM core are omitted. (b) The matchline sense amplifier (MLSA) [34], that employs the current-save sensing scheme for the search operation of the CAM array. (c) The CAM cell layout.

In Figure 2b, an example of current-saving scheme [32] is presented. This consists of a current source used to charge the matchline; when a match occurs, the matchline behaves as a capacitance; as a consequence, the capacitance gets charged resulting in a matchline voltage variation, and a match is registered in output. In case of a mismatch, instead, the ML connects the current source to ground and it does not get charged, preventing

a variation in the matchline electric potential, which would lead to additional energy consumption. A feedback control circuit is employed to limit the current that is injected to ground in the mismatch case, in order to save power during the search operation; this circuit allows to deliver as few current as possible to the mismatch line, while providing the match ones with as much current as possible to speed up the match sensing.

In order to limit the conduction time of the MLSAs current sources, the circuit of the MLSA and the architecture are modified. To turn off all the current sources as soon as all the matchlines values are correctly sensed, i.e., all the matching lines are charged to the MLSAs input threshold, so that no current is wasted in the mismatch lines, the “dummy matchline” scheme, shown in Figure 3, is employed.

In Figure 3a, a dummy CAM cell is shown. This consists of a CAM cell from which all the transistors that are not connected to the matchline are removed. The gates electric potentials of the remaining MOSFETs are chosen so that the cell always provides a match, i.e., it behaves as a capacitance. In fact, since the result that involves a voltage variation on the line is the match one, the latter determines the search operation performance.

In Figure 3b, a dummy ML is shown. The dummy cells are arranged in a row, which is connected to an MLSA that provides in output a “dummy match result”, denoted with *Dummy_MLSAO*, at each search operation. This signal is used in the architecture to disable all the real MLSAs, as soon as a match is detected on the dummy ML.

In Figure 3c, the circuit of the MLSA is depicted. An OR gate is added to each MLSA, and its output is used as an internal enable signal inside this. In particular, since the enable signal is low-active, the output of the OR gate should switch to ‘1’ as soon as *Dummy_MLSAO* switches to ‘1’, i.e., a match is detected on the matchline, in order to disable the MLSA current source. As a consequence, the global enable signal \overline{EN} is connected using a logic OR with *Dummy_MLSAO*.

In Figure 3d, the whole CAM architecture is shown. As explained above, the output of the dummy MLSA is connected to all the MLSAs, together with the global enable signal. Since the dummy matchline sensing delay determines the time available to correctly sense the matchline potential for each MLSA, its position in the memory array is crucial for the circuit timing. This means that the worst-case delay has to be associated to the dummy matchline position, i.e., it has to be placed as far as possible from the enable signals drivers in the circuit.

In Figure 4 it is shown a section of the layout for the dummy line. One can notice that some transistors are missing from the original layout depicted in Figure 2c: in fact, the SRAM core is modified so that the cell stores a logic ‘1’ without needing to explicitly write this value to each cell of the dummy line.

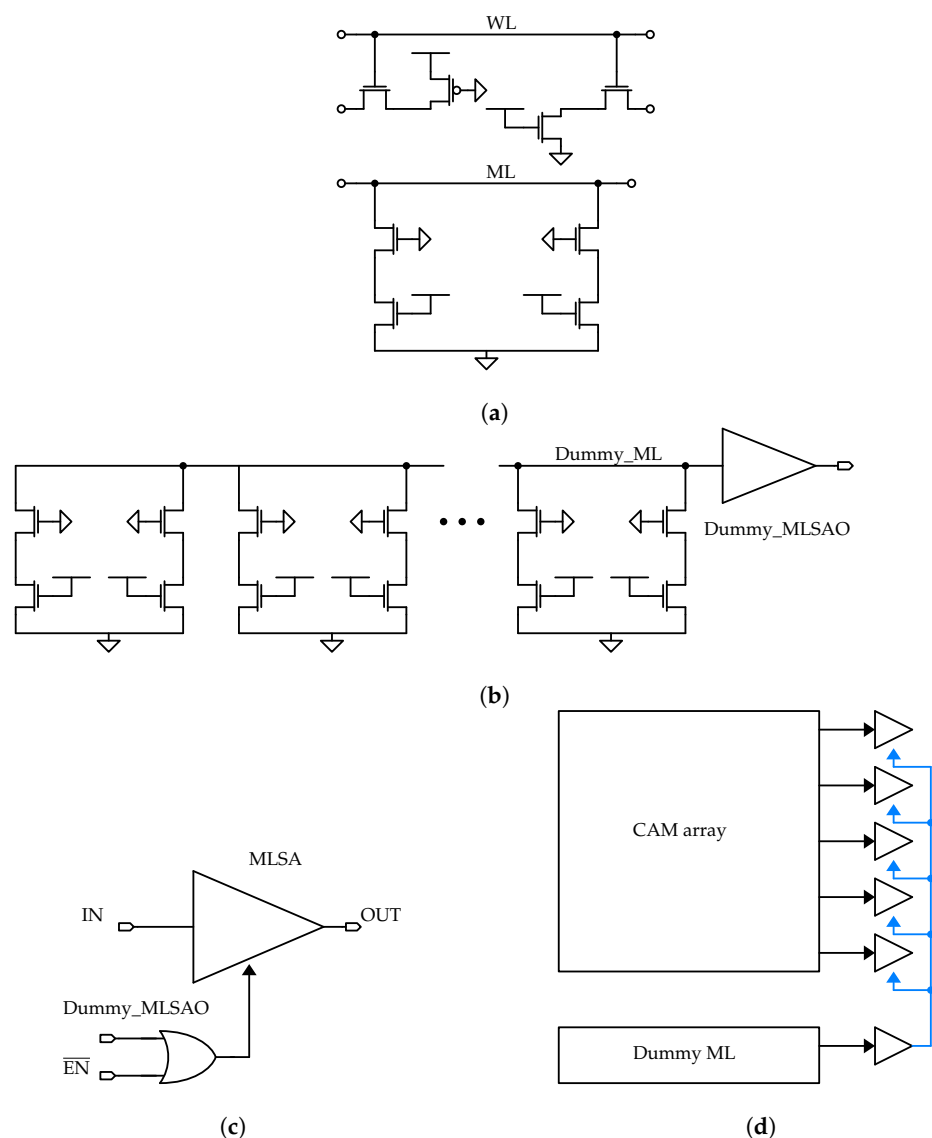


Figure 3. The dummy matchline scheme. A dummy MLSA is used to disable the current sources of the real MLSAs to save power during a search operation. (a) The dummy cell of the CAM. Only the matchline and wordline transistors are kept in the circuit, together with a dummy SRAM core that stores a logic '1'. (b) The dummy matchline. The dummy cells are arranged in a matchline of length equal to the memory width one, and that is connected to a dummy MLSA. Part of the dummy CAM cell is omitted for the sake of clarity. (c) The output of the dummy MLSA is used to disable the other MLSAs: as soon as the dummy MLSA output changes, it means that the time needed for the match sensing has passed, and the current sources of the real MLSAs can be disabled. In order to achieve this, an OR gate is added inside each MLSA, and its output is used as internal enable signal. (d) The output of the dummy MLSA is connected to all the other MLSAs. The position of the dummy matchline is critical: since the dummy MLSA determines the timing of the memory, the line position has to be associated to the worst case for the sensing delay.

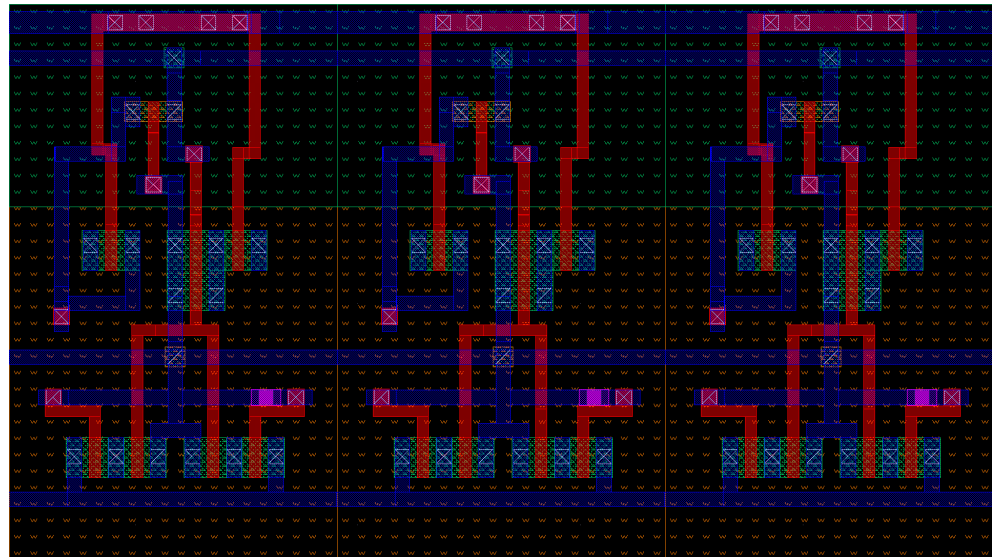


Figure 4. A layout section of the dummy line for the CAM architecture.

3. The LiM Array

As a case of study, an architecture [27] for in-Memory maximum/minimum computation designed by the authors is chosen, since it combines a general-purpose modification (bit wise in-Memory AND logic operation) with a special-purpose Near-Memory logic circuitry for the maximum/minimum computation.

Therefore, it represents a proper case of study to quantify the impact of this particular approach to in-Memory computing, which is the goal of this work. The architecture is not intended as a CPU substitute, but as a hardware accelerator for particular tasks, such as the maximum/minimum computation or bit-wise memory operations.

The algorithm for in-Memory maximum/minimum value search is based on the bitwise AND operation. All the words stored in memory are AND-ed with an external word called “mask vector”, which is put on the memory bitlines one bit at a time until the whole word width is scanned; the results of these AND operations are then elaborated by the Near-Memory logic to choose the words to be discarded at each step, until only the maximum/minimum value remains.

Consider the case in which unsigned words are stored in memory and the maximum value among these has to be found: in this case, at each iteration, only one bit of the mask is set to ‘1’ starting from the MSB, and all the words for which the result of the AND is equal to ‘0’ are discarded. In fact, if the bit of a word A is equal to ‘0’, while the same bit of a word B is equal to ‘1’, then B is larger than A ; hence, A is discarded from the search.

An example of the maximum search for unsigned words is provided in Figure 5. At each step, depending on the result of the AND operation, a word is discarded, until the whole memory width is processed or only one word remains. For minimum search and/or signed words, as well as other types of data encoding, it is enough to change the bits of the mask and to program the Near-Memory logic.

Mask	1	0	0	0	AND	Mask	0	1	0	0	AND
Word #0	0	1	1	1	0	Word #0	0	1	1	1	-
Word #1	1	0	1	0	1	Word #1	1	1	1	0	1
Word #N	1	1	0	0	1	Word #N	1	1	0	0	1

Mask	0	0	0	1	AND	Mask	0	0	1	0	AND
Word #0	0	1	1	1	-	Word #0	0	1	1	1	-
Word #1	1	1	1	0	0	Word #1	1	1	1	0	1
Word #N	1	1	0	0	-	Word #N	1	1	0	0	0

Figure 5. All the words are scanned through a bitwise AND with an external word called “mask vector”. The ones for which a logic ‘0’ is obtained as result, are discarded; the remaining ones at the end are selected as maximum values, and a priority mechanism can be applied to choose among these. In the example, the selected word is highlighted in green.

The memory architecture consists of a standard NOR CAM, as the one presented in Section 2, to which the capability to perform the AND operation is added; the circuit is presented in Figure 6. It has to be remarked that, in this work, only the LiM array schematic is presented, without including the Near-Memory logic circuitry that is described in [27].

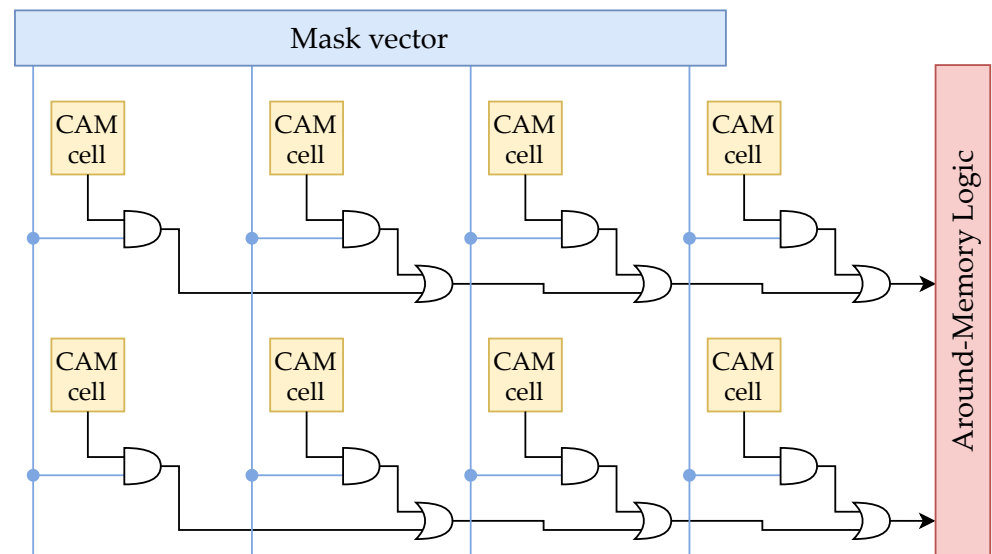


Figure 6. The LiM array. It consists of CAM cells to which the AND logic function capability is added; the results of the AND operations are OR-ed on the rows and provided to the Near-Memory logic.

As previously explained, the AND operations between the memory content and the mask are performed in parallel on all the rows, one bit at a time; then, the results of these are collected through OR operations on the rows by the sense amplifiers and provided to the peripheral logic circuits. Hence, the single cell includes two additional functionalities: AND and OR. The AND is a proper logic gate inserted into the cell, while the OR is implemented through a wired-OR line across the row, which result is handled on the periphery by a sense amplifier, denoted with “ANDSA”. The AND line schematic is depicted in Figure 7.

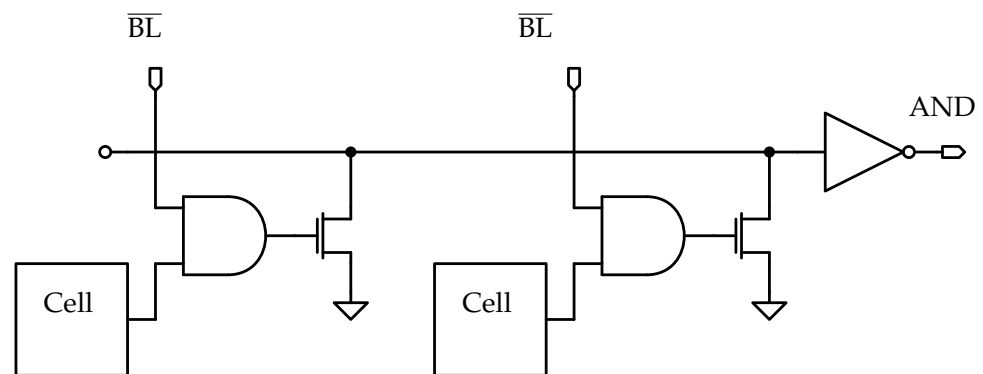


Figure 7. The AND line. The AND gates outputs are connected to a wired-OR line through pull-down transistors. The signal on the line is then inverted to the AND result. The AND gates results are selected through \overline{BL} by the mask.

To select the column on which the AND operation has to be performed, all the bits of the mask vector have to be set to '0' except the one corresponding to the selected column: in this way, all the AND operations on the other columns give '0' as result, disabling the corresponding pull-down transistors, and the logic value sensed on the line depends exclusively on the output of the AND on the selected cell. This can be clarified with an example. Denoting with D_i the content of the cell on the i -th column, with M_i the mask bit in the same position and with O the result obtained in output, when considering the bitwise AND implemented on the row, Equation (1) is obtained:

$$O = \sum_{i=0}^{N-1} D_i \cdot M_i \quad (1)$$

A non-rigorous notation is used in the equation, associating to the sum sign '+' the OR operation and the product sign '.' to the AND one. Indicating with the index j the selected column, the formula can be rewritten in the following way:

$$M_i = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

$$\begin{aligned} O &= \sum_{i=0}^{N-1} D_i \cdot M_i \\ &= \sum_{i=0, i \neq j}^{N-1} D_i \cdot M_i + D_j \cdot M_j \\ &= \sum_{i=0, i \neq j}^{N-1} D_i \cdot 0 + D_j \cdot 1 \\ &= D_j \end{aligned}$$

Hence, the output of the OR operation is determined only by the selected cell content.

The AND logic function is implemented by embedding a logic gate inside the memory cell. Three variants of this are presented:

- a dynamic CMOS logic AND gate.
- a static CMOS logic AND gate.
- a special purpose AND gate, designed appositely for the algorithm to be implemented in order to reduce as much as possible the cell area.

3.1. Dynamic CMOS Logic AND

In Figure 8a, the circuit of the AND gate is shown. It takes in input the negated values of the cell content, \overline{D} , the mask bit on the bitlines, \overline{BL} , and an additional external signal, \overline{PRE} , used to precharge the output node of the gate, O. It can be noticed that an AND function is obtained without adding an inverting stage on the output of the inner gate: since the negated values of the cell content and mask bit are available, one can use De Morgan's laws to avoid the inverting stage. In fact, since the gate in Figure 8a takes in input \overline{D} and \overline{BL} , the logic NOR between these, implemented by the logic gate, can be rewritten in the following way:

$$\overline{\overline{D} + \overline{BL}} = D \cdot BL$$

Hence, the inverting stage is not needed to implement the AND function. This logic gate is embedded in the cell, obtaining the circuit show in Figure 8b. One can notice that a pull-down transistor is added on the output of the AND gate and connected to the row line.

The AND line is an implementation of dynamic CMOS logic: the line is precharged to the logic '1' and then, if one of the pull-down transistors connected to it is enabled, discharged during the evaluation phase. In order to properly carry out the precharge phase, all the pull-down transistors must be disabled. This is usually achieved by adding a footer transistor on the source of the pull-down of each cell, that is disabled during the precharge phase through a dedicated row signal, preventing the pull-downs from discharging the line independently of the output values of the AND gates. A possible circuit is highlighted in Figure 9a.

In this work, a different approach is used to disable the pull-down transistors during the precharge phase: the same current-saving sensing scheme of the CAM is adopted for the AND line. In this way, since the line is pre-discharged and not pre-charged, there is no need to disable the pull-downs and, hence, additional transistors and signals are not required, allowing for smaller cell and row footprints.

A circuit is proposed in Figure 9b. A truth table for the logic gate, which takes into account the implementation of the current-saving scheme, is shown in Table 1.

Table 1. The truth table of the dynamic AND cell of Figure 8b. One can notice that, using the current-saving scheme, the \overline{AND} output is charged to '1' when AND is discharged to '0', while it remains at the ground voltage when $AND = '1'$.

D	BL	\overline{D}	\overline{BL}	AND	\overline{AND}
0	0	1	1	1 → 0	0 → 1
0	1	1	0	1 → 0	0 → 1
1	0	0	1	1 → 0	0 → 1
1	1	0	0	1	0

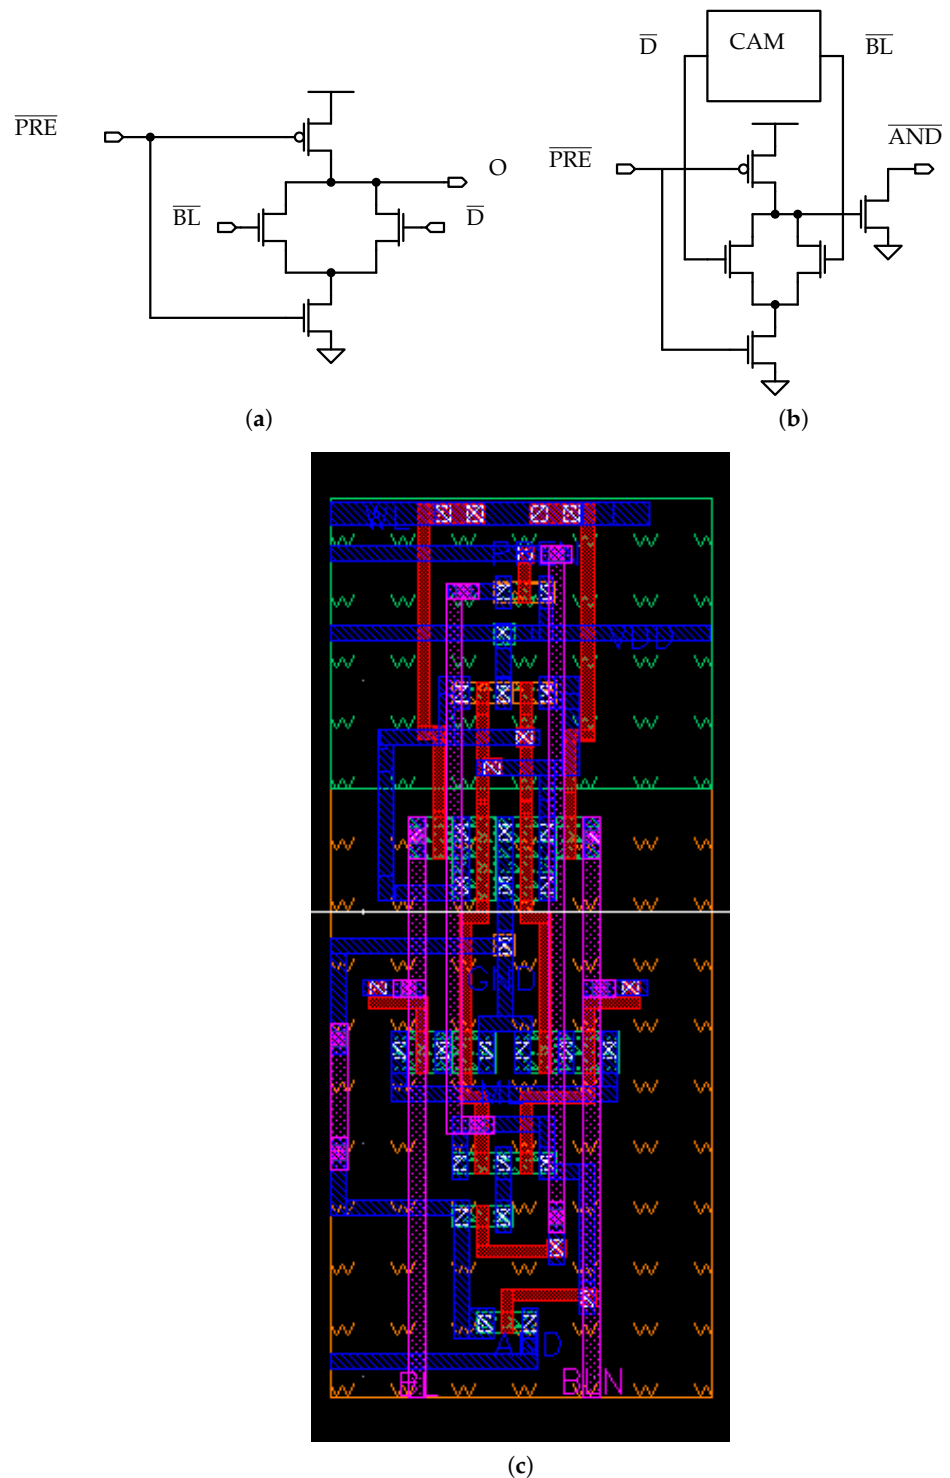


Figure 8. The dynamic AND gate and the cell in which it is integrated. (a) The dynamic AND gate. By using the negated values of the cell content, \bar{D} , and the mask bit, \bar{BL} , it is possible to take advantage of boolean logic laws to obtain an AND gate without adding an inverting stage on the output. (b) The memory cell that embeds the dynamic AND gate and the pull-down transistor of the AND line. It can be noticed that the output of the AND gate is negated using a single pull-down transistor, which output corresponds to the \overline{AND} signal associated to the line. (c) The cell layout.

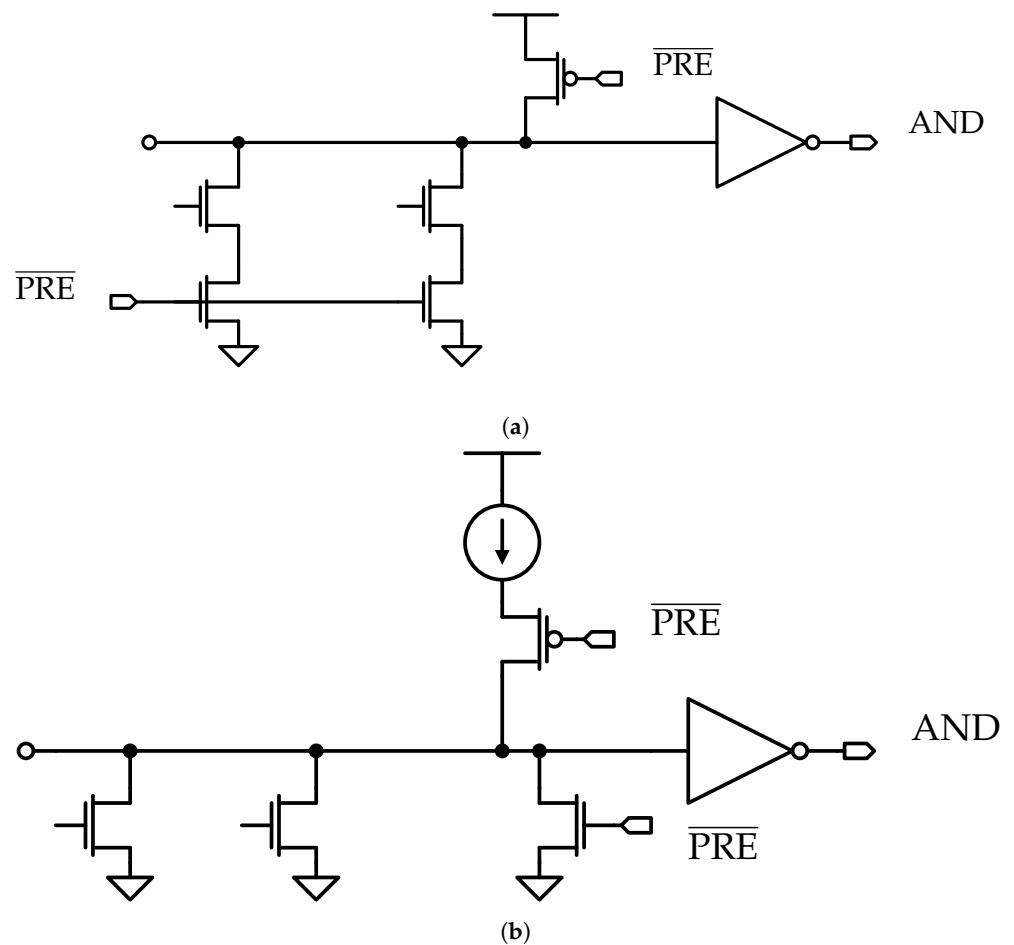


Figure 9. Standard and current-saving schemes. (a) A standard precharge line. The line is precharged to the logic ‘1’ using a pull-up transistor, while all the pull-downs are disabled using footer transistors; then, these are enabled by deactivating $\overline{\text{PRE}}$ during the evaluation phase. (b) The current-saving line. In this scheme, footer transistors are not needed for disabling the pull-downs, since the line is pre-discharged instead of being pre-charged; then, during the evaluation phase, the line gets charged if there are not conducting pull-downs.

3.2. Static CMOS Logic AND

A second cell embedding a static CMOS logic AND gate is proposed. The circuits of the gate and the cell are depicted in Figure 10.

The static AND gate is presented in Figure 10a. With respect to the dynamic AND (Section 3.1), a larger cell footprint is required, since the additional pMOS transistors have to be sized with a width larger than the precharge transistor in Figure 8a, following the rules of standard microelectronic design [35]. However, the addition of these allows to remove the precharge signal $\overline{\text{PRE}}$ of Figure 8a, which is required for the dynamic logic functioning. The gate is embedded in the memory cell, as it is shown in Figure 10b, and its output is connected to the pull-down transistor of the AND line. The truth table for the gate is the same of the dynamic AND cell, which is reported in Table 1, except for the fact that, for the static cell, the AND output signal is a static CMOS one.

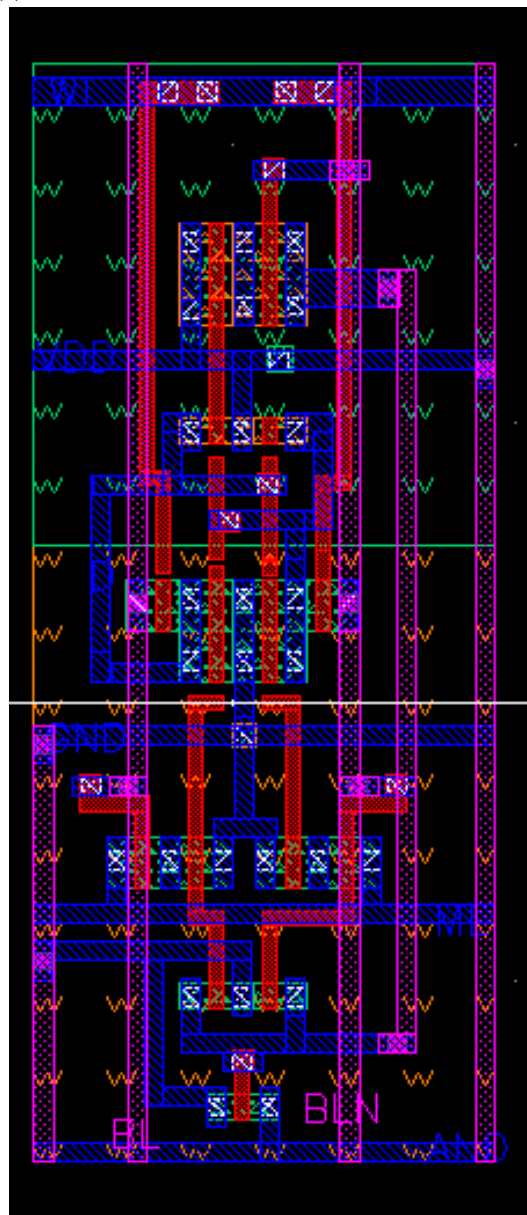
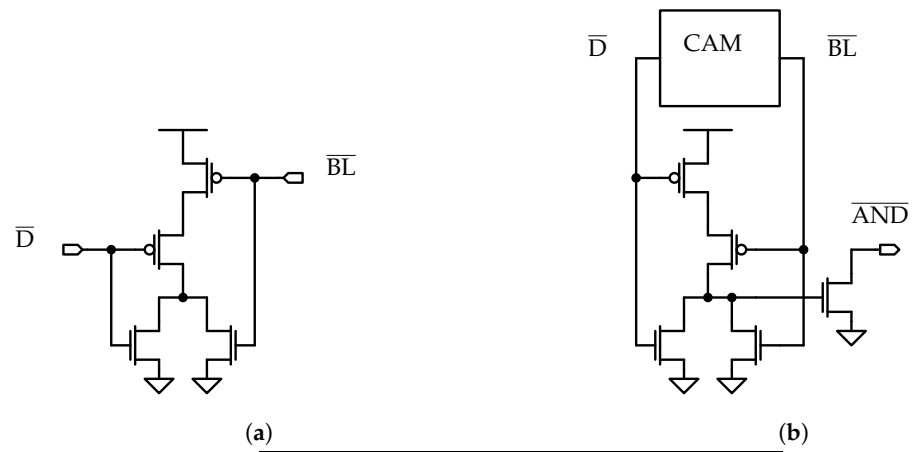


Figure 10. The static AND gate and memory cell. (a) The static CMOS AND gate. (b) The memory cell that embeds the static AND gate. (c) The cell layout.

3.3. Special Purpose AND

A third variant of the cell is proposed. The objective of this cell design is to reduce as much as possible the cell area overhead resulting from the addition of the AND gate, by making design choices tuned on the characteristics of the algorithm. The schematics of the gate and the cell are depicted in Figure 11.

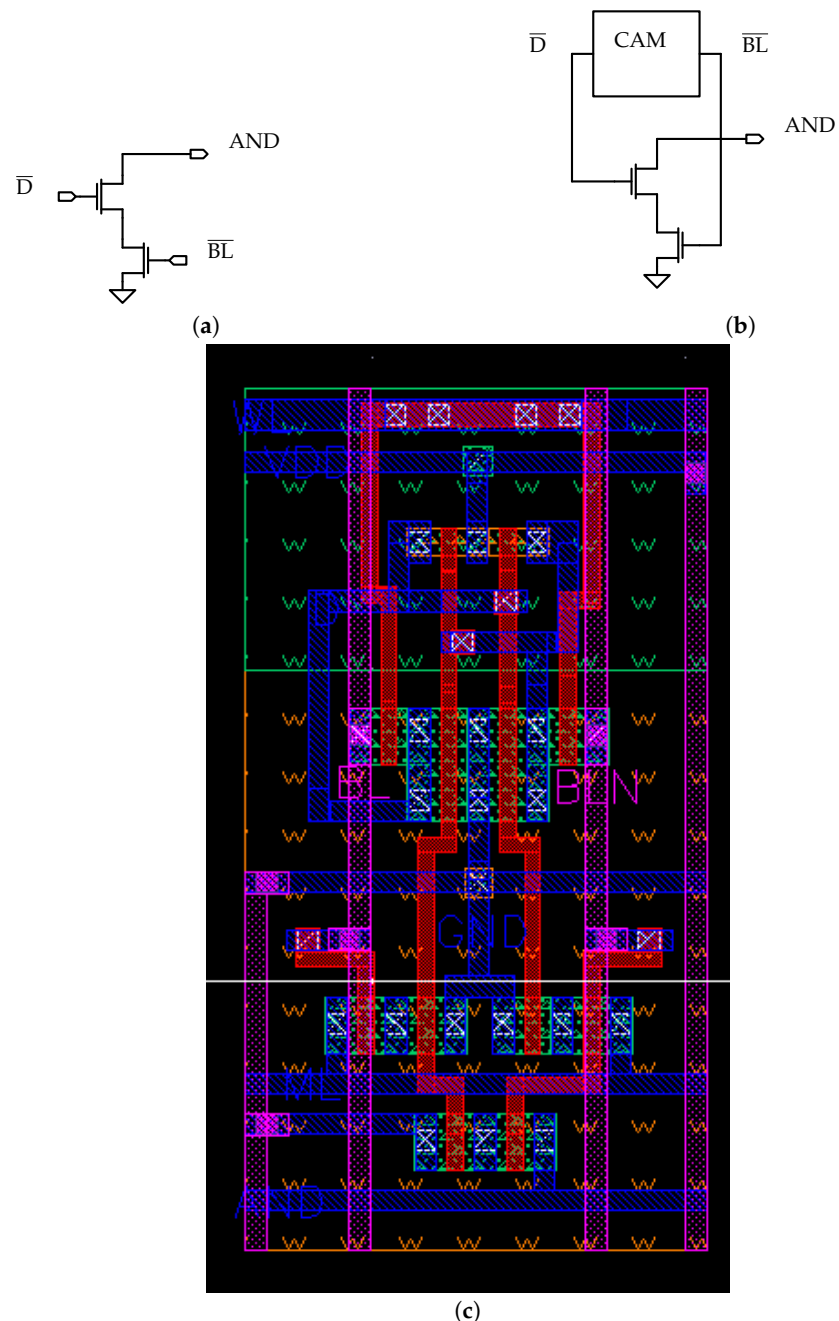


Figure 11. The special-purpose AND gate and memory cell. Taking into account the algorithm characteristics, it is possible to implement the AND operation by using only two additional transistors. (a) The special-purpose AND gate. (b) The memory cell that embeds the logic gate. (c) The cell layout.

As it is highlighted in Figure 5, the mask vector is used to select a memory column at each iteration by setting the corresponding bit in the mask to '1', while all the other cells are disabled. Since the AND operation is computed between a bit equal to '1' and the cell content, the result of this is determined by the cell, as it is shown in Equation (1); hence, it is more a selection operation than an AND one. For this reason, the cell circuit can be simplified to only implement the cell selection operation using the bitlines on which

the mask vector is put, instead of a proper AND function, and to allow the selected cell content to be reflected on the AND line. This result can be achieved by connecting a single pull-down transistor with the input on the cell content and the output on the AND line, as it is depicted in Figure 11a.

Since the cell has to be selected only when the mask bit M is equal to the logic '1' (i.e., $BL = '1', \overline{BL} = '0'$), it should be disconnected from the AND line when $M = '0'$ (i.e., $BL = '0', \overline{BL} = '1'$); hence, it would be enough to add a footer transistor, which gate is connected to BL , on the source of the pull-down one in order to disable this. However, since the static (Figure 10a) and dynamic (Figure 8a) gates have one of their inputs connected to \overline{BL} instead of BL , a different encoding of the mask vector is used in this case, using the logic '0' as active value for the mask bit instead of the logic '1'; in this way, the footer transistor in Figure 11a can be connected to \overline{BL} ; therefore, the three variants are equivalent in terms of connections to the memory signal lines and, hence, can be properly compared.

For what concerns the pull-down transistor, its gate is connected to the output of an AND logic gate in the static (Figure 10a) and dynamic (Figure 8a) gates; in Figure 11a, instead, it is connected to the negated value of the cell content \overline{D} ; in fact, once the cell is selected, the algorithm needs only to know if the cell content is equal to '0' or '1', and the latter can be connected directly to the pull-down transistor gate. In this way, when $D = '1'$ ($\overline{D} = '0'$), the AND logic gate is disabled, the line is charged to the logic '1'; when $D = '0'$ ($\overline{D} = '1'$), the pull-down transistor is enabled, the line is not charged and a logic '0' is sensed.

One can notice that the output pin of the cell is denoted with \overline{AND} instead of AND , in Figure 11b: this is due to the fact that the AND result is not inverted by the pull-down transistor. In fact, the pull-down transistors of the unselected columns are disabled using the mechanism presented in Figure 9b and, hence, the AND result on the selected column can be directly reported on the line. If the selected cell content D_i is equal to '1', the line is charged and $D_i \cdot M_i = '1'$ (M_i is the active mask bit) is registered in output; otherwise, the line does not get charged and $D_i \cdot M_i = '0'$. Hence, there is no need for an additional separation stage between cell core and AND line, while there is for the static and dynamic implementations of Figures 8b and 10b, respectively, which logic gates outputs have to be disconnected from the line when the corresponding cells are not selected. The truth table for the special-purpose AND cell of Figure 11b is shown in Table 2.

Table 2. The truth table of the special-purpose AND cell of Figure 11b. When evaluating this function, one needs to remember that \overline{BL} is not a proper data signal but a selection one that allows to report the cell content D on the line. Every time $\overline{BL} = '0'$, the AND logic gate is disabled and the line is charged to '1' (in particular, the pull-down is prevented from discharging the line in the case in which $D = '0'$).

D	BL	\overline{D}	\overline{BL}	\overline{AND}
0	0	1	1	0
-	1	-	0	$0 \rightarrow 1$
1	0	0	1	$0 \rightarrow 1$

The special-purpose cell in Figure 11b is characterised by the lowest area overhead (lowest number of additional transistors) among the cells. However, these are able to perform a proper AND logic operation, which can be useful for implementing other algorithms; nevertheless, in the special-purpose cell circuit it is demonstrated that, with proper optimisations, it is possible to greatly reduce the area overhead introduced by the logic circuits.

The dynamic and static cells, in Figures 8b and 10b respectively, are characterised by the same number of transistors, but the static one occupies a larger area due to the pull-up pMOS transistors in the logic gate, that are much larger than the precharge pMOS of the dynamic cell; however, the static cell does not require the precharge (\overline{PRE}) signal for its functioning, which leads to smaller cell and row areas.

3.4. Dummy Line Sensing Scheme

For the LiM array, the same dummy line sensing scheme of the CAM is adopted: dummy cells are used to create a dummy memory line that acts as reference for all the AND sense amplifiers (ANDSAs).

In Figure 12, the dummy cells for the LiM variants are presented:

- in Figure 12a, the dummy cell for the dynamic logic version is depicted. In this gate, two row signals are connected to each cell: the AND line $\overline{\text{AND}}$ signal and the precharge signal $\overline{\text{PRE}}$; for this reason, the transistors connected to these signals have to be included.
- in Figure 12a,b, the static and special-purpose variants are presented. Since these do not require an additional row signal, only the AND line pin is present in the circuit.

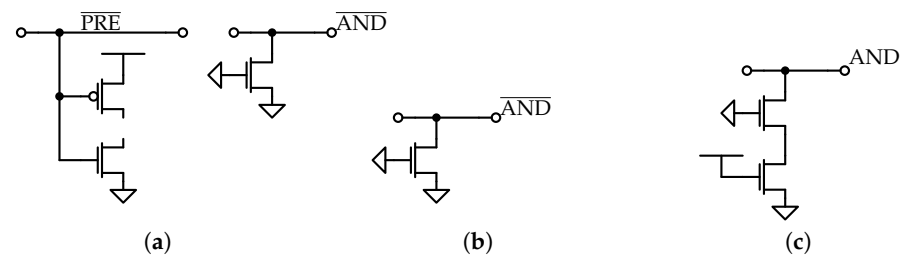


Figure 12. The dummy cells of the LiM array. These are used to mimic a dummy memory row, which is sensed by a special sense amplifier that drives the other ones in order to reduce the overall energy consumption involved in the AND operation, which is performed on the whole array. (a) The dynamic AND dummy cell. (b) The static AND dummy cell. (c) The special-purpose AND dummy cell.

4. Memory Arrays Characterisation

The cells are organised in memory arrays in order to evaluate their performance. The memory circuits are simulated for different values of height and width of the array, in order to obtain measurements valid for a wide range of memory sizes. All the simulations are performed at schematic level in Cadence Virtuoso, using the SPECTRE simulation engine.

In order to take into account the interconnections parasitics contributions, the layouts of the dummy rows and columns are produced and included in the simulated netlist.

In particular, 32-bits wide rows and columns are used as basic blocks to create the array: their layouts are extracted and converted in netlists which are, then, included in the testbench.

When considering the read operation, the distances of the cell to be read from the wordline driver and the sense amplifier have to be taken into account to measure how much the cell position affects the performance. Consider the schematic shown in Figure 13:

- when activating the wordline for selecting a cell, the farthest this is from the driver (i.e., on the last column in Figure 13), the larger the selection delay results to be, due to the higher capacitive-resistive load that the driver has to supply; hence, the read delay associated this cell is the largest possible in the array.
- when sensing the bitlines with the sense amplifier (SA), the farthest the cell is from the SA inputs (i.e., on the first row in Figure 13), the longer the time needed by the cell to generate a voltage variation on the SA pins is.

For these reasons, the cell to which the worst case read delay is associated is the one on the first row and last column in Figure 13 (highlighted in blue), and the read operation performance is evaluated on this cell. For what concerns the worst case for the write operation, a similar analysis can be conducted, referring to the schematic in Figure 13:

- for the wordline activation and cell selection, the considerations made for the read operation hold true: the cell to which the largest selection delay is associated is the one on the last column.

- when putting the datum to be written on the bitlines, to evaluate the worst case one needs to consider the farthest cells from the bitlines drivers outputs. In Figure 13, these are the ones placed on the first row.

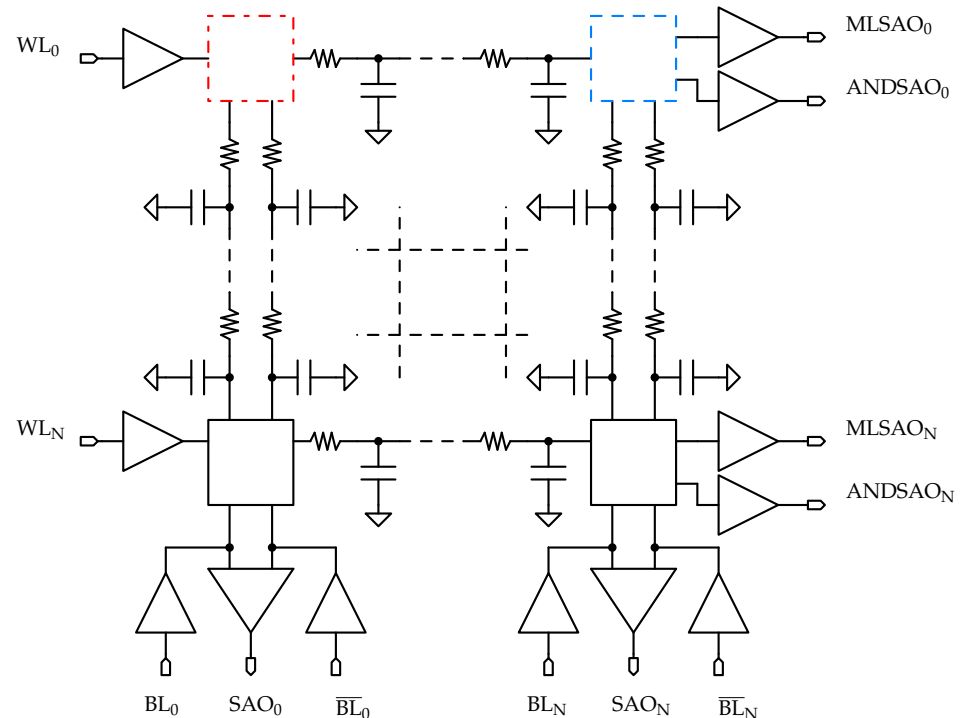


Figure 13. Worst case delays for each memory operation. Most of the memory cells are omitted for the sake of clarity, and the interconnections are represented by the RC circuits, that are substituted by the extracted rows/columns netlists in the testbench. The cell associated to the read and write operations, highlighted in blue and denoted with a dashed trait, is the farthest one from wordline and bitlines drivers, and sense amplifier; the cell associated to the worst case for the search and AND operation, highlighted in red and denoted with a dashed-and-dotted trait, is the farthest one from the MLSA and ANDSA.

For these reasons, the cell associated to the worst case sensing delay for the write operation is the one on the first row and last column (highlighted in blue) in Figure 13.

For what concerns the AND and search operations, consider the schematic in Figure 13: since both MLSA and ANDSA are placed at the end of the row, the farthest cell from these is the one on the first column, highlighted in red. Hence, to this cell it is associated the worst case for both AND and search operations. The row position does not affect the performance of the search and AND operations, even if these are associated to the bitline drivers: this is due to the particular sensing scheme employed for the architecture. In fact, since with the current-saving scheme the pull-down transistors of the cells do not require to be disabled during the pre-discharge phase, one can load the mask vector on the bitlines during this cycle, so that all the cells are already configured before the evaluation phase; in this way, the performance of the search and AND operations do not depend on the distance of the row from the bitline drivers outputs.

Since the cells required to properly test the memory array are very few, it is not necessary to include all the memory cells in the simulation testbench: the array is reduced to a worst case model, based on the considerations made before, by removing all the cells that are not tested from the array, which leads to shorter simulation time and, hence, faster tuning of the design parameters; consecutively, the circuit model depicted in Figure 14 is derived and used during the simulations.

Only two memory lines are considered in the model: the first row and the last column. This is due to the fact that the critical cells for all the memory operations are placed on these lines; moreover, since only two cells are tested, the remaining ones can be replaced with dummy versions, which circuits are depicted in Figure 14.

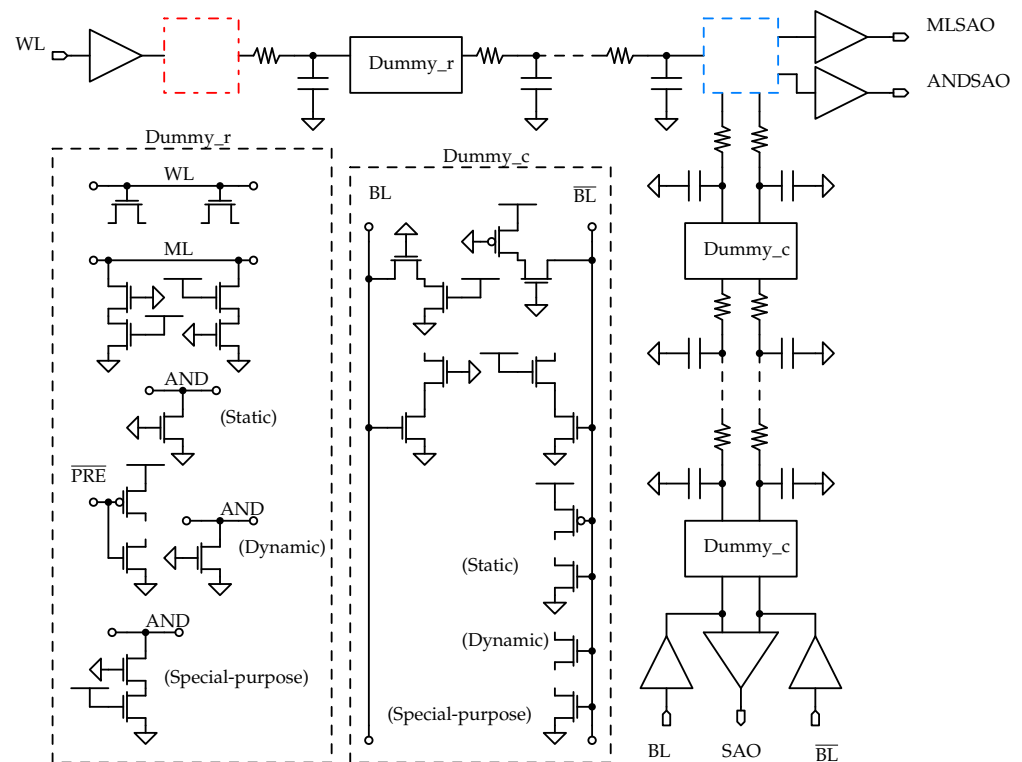


Figure 14. The array model. The first row correspond to the top one, while the first column with the leftmost one. Only the critical cells for the read, write, search and AND operations are actually included in the array, as layout-extracted circuits, while all the others are substituted by dummy rows and columns, extracted from the layout, that contain only the significant transistors (i.e., the ones connected to the row/column signals in the original array).

The dummy cells are distinguished in row and column ones: in the dummy row cells, only the transistors that are connected to the row signals (wordline; matchline; AND line; precharge line only for the dynamic AND cell) are included in the cell circuit; in the dummy column ones, instead, only the transistors that are connected to the bitlines are kept. In this way, the presence of a memory cell on the line signals is still taken into account while many transistors are removed from the circuit, which leads to a big reduction of the simulation time for large memory arrays.

In Cadence Virtuoso, the testbench shown in Figure 15 is employed. This schematic is valid for the LiM array, but it can be simplified and adapted to the CAM and SRAM architectures, since the LiM memory embeds these functionalities, by removing some blocks and substituting the cells circuits.

In Figure 15, it can be noticed that the bitline drivers are included only for the last column, since only on this line the read and write operations and tested; for the first column, instead, ideal switches and voltage generators are employed to modify the cell content, since only row operations, such as the AND and search ones, are tested on it.

In the schematic shown in Figure 15, one can also notice that a block called “dummy load” is added on the output of each dummy sense amplifier: these blocks are needed to emulate the presence of all the sense amplifiers of the rows of an actual memory array. As it is discussed in Section 3, the dummy sense amplifier has to drive all the OR logic gates embedded in each real sense amplifier; since in the model presented in Figure 14 only

one row is equipped with MLSA and ANDSA, the other rows SAs have to be modeled to take into account their influence on performance in an actual memory array. For this reason, dummy loads made by OR gates input sections are connected to the output of the sense amplifiers.




Figure 15. The testbench. The wires related to the CAM and AND functionalities are highlighted in blue and orange, respectively, for the sake of clarity.

The circuit of the dummy load block is shown in Figure 16. It consists of multiple OR logic gates which share the same input, and the number of OR gates coincides with the number of rows in the array. These are not actual gates: only the transistors connected to the input are included, in order to reduce as much as possible the number of elements in the testbench netlist.

Some additional blocks are shown in Figure 15: the precharge circuit is used to precharge the bitlines before a read operation; the “Delay SA” circuit is used to delay the enable signal of the sense amplifier used to test the read operation, since a voltage latch SA [33] is employed.

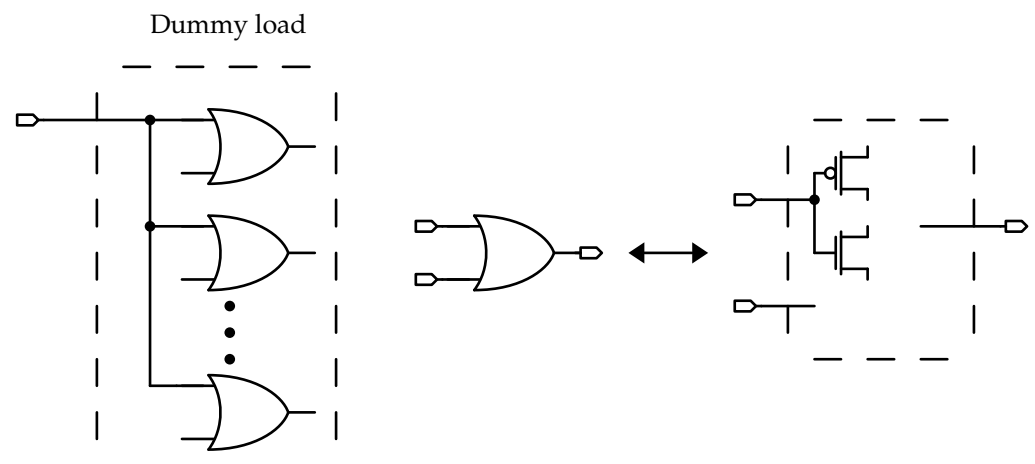


Figure 16. The dummy load for the dummy SA. This is used to emulate the input sections of multiple OR gates, which are embedded in each real MLSA/ANDSA, in order to take into account their influence on the sensing performance in the array.

5. The Simulation Framework

To characterise large memory arrays, a scripting approach is adopted, generating the circuit netlists automatically after an initial by-hand characterisation of the design.

The approach adopted for the simulation of large arrays is presented in Figure 17, and it consists of the following steps:

- the memory array and the sensing circuitry are designed by-hand and characterised by simulating small arrays (32×32 cells).
- the cells and rows layouts are produced and extracted. 32-bits wide rows and columns are used as basic blocks to create the final array.
- after the circuits netlists have been extracted, a script is written, following precise guidelines, to make the circuit parametric with respect to its size (array height and width).
- a script is used to generate a parametric Cadence Virtuoso testbench that allows to characterise the circuit for arbitrary values of width and height, by using the SPECTRE simulation engine.
- the input stimuli of the testbench are automatically generated, starting from the operations sequence to be simulated provided by the user.
- the circuit is simulated using the SPECTRE engine of Cadence Virtuoso.
- the array performance are extracted by measuring the energy consumption and the delay associated to each memory operation.

In Figure 18, the scripting workflow, called ALiAS (Analog Logic-in-Memory Arrays Simulation), is presented. ALiAS takes in input:

- the netlists of the fundamental blocks, which are the memory cells and the sense amplifiers, that have to be designed by-hand.
- the desired characteristics for the array to be simulated: type (SRAM, CAM, the three LiM variants) and size (width and height).
- simulation parameters for SPECTRE (such as the maximum number of computational threads associated to the simulation, GUI mode etc.).
- the clock period selected for the simulation, which is equal to 1 ns by default.

Given this information, the netlist of the array and the testbench are generated, the SPECTRE simulation is run, performance measurements are extracted (in particular, energy consumption and delay associated to each memory operation) and saved in different formats (bar diagrams and CSV files). With this approach, ALiAS allows to speed up the design and simulation of memory arrays with custom cell topologies at schematic level.

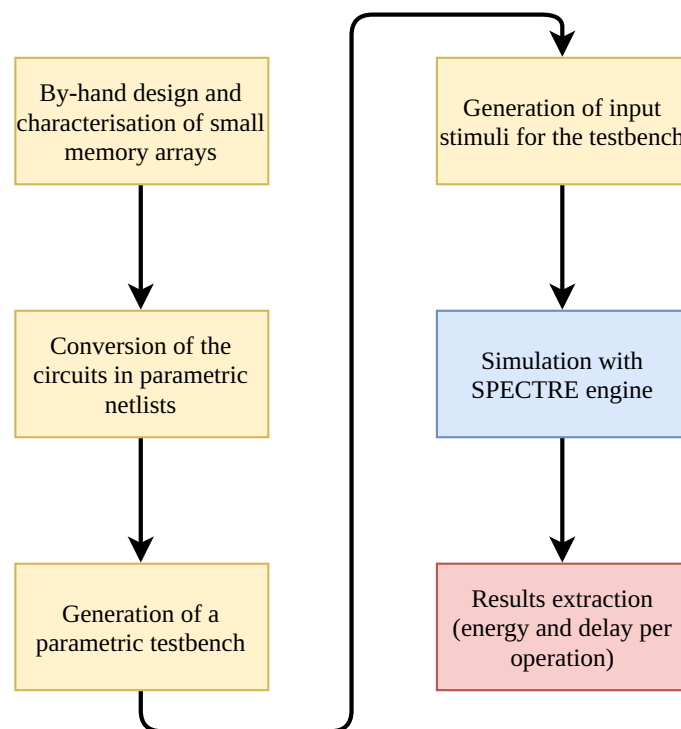


Figure 17. The simulation flow. Starting from the by-hand designed circuits of small arrays, the simulation of large ones is achieved through the algorithm shown in the figure.

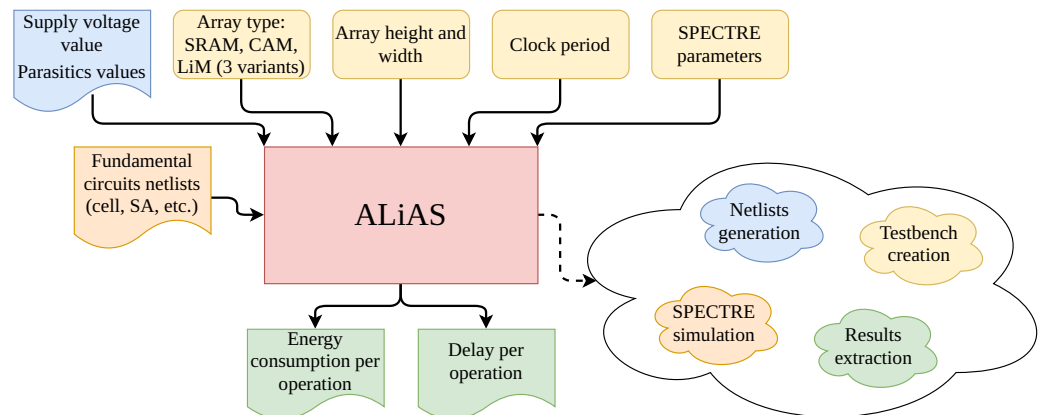


Figure 18. The scripting approach adopted, called ALiAS (Analog Logic-in-Memory Arrays Simulation). Starting from the array characteristics (type and dimensions), the simulation conditions (circuit parameters, clock period, SPECTRE configuration) and the layout extracted netlists of the basic circuits (memory cells, rows and columns), a simulation is performed in SPECTRE and the array performance is evaluated.

6. Results and Discussion

To evaluate the memory arrays performance, energy consumption and latency of each memory operation are extracted from SPECTRE simulations. The energy consumption is measured by integrating the array instantaneous power consumption over each simulation cycle:

$$E_{operation} = \int_{cycle} p(t)dt$$

Each array is simulated with a supply voltage $V_{DD} = 1$ V and a clock period $t_{ck} = 4$ ns using the SPECTRE simulator in Cadence Virtuoso.

In Figure 19, the energy-delays product per operation of each memory array, are presented. Four different memory sizes are considered: 64×64 , 128×128 , 192×192 and 256×256 , intended as rows and columns. These values have been chosen to estimate how the array performance scales with its size, with size values usually adopted in literature for test chips [32,36–38]. In Table 3, the energy-delay products values are shown, using as reference case the 256×256 array of Figure 19.

In the following, each operation is analysed and compared to the others.

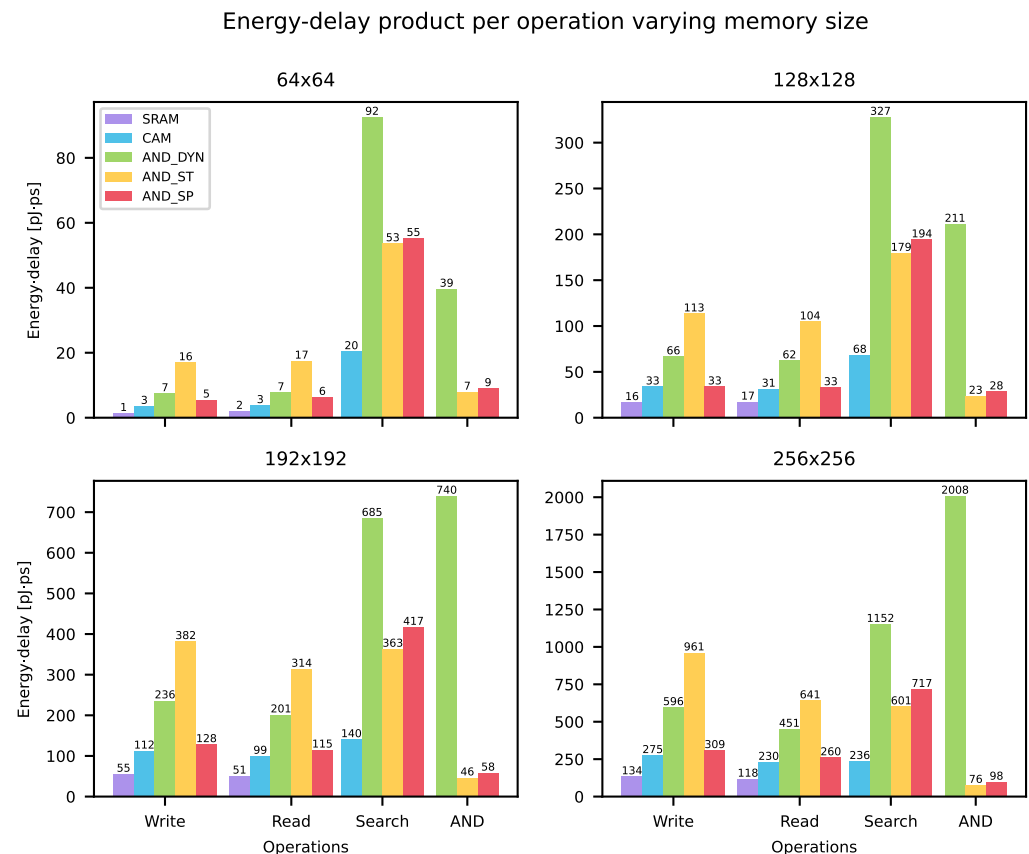


Figure 19. The energy-delay product associated to each memory operation in each array, for different values of the memory size.

Table 3. The energy-products associated to each memory operation, for each memory array. Data are extracted from the 256×256 array of Figure 19, which is used as case study.

Memory	Energy-Delay Products [pJ·ps]			
	Write	Read	Search	AND
SRAM	134	118	–	–
CAM	275	230	236	–
AND SP	309	260	717	98
AND DYN	596	451	1152	2008
AND ST	961	641	601	76

6.1. Read Operation

From Figure 19, one can observe that the LiM (in the figure *AND_SP*, *AND_DYN*, *AND_ST* for special-purpose, dynamic and static AND cells, respectively) and CAM memories perform worse than the SRAM array for every value of the memory size. This is due to the fact that these architectures employ cell circuits that are much more complex (i.e., higher number of transistors, wider transistors and more interconnections) than the SRAM one.

In Table 4, the differences in the energy-delay products associated to the read operation, expressed in percentage, among the arrays, are shown. For instance, for the CAM memory an energy-delay product value 94.41% higher than the SRAM one is measured; for the static AND memory, an energy-delay product value 40.57% higher than the special-purpose AND one is obtained. The data are extracted from the 256×256 array of Figure 19, which is used as case study in the following.

Table 4. Percentage differences in the read energy-delay product among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. Some values are omitted for the avoid ambiguities in the table interpretation (i.e., each percentage value is calculated using as reference the memory on the column and each comparison is made only one time per memory). The data are extracted from Table 3.

Energy-Delay Products Relative Variations for Read					
	SRAM	CAM	AND SP	AND DYN	AND ST
SRAM	—	—	—	—	—
CAM	+94.91%	—	—	—	—
AND SP	+120.34%	+13.04%	—	—	—
AND DYN	+282.2%	+96.08%	+73.46%	—	—
AND ST	+443.22%	+178.69%	+146.54%	+42.13%	—

The differences among the memories performance can be explained by investigating their circuits. In Figure 20, these are depicted showing only the cell transistors connected to the bitlines. In fact, it is well known that to read from a SRAM-like memory cell, one needs to access it through a wordline and to let the cell discharge one of the bitlines to determine its content; the higher the equivalent capacitive load corresponding to the bitlines, the longer the discharge time is, given the same discharge current. Since the bitlines capacitance is determined by the interconnections layout and the transistors connected to these, it follows that the higher the number of the cell transistors linked to the bitlines, the worse the read performance is.

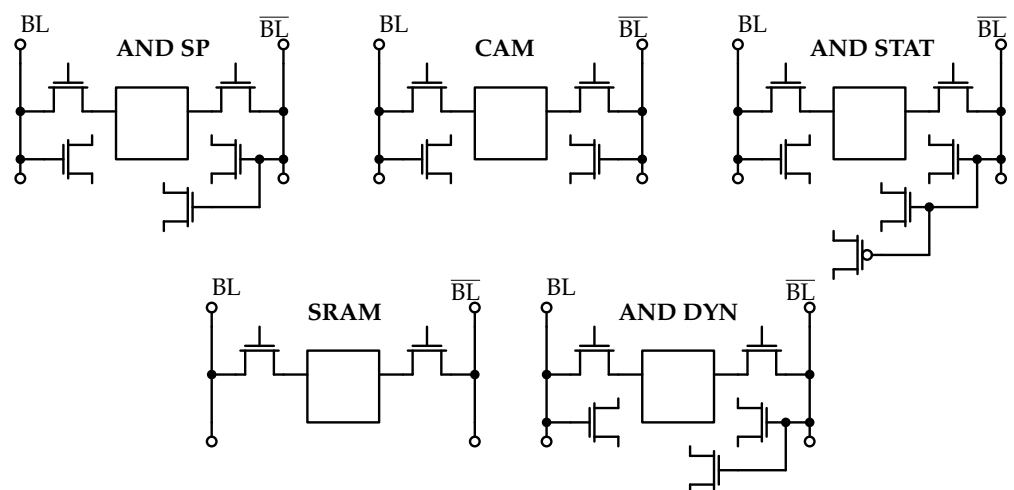


Figure 20. The transistor load on bitlines for each cell type.

Considering the data in Table 4, one can notice that the worst-performing memory is the static AND memory, which is also the one with the highest number of transistors connected to the bitlines (Figure 20). This explains why the best performing memory is the SRAM: being the simplest from a circuital point of view, it has the lowest bitlines capacitance associated to it. Similar considerations can be made to explain the differences among the other cells.

One may notice from Figure 20 that even if the special-purpose and dynamic cell have the same number of transistors connected to the bitlines (in particular, to \overline{BL}), the second one performs worse than the first one; this is because one has to take into account also the layouts of these cells, depicted in Figures 8c and 11c, for the dynamic and special-purpose AND cell, respectively. It can be observed that the dynamic AND circuit is more complex, having a higher number of transistors and interconnection, which lead to more parasitics in the resultant circuit that slow down the cell read operation, increasing also the corresponding power consumption.

6.2. Write Operation

In Table 5, the differences in the write operation energy-delay products, expressed in percentage, among the arrays are shown. The same considerations made for the read operation apply, since write and read performance are both approximately determined by the memory circuit and layout.

Table 5. Percentage differences in the write energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the write energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations for Write					
	SRAM	CAM	AND SP	AND DYN	AND ST
SRAM	—	—	—	—	—
CAM	+105%	—	—	—	—
AND SP	+130.6%	+12.36%	—	—	—
AND DYN	+344.78%	+116.72%	+92.88%	—	—
AND ST	+617%	249.45%	+211%	+61.24%	—

6.3. Search Operation

In Table 6, the differences in the search operation energy-delay products, expressed in percentage, among the arrays, are shown.

Table 6. Percentage differences in the search energy-delay product among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations for Search				
	CAM	AND SP	AND DYN	AND ST
CAM	—	—	—	—
AND SP	+203.81%	—	—	—
AND DYN	+388.13%	+60.67%	—	—
AND ST	+154.66%	−19.30%	−91.68%	—

One can notice that the LiM arrays perform worse than the CAM one in the search operation. This can be explained considering the layout of the cells: being the LiM cells more complex, their search functionalities are affected by more parasitics.

Consider the case of the dynamic AND cell, which layout lower section is shown in Figure 8c. One can notice that the CAM circuitry is placed very close to the AND one; as a consequence, the parasitics values associated with the matchline are increased with respect to the original CAM cell, which leads to higher latency and power consumption for the search operation. Similar considerations can be made for the special-purpose and static AND cells.

It can be observed that, among the LiM arrays, the best performing one for the search operation is the static AND array. This seems counter-intuitive, since the static AND

gate is the most complex one among the AND cells; however, this can be explained by investigating the layout of the cells.

In Figure 21, the lower sections of the static, special-purpose and dynamic AND cells are shown side to side. By considering the AND gates regions in the layouts, which are highlighted in the figure, one can notice that the most complex layout (in terms of the number of transistors and local interconnections) is the dynamic AND one, highlighted in orange, followed by the special-purpose one (there are less transistors but these are wider), highlighted in cyan, and, then, the static one, highlighted in pink. For this reason, the worst performance is associated with this cell.

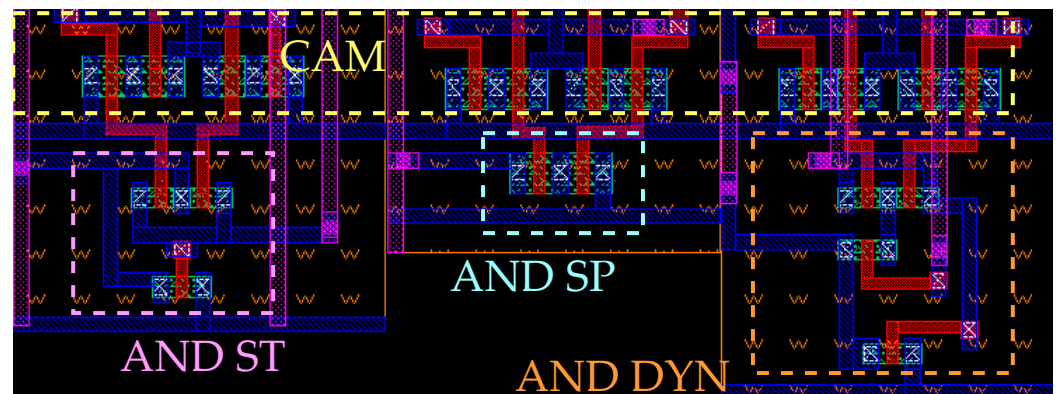


Figure 21. Layout bottom sections of static, special-purpose and dynamic AND cells.

For what concerns the special-purpose cell, its circuit seems to be less complex than the static one, but it should be noted that the transistors of the special-purpose circuit are wider than the ones of the static cell; this leads to larger parasitic capacitances, that lead to a worsening in performance for the search operation, being these transistors connected through the gates to the CAM functionality ones.

6.4. AND Operation

In Table 7, the differences in the AND operation energy-delay products, expressed in percentage, among the arrays are shown.

Table 7. Percentage differences in the AND energy-delay product among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations for AND			
	AND SP	AND DYN	AND ST
AND SP	—	—	—
AND DYN	+1948.98%	—	—
AND ST	−28.95%	−2542.1%	—

One can notice that the best performing array is the static AND one. This can be explained by referring to the cells circuits.

The static cell performs better than the special-purpose one due to its simpler output circuit (Figure 10 for the static AND, Figure 11 for the special-purpose AND): while the static gate has only one transistor connected to the AND line, the special-purpose one has two NMOSFETs in series linked to it; this leads to higher latency and power consumption.

The static AND cell performs better also than the dynamic cell, since the latter is implemented in dynamic CMOS logic, while the first one in static CMOS logic. In fact, considering the circuit of the dynamic AND cell in Figure 8, it can be noticed that, once the sensing of the AND line is enabled through \overline{EN} , it takes a certain amount of time for

the dynamic gate to discharge its output, denoted with AND , and, hence, disable the pull-down. During this time interval, the pull-down is conducting and prevents the AND line, denoted with \overline{AND} , from getting charged by the ANDSA. This leads to an increase in both energy consumption and sensing delay.

Considering the circuit of the static AND cell in Figure 10, one can notice that the output of the AND gate is already at ground voltage before the sensing enabling, for the reasons discussed in Section 4. At the beginning of the AND operation, the pull-down is already disabled, which means that the line starts immediately to get charged, without having any current flowing to ground. Moreover, at each AND execution, all the AND gates invert their outputs to turn off the pull-down transistor connected to the AND line; this leads to a large increase in the energy consumption, as it can be observed from Table 7.

6.5. Comparison among Different Operations

In this section, the operations performed are compared and analysed in relation to each other.

From Figure 19, one can notice that write performance worsens more than the read one, as the array size is increased. This is mainly due to the fact that while a read operation does not imply the complete commutation of one of the bitlines (one of the two lines needs to discharge just enough for the sense amplifier to properly read the cell value), a write one does, since a “strong” logic ‘0’ has to be put on one of the bitlines to force the desired value to be written to the cell; as a consequence, larger energy consumption for the write operation with respect the read one is measured.

In Table 8, the read and write performance, in terms of energy-delay product, are compared in each memory. One can notice that the largest difference between reading and write performance is associated to the static AND memory. This is due to the fact that, as the array size is enlarged, the corresponding bitlines capacitive load increases more than linearly; since the static AND cell is the most complex one, a larger difference in write and read performance is measured for large arrays (e.g., the 256×256 one in Figure 19), while in the other ones a smaller one is obtained. In fact, in Table 8, the write/read discrepancy value follows the cell circuit complexity: the best performing memory is the SRAM, followed by CAM, special-purpose, dynamic and static AND .

Table 8. Percentage differences of the write and read energy-delay products in each memory. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: Write vs. Read						
		Read				
		SRAM	CAM	AND SP	AND ST	AND DYN
Write	SRAM	+13.56%	—	—	—	—
	CAM	—	+19.56%	—	—	—
	AND SP	—	—	+18.85%	—	—
	AND ST	—	—	—	+49.92%	—
	AND DYN	—	—	—	—	+32.15%

In Tables 9 and 10, the comparisons between the search operation and write and read operations, respectively, energy-delay products are reported. One can notice that in all the cases the search operations perform worse than the read/write one of the SRAM array. However, for the static AND and CAM arrays, the search operation is characterised by 16.52% and 59.9%, respectively, lower energy-delay products when compared to the same array write operation; for what concerns the read one, the CAM search operation performs just 2.61% worse, while the static array performs 6.65% better.

From Figure 19, it can be observed that the AND operation performs better than the search one for the static and special-purpose AND arrays. This is due to the fact that

the hardware involved in the AND operation is less complex than the one of the search operation: while in the CAM cell (Figure 2a) there are two pull-down paths of two series transistors connected to the matchline, in the AND cells (Figures 8b, 10b and 11b) there is only one pull-down path. This leads to lower power consumption and latency.

Table 9. Percentage differences between the read and search energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: Search vs. Write						
				Write		
		SRAM	CAM	AND SP	AND DYN	AND ST
Search	CAM	+76.12%	-16.52%	—	—	—
	AND SP	+435.07%	+160.72%	+132.04%	—	—
	AND DYN	+759.7%	+318.91%	+272.81%	+93.29%	—
	AND ST	+348.51%	+118.54%	+94.5%	-91.68%	-59.9%

Table 10. Percentage differences between the write and search energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: Search vs. Read						
				Read		
		SRAM	CAM	AND SP	AND DYN	AND ST
Search	CAM	+100%	+2.61%	—	—	—
	AND SP	+507.63%	+211.74%	+175.77%	—	—
	AND DYN	+876.27%	+400.86%	+343.08%	+115.43%	—
	AND ST	+409.32%	+161.3%	+131.15%	+33.26%	-6.65%

In Table 11, the AND and search operations energy-delay products values are compared. It can be observed that, apart from the dynamic AND case, the AND operation performs always better than the search one. In the dynamic AND case, this does not hold true due to the dynamic CMOS logic implementation of the gate, which leads to the commutation of all the row cells AND gates every time an AND operation is performed. This leads to a large increase in the energy consumption associated with the AND functionality.

Table 11. Percentage differences in the AND and Search energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: AND vs. Search						
				Search		
		CAM	AND SP	AND DYN	AND ST	
AND	AND SP	-140.81%	-631.63%	—	—	—
	AND DYN	+750.85%	+180.05%	+74.3%	—	—
	AND ST	-210.52%	-843.42%	-1415.79%	-690.79%	—

For what concerns the AND operation and the conventional ones, one can notice from Figure 19 that the AND operation, in the static and special-purpose arrays, performs better than both read and write ones in the SRAM array, for an array size equal to 256×256 . This is due to the fact to perform the AND operation there is no need to access the cell content, thank to the additional cell circuitry, which allows for lower latency and energy

consumption; in fact, the SRAM core circuit is highly inefficient, as observed in the previous discussion.

In Table 12 the comparison between AND and write performance is detailed. One can notice that, apart from the dynamic AND case, the AND operation always outperforms the write one, even when comparing it with the conventional SRAM architecture: for the special-purpose case, a 36.7% reduction in the AND energy-delay product is measured with respect the SRAM write one, while in the static AND case the reduction is equal to 76.31%.

In Table 13, the comparison between AND and read performance is analysed. Also in this case, the AND operation always outperforms the read one, apart from the dynamic AND case, even when compared with the SRAM: for the special-purpose AND, a 20.41% reduction in the AND energy-delay product with respect to the SRAM read one is measured; for the static AND case, a reduction of 55.26% is obtained.

Table 12. Percentage differences in the write and AND energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: AND vs. Write						
		Write				
		SRAM	CAM	AND SP	AND DYN	AND ST
AND	AND SP	−36.7%	−180.61%	−215%	—	—
	AND DYN	+1398%	+630.2%	+549.84%	+236.91%	—
	AND ST	−76.31%	−261.84%	−306.57%	−684.21%	−1164%

Table 13. Percentage differences in the read and AND energy-delay products among the arrays. Each value corresponds to the increase, expressed in percentage, in the energy-delay product of the memory on the corresponding row with respect to the one of the memory on the corresponding column. The data are extracted from Table 3.

Energy-Delay Products Relative Variations: AND vs. Read						
		Read				
		SRAM	CAM	AND SP	AND DYN	AND ST
AND	AND SP	−20.41%	−134.69%	−165.31%	—	—
	AND DYN	+1601.7%	+773.04%	+672.31%	+345.23%	—
	AND ST	−55.26%	−202.63%	−242.1%	−1164.47%	−743.2%

This implies that performing an in-Memory operation, such as the AND one, is more convenient from both energetic and latency points of view, even when compared with a conventional SRAM memory. It has to be highlighted that in this analysis the overhead associated to the extraction of the data from the array—i.e., the energy and latency contributions due to the data transfer between the memory and the CPU, and due to the data processing inside the process—is not taken into account; as a consequence, the advantages resulting from the in-Memory approach are being heavily underestimated.

7. Conclusions

In this work, a LiM array with 3 memory cell variants is designed and implemented at physical level in Cadence Virtuoso, by implementing the cells layout and extracting the parasitic netlists from these. The resulting circuit is compared against conventional memory arrays, such as SRAM and CAM ones, by evaluating the overheads associated to the LiM hardware on the standard memory operations.

From the results, an increase in energy consumption and latency is observed for the read and write memory operations in the LiM array (+120.34% and +13.04% for the read operation w.r.t. SRAM and CAM, respectively, in the best case). The results also highlight that the in-Memory processing cost, represented by the energy-delay product associated with the LiM operation, is 55.26% lower than the one associated to the read operation of an SRAM memory, in the best case, even without considering the energy and delay contributions due to the out-of-chip transfer of the data to the CPU. This implies that processing the data directly in memory is much more convenient than extracting them from the array and performing the computations in the CPU, despite the previously discussed drawbacks due to the additional hardware complexity.

These results highlight that Logic-in-Memory arrays, in which the memory cell is modified by adding computational elements to it, are best suited for applications with a low number of reading and write operations and a large number of in-Memory logic operations. These represent a suitable alternative for the design of algorithm accelerators, that can be also used as secondary low-density conventional memories for data storage.

Author Contributions: Conceptualization: M.V. and F.O.; methodology: F.O.; validation: F.O.; writing—original draft preparation: F.O.; writing—review and editing: F.O., M.V., G.M. and G.T.; visualization: F.O. and M.V.; supervision: M.V., G.M. and G.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available upon request to the corresponding author.

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

References

1. Santoro, G.; Turvani, G.; Graziano, M. New Logic-In-Memory Paradigms: An Architectural and Technological Perspective. *Micromachines* **2019**, *10*, 368. [[CrossRef](#)] [[PubMed](#)]
2. Angizi, S.; Fan, D. ReDRAM: A Reconfigurable Processing-in-DRAM Platform for Accelerating Bulk Bit-Wise Operations. In Proceedings of the 2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), Westminster, CO, USA, 4–7 November 2019; pp. 1–8.
3. Simon, W.A.; Qureshi, Y.M.; Rios, M.; Levisse, A.; Zapater, M.; Atienza, D. BLADE: An in-Cache Computing Architecture for Edge Devices. *IEEE Trans. Comput.* **2020**, *69*, 1349–1363. [[CrossRef](#)]
4. Jiang, H.; Huang, S.; Peng, X.; Su, J.-W.; Chou, Y.-C.; Huang, W.-H.; Liu, T.-W.; Liu, R.; Chang, M.-F.; Yu, S. A Two-way SRAM Array based Accelerator for Deep Neural Network On-chip Training. In Proceedings of the 2020 57th ACM/IEEE Design Automation Conference (DAC), San Francisco, CA, USA, 20–24 July 2020; pp. 1–6.
5. Lee, K.; Jeong, J.; Cheon, S.; Choi, W.; Park, J. Bit Parallel 6T SRAM In-memory Computing with Reconfigurable Bit-Precision. In Proceedings of the 2020 57th ACM/IEEE Design Automation Conference (DAC), San Francisco, CA, USA, 20–24 July 2020; pp. 1–6.
6. Jiang, H.; Peng, X.; Huang, S.; Yu, S. CIMAT: A Compute-In-Memory Architecture for On-chip Training Based on Transpose SRAM Arrays. *IEEE Trans. Comput.* **2020**, *69*, 1. [[CrossRef](#)]
7. Rajput, A.K.; Pattanaik, M. Implementation of Boolean and Arithmetic Functions with 8T SRAM Cell for In-Memory Computation. In Proceedings of the 2020 International Conference for Emerging Technology (INCET), Belgaum, India, 5–7 June 2020; pp. 1–5.
8. Yin, S.; Jiang, Z.; Seo, J.; Seok, M. XNOR-SRAM: In-Memory Computing SRAM Macro for Binary/Ternary Deep Neural Networks. *IEEE J. Solid-State Circuits* **2020**, *55*, 1733–1743. [[CrossRef](#)]
9. Agrawal, A.; Kosta, A.; Kodge, S.; Kim, D.E.; Roy, K. CASH-RAM: Enabling In-Memory Computations for Edge Inference Using Charge Accumulation and Sharing in Standard 8T-SRAM Arrays. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **2020**, *10*, 295–305. [[CrossRef](#)]
10. Shin, H.; Sim, J.; Lee, D.; Kim, L.-S. A PVT-robust Customized 4T Embedded DRAM Cell Array for Accelerating Binary Neural Networks. In Proceedings of the 2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), Westminster, CO, USA, 4–7 November 2019; pp. 1–8.
11. Kim, H.; Oh, H.; Kim, J.-J.; Postech, H.K.; Postech, H.O.; Postech, J.-J.K. Energy-efficient XNOR-free in-memory BNN accelerator with input distribution regularization. In Proceedings of the 39th International Conference on Computer-Aided Design, Virtual, 2–5 November 2020.
12. Huang, S.; Jiang, H.; Peng, X.; Li, W.; Yu, S. XOR-CIM: Compute-in-Memory SRAM Architecture with Embedded XOR Encryption. In Proceedings of the 39th International Conference on Computer-Aided Design, ICCAD '20, San Diego, CA, USA, 2–5 November 2020.

13. Ali, M.; Jaiswal, A.; Kodge, S.; Agrawal, A.; Chakraborty, I.; Roy, K. IMAC: In-Memory Multi-Bit Multiplication and Accumulation in 6T SRAM Array. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2020**, *67*, 2521–2531. [[CrossRef](#)]
14. Jiang, H.; Liu, R.; Yu, S. 8T XNOR-SRAM based Parallel Compute-in-Memory for Deep Neural Network Accelerator. In Proceedings of the 2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS), Springfield, MA, USA, 9–12 August 2020; pp. 257–260.
15. Biswas, A.; Chandrakasan, A.P. CONV-SRAM: An Energy-Efficient SRAM With In-Memory Dot-Product Computation for Low-Power Convolutional Neural Networks. *IEEE J. Solid-State Circuits* **2019**, *54*, 217–230. [[CrossRef](#)]
16. Yang, K.; Karam, R.; Bhunia, S. Interleaved logic-in-memory architecture for energy-efficient fine-grained data processing. In Proceedings of the 2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS), Boston, MA, USA, 6–9 August 2017; pp. 409–412.
17. Akyel, K.C.; Charles, H.-P.; Mottin, J.; Giraud, B.; Suraci, G.; Thuries, S.; Noel, J.-P. DRC2: Dynamically Reconfigurable Computing Circuit based on memory architecture. In Proceedings of the 2016 IEEE International Conference on Rebooting Computing (ICRC), San Diego, CA, USA, 17–19 October 2016; pp. 1–8.
18. Jeloka, S.; Akesh, N.B.; Sylvester, D.; Blaauw, D. A 28 nm Configurable Memory (TCAM/BCAM/SRAM) Using Push-Rule 6T Bit Cell Enabling Logic-in-Memory. *IEEE J. Solid-State Circuits* **2016**, *51*, 1009–1021. [[CrossRef](#)]
19. Jiang, Z.; Yin, S.; Seo, J.-S.; Seok, M. C3SRAM: An In-Memory-Computing SRAM Macro Based on Robust Capacitive Coupling Computing Mechanism. *IEEE J. Solid-State Circuits* **2020**, *55*, 1888–1897. [[CrossRef](#)]
20. Jaiswal, A.; Agrawal, A.; Ali, M.F.; Sharmin, S.; Roy, K. i-SRAM: Interleaved Wordlines for Vector Boolean Operations Using SRAMs. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2020**, *67*, 4651–4659. [[CrossRef](#)]
21. Saha, G.; Jiang, Z.; Parihar, S.; Xi, C.; Higman, J.; Karim, M.A.U. An Energy-Efficient and High Throughput in-Memory Computing Bit-Cell With Excellent Robustness Under Process Variations for Binary Neural Network. *IEEE Access* **2020**, *8*, 91405–91414. [[CrossRef](#)]
22. Kim, H.; Chen, Q.; Yoo, T.; Kim, T.T.-H.; Kim, B. A Bit-Precision Reconfigurable Digital In-Memory Computing Macro for Energy-Efficient Processing of Artificial Neural Networks. In Proceedings of the 2019 International SoC Design Conference (ISOCC), Jeju, Korea, 6–9 October 2019; pp. 166–167.
23. Agrawal, A.; Jaiswal, A.; Roy, D.; Han, B.; Srinivasan, G.; Ankit, A.; Roy, K. Xcel-RAM: Accelerating Binary Neural Networks in High-Throughput SRAM Compute Arrays. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2019**, *66*, 3064–3076. [[CrossRef](#)]
24. Saikia, J.; Yin, S.; Jiang, Z.; Seok, M.; Seo, J.-S. K-Nearest Neighbor Hardware Accelerator Using In-Memory Computing SRAM. In Proceedings of the 2019 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED), Lausanne, Switzerland, 29–31 July 2019; pp. 1–6.
25. Agrawal, A.; Jaiswal, A.; Lee, C.; Roy, K. X-SRAM: Enabling In-Memory Boolean Computations in CMOS Static Random Access Memories. *IEEE Trans. Circuits Syst. I Regul. Pap.* **2018**, *65*, 4219–4232. [[CrossRef](#)]
26. Dong, Q.; Jeloka, S.; Saligane, M.; Kim, Y.; Kawaminami, M.; Harada, A.; Miyoshi, S.; Yasuda, M.; Blaauw, D.; Sylvester, D. A 4 + 2T SRAM for Searching and In-Memory Computing With 0.3 V VDDmin. *IEEE J. Solid-State Circuits* **2018**, *53*, 1006–1015. [[CrossRef](#)]
27. Vacca, M.; Tavva, Y.; Chattopadhyay, A.; Calimera, A. Logic-In-Memory Architecture For Min/Max Search. In Proceedings of the 2018 25th IEEE International Conference on Electronics, Circuits and Systems (ICECS), Bordeaux, France, 9–12 December 2018; pp. 853–856.
28. Coluccio, A.; Vacca, M.; Turvani, G. Logic-in-Memory Computation: Is It Worth It? A Binary Neural Network Case Study. *J. Low Power Electron. Appl.* **2020**, *10*, 7. [[CrossRef](#)]
29. Le, M.; Pham, T.K.H.; Truong, S.N. Noise and memristance variation tolerance of single crossbar architectures for neuro-morphic image recognition. *Micromachines* **2021**, *12*, 690. [[CrossRef](#)]
30. Abbas, H.; Abbas, Y.; Hassan, G.; Sokolov, A.S.; Jeon, Y.-R.; Ku, B.; Kang, C.J.; Choi, C. The coexistence of threshold and memory switching characteristics of ALD HfO₂ memristor synaptic arrays for energy-efficient neuromorphic computing. *Nanoscale* **2020**, *12*, 14120–14134. [[CrossRef](#)] [[PubMed](#)]
31. Alimkhanuly, B.; Sohn, J.; Chang, I.-J.; Lee, S. Graphene-based 3D XNOR-VRRAM with ternary precision for neuromorphic computing. *NPJ 2D Mater. Appl.* **2021**, *5*, 55. [[CrossRef](#)]
32. Pagiampzis, K.; Sheikholeslami, A. Content-Addressable Memory (CAM) Circuits and Architectures: A Tutorial and Survey. *IEEE J. Solid-State Circuits* **2006**, *41*, 712–727. [[CrossRef](#)]
33. Kobayashi, T.; Nogami, K.; Shirotori, T.; Fujimoto, Y. A current-controlled latch sense amplifier and a static power-saving input buffer for low-power architecture. *IEEE J. Solid-State Circuits* **1993**, *28*, 523–527. [[CrossRef](#)]
34. 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](#)]
35. Rabaey, J.M.; Chandrakasan, A.; Nikolic, B. *Digital Integrated Circuits*, 3rd ed.; Prentice Hall Press: Hoboken, NJ, USA, 2008.
36. Fritsch, A.; Kugel, M.; Sautter, R.; Wendel, D.; Pille, J.; Torreiter, O.; Kalyanasundaram, S.; Dobson, D.A. A 4 GHz, low latency TCAM in 14 nm SOI FinFET technology using a high performance current sense amplifier for AC current surge reduction. In Proceedings of the ESSCIRC Conference 2015—41st European Solid-State Circuits Conference (ESSCIRC), Graz, Austria, 14–18 September 2015; pp. 343–346.

-
37. Seevinck, E.; Van Beers, P.; Ontrop, H. Current-mode techniques for high-speed VLSI circuits with application to current sense amplifier for CMOS SRAM's. *IEEE J. Solid-State Circuits* **1991**, *26*, 525–536. [[CrossRef](#)]
 38. NMohan, N.; Fung, W.; Wright, D.; Sachdev, M. Design techniques and testmethodology for low-power tcams. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2006**, *14*, 573–586.