

# CAT: Cellular Automata on Tensor cores

Cristóbal A. Navarro, Felipe A. Quezada, Enzo Meneses, Héctor Ferrada, Nancy Hitschfeld

**Abstract**—Cellular automata (CA) are simulation models that can produce complex emergent behaviors from simple local rules. Although state-of-the-art GPU solutions are already fast due to their data-parallel nature, their performance can rapidly degrade in CA with a large neighborhood radius. With the inclusion of tensor cores across the entire GPU ecosystem, interest has grown in finding ways to leverage these fast units outside the field of artificial intelligence, which was their original purpose. In this work, we present CAT, a GPU tensor core approach that can accelerate CA in which the cell transition function acts on a weighted summation of its neighborhood. CAT is evaluated theoretically, using an extended PRAM cost model, as well as empirically using the Larger Than Life (LTL) family of CA as case studies. The results confirm that the cost model is accurate, showing that CAT exhibits constant time throughout the entire radius range  $1 \leq r \leq 16$ , and its theoretical speedups agree with the empirical results. At low radius  $r = 1, 2$ , CAT is competitive and is only surpassed by the fastest state-of-the-art GPU solution. Starting from  $r = 3$ , CAT progressively outperforms all other approaches, reaching speedups of up to  $101\times$  over a GPU baseline and up to  $\sim 14\times$  over the fastest state-of-the-art GPU approach. In terms of energy efficiency, CAT is competitive in the range  $1 \leq r \leq 4$  and from  $r \geq 5$  it is the most energy efficient approach. As for performance scaling across GPU architectures, CAT shows a promising trend that if continues for future generations, it would increase its performance at a higher rate than classical GPU solutions. A CPU version of CAT was also explored, using the recently introduced AMX instructions. Although its performance is still below GPU tensor cores, it is a promising approach as it can still outperform some GPU approaches at large radius. The results obtained in this work put CAT as an approach with great potential for scientists that need to study emerging phenomena on CA with large neighborhood radius, both in GPU and CPU.

**Index Terms**—Cellular Automata, Tensor Cores, Game of Life, Larger than life, GPU Computing, Energy Efficiency.

## I. INTRODUCTION

Cellular automata (CA) are discrete dynamical systems capable of producing complex emergent phenomena [1]. Originally conceived as abstract mathematical systems, CA have found utility in several fields of science and technology; from artificial life systems [2], [3], [4], fluid simulation [5] to highly complex ecological and urban systems [6], [7], among many others. In terms of computation, a key property of CA is that they are highly data-parallel, making them ideal candidates for parallel computing.

With the advent of Graphics Processing Units (GPUs), current implementations of CA have harnessed their parallel processing capabilities [8], enabling the study of larger-scale and more intricate models. State-of-the-art implementations employ several techniques that range from using the

programmable cache memory (known as *shared memory* in CUDA) [9], multi-GPU [10], to more specific ones such as using multiple cells per thread [11], packet coding [12] and multiple step simulation [13] in programmable cache. These state-of-the-art techniques build on top of the traditional data-parallel scheme for cellular automata where each GPU thread is in charge of one (or more) cells and its entire neighborhood of radius  $r$ , typically Moore or Von Neumann. An example illustration is shown in Figure 1.




Fig. 1. Traditional data-parallel approach for simulating Cellular Automata (CA) using a global halo of ghost cells which is common for avoiding complex logic on the boundary threads. In this example each thread  $t_{i,j}$  is in charge of one cell and must explore its Moore neighborhood of radius  $r = 1$ . In general, with this approach one simulation step costs at least  $\Omega(r^2)$  time.

One of the main limitations of the traditional data-parallel scheme is that performance decreases as the CA neighborhood radius  $r$  increases. This is because for each thread, a potential neighborhood of  $(1 + 2r) \times (1 + 2r)$  cells must be explored, producing a memory bound scenario that costs at least  $\Omega(r^2)$  memory accesses. This cost per thread limits the possibilities of researching and applying CA models such as Larger than Life [14] which exhibits emergent complex phenomena at large neighborhood radius. Therefore, today a key challenge is to find new ways to achieve fast and efficient simulation of CAs with large neighborhood radius.

In the last decade, GPUs have shifted from pure general purpose parallel processors (FP32 / FP64 / INT32 units), to hybrid ones that also include specific purpose units which are significantly faster for the task they were designed for. One type of specialized units are the tensor cores [15], which are part of the GPU chip and offer a hardware accelerated Matrix Multiply Accumulate (MMA) function that runs in one GPU cycle. These tensor core units were designed to keep up with the demands of applications of artificial intelligence (AI) that require training very large models, such as large language

Cristóbal A. Navarro, Felipe A. Quezada, Enzo Meneses, and Héctor Ferrada are with the Department of Informatics of Austral University of Chile and the Temporal research group (<http://temporal.uach.cl>).

Nancy Hitschfeld is with the Computer Science Department of University of Chile.

models (LLMs) [16] or computer vision models (CVM) [17], among many others. Considering the high performance of tensor cores and the fact that modern GPU chips can contain hundreds of them, it is relevant to explore whether or not tensor cores can be used to accelerate CA simulations beyond current GPU solutions, and mitigate the performance degradation identified when increasing the neighborhood radius.

In this work, we present CAT (Cellular Automata on Tensor cores); a GPU method that uses tensor cores to simulate Cellular Automata (CA) at different neighborhood radii. The main feature of CAT is that its cost per simulation step remains constant, *i.e.*,  $\mathcal{O}(1)$ , as long as the radius is within the dimension of the square tensor core matrix. For current GPUs, the matrix dimensions are  $16 \times 16$ , thus the supported constant cost range is  $1 \leq r \leq 16$ . With this design, CAT avoids performance degradation that is found in other state-of-the-art approaches when the radius is increased. In terms of requirements, CAT can accelerate any CA where the cell transition function acts on a weighted summation of the cell's neighborhood. A CPU version of CAT was also explored using the recent AMX instructions, showing promising results at a large radius, even outperforming some GPU approaches.

The rest of the manuscript is organized as follows. Section II covers related work, including the state-of-the-art approaches for which CAT will be compared. Section III contains the formulation and analysis of theoretical performance of CAT. Section V presents experimental results using the Larger than Life (LTL) family of CA at different neighborhood radius. Section VI discusses the main results and concludes the work.

## II. RELATED WORK

Up to date, several works have provided performance comparisons between state-of-the-art techniques to simulate CA on multicore CPU and GPU [18], [8]. In general, it is well known that GPU implementations can run faster than multicore CPU ones (assuming high-end hardware on both cases); however, it is worth noting that certain CA features may affect the performance of one or the other significantly. For example, more arithmetic operations in the transition function have been found to favor more GPU parallelism [8], while complex neighborhood accesses or avoiding the quiescent states of a CA is a big challenge for GPUs and less so for CPUs [19], [18].

Regarding research on GPU approaches, Fujita *et al.* presented a GPU implementation to accelerate the simulation of game of life (GOL) [20] using a multistep scheme. This scheme, also known as temporal blocking [21], consists of loading tiles of the CA into CUDA's programmable shared memory (one region per CUDA block), and simulating several time steps per kernel call in order to reuse cache memory. During each kernel call, a halo of outdated cells progressively grows from the perimeter of the CA region to the interior, for  $t$ -steps. The cells that are further inside (not touched by this halo) become valid simulated cells after all the steps. The authors report large speedups of up to two orders of magnitude over a sequential CPU implementation. This approach is best suited for low neighborhood radius, as the *outdated cells*

boundary propagates at a rate of  $r$  cells per time step, limiting the effectiveness of the technique at larger radius. The authors did not link any source code in their manuscript.

Millan *et al.* [11] implemented and compared the performance of classic and state-of-the-art GPU approaches for simulating CA, using the game of life (GOL) as case study as well as a compute-bound variant with additional synthetic computation per time step (in order to manifest maximum parallelism in the GPU). Their results showed that with modern GPU architectures (2017+) the classic global memory implementation is often the fastest approach, followed by a *multi-cell* approach [22] which consists of simulating two (or more) cells per thread in order to re-utilize part of the thread's neighborhood exploration effort on the adjacent cell. The authors also consider radius  $r > 1$ , and show performance results in the range  $1 \leq r \leq 5$ . Their results reveal that indeed the running times increase with the radius because the work per cell increases at least as  $\Omega(r^2)$ . These results support the fact that running efficient CA simulations with larger neighborhood radius is a known challenge that still requires more research [23]. In the same work the authors also compared their results with a shared memory approach [24], [25] where for each block, the first four rows of threads are in charge of filling the halo of the shared memory tile. They found that a *shared-memory* approach on stencil-like patterns [26] such as CA does not necessarily improve the performance anymore as it did in the past. This behavior has also been reported in another work [27] and agrees with our experimental results when comparing global vs shared memory baselines at low radius. This phenomenon can be attributed in part to the L1/L2 caches now being much more effective than before as well as the increased size of the L2 which automatically caches the global memory accesses of threads. Source codes of both the multi-cell and shared memory implementations were made available by Millan *et al.* [11]<sup>1</sup> and were included in our experimental comparison, named MCELL and SHARED respectively. These implementations were extended by our team to support  $r \in [1..16]$  and performance optimizations were made for the multi-cell (MCELL) approach.

Cagigas *et al.* [12] proposed an efficient approach for simulating CA using regular GPU Computing, *i.e.*, no use of tensor cores. The core idea is the application of *packet coding*, a technique where each GPU thread reads a 64-bit word from GPU global memory and simulates 8-bit cells codified inside with bitwise operations. Two benefits arise from this approach; i) just one global coalesced memory access is performed for multiple cells, and ii) the computational effort for accessing the neighborhood is shared among the codified cells. The authors report that doing packet coding is between  $3 \times$  to  $4 \times$  faster than the baseline GPU approach, and faster than other state of the art techniques such as *lookup-table* (*i.e.*, to precompute possible outcomes of the transition function, and access the table with the cell's current state and neighbor information) and *temporal blocking* [21] implemented with the AN5D framework [28]. The work was focused on radius  $r = 1$ ,

<sup>1</sup>The classic global memory approach is also included in the experiments as part of our own baseline implementations.

however the approach can be extended to larger radius. The authors made their source code available but only working for radius  $r = 1$ , therefore our team extended the implementation, named as PACK, to support radiuses in the range  $r \in [1..16]$ .

Regarding the use of tensor math, Zhuang *et al.* proposed a Python deep learning framework based on high-level tensor computation on GPU to accelerate land simulation CA [29]. The authors report that by using their approach, they achieve up to  $\sim 50\times$  of speedup over a sequential counterpart running on a CPU. No details are given on the use of tensor-cores, neither tests on other well known cellular automata such as game of life (GOL), as the scope of the work is more oriented at formulating the problem in a high-level tensor framework for a specific application.

Regarding the use of GPU tensor cores, Liu *et al.* [30] used them to accelerate the finite difference method (FDM) for partial differential equations (PDEs), which is a stencil-like pattern. The authors report average speedups of up to  $\sim 4.55\times$  over a highly optimized GPU baseline. This work is relevant to mention because it puts a precedent on the benefits of using tensor cores in stencil-like patterns, and it can even be used to simulate cellular automata as well, although best suited for CA with small neighborhood radius. This is because it was formulated using an *in-halo* scheme. Although this scheme has the benefit of using fewer matrix operations than the method we propose (two products instead of six), it reduces the effective tensor core fragment size  $p \times q$  cells, to  $(p - 2r) \times (q - 2r)$  cells. This creates the need to overlap fragments according to  $r$  and limits the range of  $r$  to  $r < \frac{p}{2}, r < \frac{q}{2}$ . Considering that current tensor core fragments can have sizes in the order of  $p \times q = 16 \times 16$  elements, the *in-halo* approach at  $r = 1$  would produce an effective simulation area of  $14 \times 14$  which is not a big issue. But as the radius increases, the *in halo* would grow to the point that only a small area in the center is effective. In contrast, CAT (our proposed method) chooses the opposite design; an *out-halo* scheme which has favorable implications for large radius simulations with tensor cores. The design of CAT, including its *out-halo* scheme, is detailed in Section III, with its possible extensions explained in Section VI.

Lastly, Chen *et al.* [31] recently proposed an efficient way of doing high-precision (FP64) stencil computations on GPU using tensor cores as well. Their matrix layout is different however, as they convert tiles of the input data to matrix rows and convolution tiles to matrix columns, among other technical improvements such as a lookup table to reduce integer operations and *dirty bits* to alleviate bank conflicts in shared memory.

### III. FORMULATION OF CAT

One of the most time consuming tasks in simulating many CA is accessing the neighboring cells, which must be done for each cell of the domain and at each time step of the simulation. Such CA are known to be memory-bound because the dominant cost is in the memory accesses (*i.e.*, neighborhood access) rather than in the arithmetic operations of the transition function. A very well known example of memory-bound CA is John Conway's Game of Life [20], [2] and

its generalization to any radius known as Larger than Life [32], [14], [4] which is even more memory-bound due to the larger neighborhood radius. Apart from being memory-bound, these CA also characterize for having a transition function that acts on the weighted summation of each cell's neighborhood. Simulating memory-bound CA in GPU has its challenges because it saturates the memory bandwidth much earlier than the computational capacity of the chip. Traditional GPU solutions often mitigate this problem by doing an efficient use of the memory hierarchy, including the programmable L1 cache known as *shared memory* in CUDA's terminology. The proposed method, *Cellular Automata on Tensor cores* (CAT), handles these types of CA from a different perspective; by adapting the exploration of neighborhood cells as a series of MMA operations executed by the tensor cores, which take one GPU cycle per MMA.

The formulation of CAT relies on a fundamental linear algebra fact which is that the weighted neighborhood summation of a CA cells (except for the boundary ones) can be computed simultaneously with two matrix products between the entire CA domain and a constant band matrix of ones. The result of these two matrix operations returns a matrix where each element is the weighted sum of all of its neighborhood data including itself twice<sup>2</sup>. For neighborhood radius  $r = 1$ , the band matrix is tridiagonal, and for general radius the diagonal band has a width of  $2r + 1$ . A frequent approach for handling the neighborhood of boundary cells more efficiently is to include a global halo of ghost cells of width  $r$ . Although this global halo has an extra memory cost of  $\approx 4nr = \mathcal{O}(n)$ , it is not significant compared to the problem size  $\mathcal{O}(n^2)$ .

Let  $\Lambda_{n' \times n'}$  be the matrix representation of an entire cellular automata of size  $n \times n$ , using a global halo of ghost cells where  $n' = n + 2r$ , and  $\Pi_{n' \times n'}$  the entire band matrix with halo as well, then the reduction of all Moore neighborhoods can be computed with two MMAs:

$$H_{n' \times n'} = \Lambda_{n' \times n'} \times \Pi_{n' \times n'} + 0_{n' \times n'} \quad (1)$$

$$R_{n' \times n'}^{\text{Moore}} = \Pi_{n' \times n'} \times H_{n' \times n'} + 0_{n' \times n'} \quad (2)$$

The first step reduces horizontally into  $H$ , while the second step does a vertical reduction using  $H$  as input, to end up with the Moore neighborhood reduction matrix  $R$ . The simplified Von Neumann neighborhood can also be computed, it only requires to redefine Eq. (2) as:

$$R_{n' \times n'}^{\text{Von Neumann}} = \Pi_{n' \times n'} \times \Lambda_{n' \times n'} + H_{n' \times n'} \quad (3)$$

Figure 2 illustrates an example of how the two MMAs (the  $0_{n' \times n'}$  matrices were omitted) apply for a Game of Life (GoL) CA of size  $n \times n = 16 \times 16$  with Moore neighborhood of  $r = 1$ . The purple bands surrounding the  $16 \times 16$  domain is the global halo of ghost cells.

In any of the two neighborhood types (Moore or Von Neumann), there will be duplicate state values added on each non-zero cell. If the CA rule needs to exclude the center cell in the neighborhood counting, then one can compute  $R - 2\Lambda$  in parallel as an extra instruction in the same GPU kernel, using a standard per-thread logic.

<sup>2</sup>The center cell value can later be subtracted by each thread if needed.



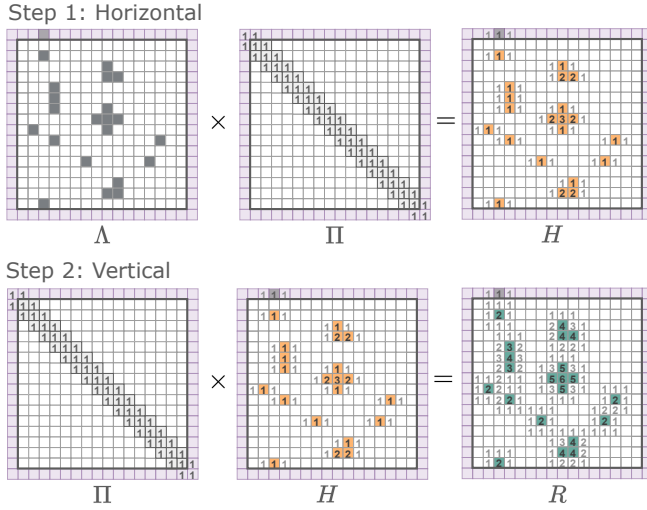


Fig. 2. Concept of how a pair of matrix products between an entire CA ( $\Lambda$ ) and a band matrix ( $\Pi$ ) can count the living neighbors of all cells (no tensor core logic introduced yet). Here, the CA includes a global halo of ghost cells, giving a total size of  $(n + 2r) \times (n + 2r) = 18 \times 18$  as the neighborhood radius is  $r = 1$ . The final cells of  $R$  contain their number of living neighbors plus its own state added twice (cells with no number have value zero).

In order to compute Eqs. (1) and (2) (or Eq. (3) if needed) efficiently with tensor cores, one has to consider that the products defined between the CA and the band matrices (as shown in Figure 2), in practice have to be programmed as several tensor core MMAs that occur at a smaller scale and in parallel between fragments of  $\Lambda$  and  $\Pi$ , similar to the blocked matrix multiply scheme. Furthermore, given that the band matrix  $\Pi$  is non-zero only on its diagonal band, then the MMAs that actually matter are the ones where the band is present, *i.e.*, all other MMAs can be skipped as the product would compute to zero.

Applying the two MMAs idea of Figure 2 directly onto each fragment would require the *in-halo* scheme, which is not efficient for a large neighborhood radius, as it would restrict the effective fragment area from  $p \times q$  to  $(p - 2r) \times (q - 2r)$  cells, and all fragments would have to be overlapped in  $r$  cells in order to cover the entire domain. Moreover, an *in-halo* scheme limits the radius range up to half the size of the tensor core fragment. To prevent this, CAT handles the problem differently; fragments do not restrict the effective area to their inside when  $r$  increases, but instead expand to the outside allowing to simulate the entire fragment without imposing limits on the radius range. This is achieved by including two more MMAs for the neighbor fragments; one for each direction (horizontal and vertical). This design requires the width of the global halo of ghost cells to be a multiple of the fragment size. Considering that CAT operates with one center fragment and two adjacent ones, and the fact that the band matrix is constant,  $\Pi$  can be just represented with three fragments;  $\pi_1, \pi_2, \pi_3$  as shown in Figure 3 (fragments are of size  $4 \times 4$  just for visual simplicity).

Representing  $\Pi$  with just three fragments not only simplifies tensor core programming inside the GPU kernels, but it also reduces the memory usage for  $\Pi$ , from  $\mathcal{O}(n^2)$  down to  $\mathcal{O}(1)$ .




Fig. 3. On the left, an explicit representation of the band matrix  $\Pi$ , which uses  $\mathcal{O}(n^2)$  memory. On the right, the CAT representation of  $\Pi$ , which uses just three fragments  $\pi_1, \pi_2, \pi_3$  to represent the entire matrix. Fragments are of size  $4 \times 4$  just for visual simplicity.

Considering that the cellular automata matrix  $\Lambda$  is composed of fragments  $F_{i,j}^\Lambda$ , and  $\Pi$  is represented by three constant fragments  $\pi_1, \pi_2, \pi_3$ , then the first step of CAT, *i.e.*, the computation of fragments  $F_{i,j}^H$  in the horizontal reduction matrix  $H$ , is

$$F_{i,j}^H = F_{i,j-1}^\Lambda \times \pi_1 + F_{i,j}^\Lambda \times \pi_2 + F_{i,j+1}^\Lambda \times \pi_3 \quad (4)$$

with  $i$  including the top and lower halo fragments and  $j$  not including them (*i.e.*, just the interior columns). All  $F_{i,j}^H$  can be computed in parallel, each one handled by a different warp of threads. Once all fragments of  $H$  have been computed, the second step (vertical reduction) consists of computing the fragments of matrix  $R$ , which are defined in terms of the fragments of  $H$  as

$$F_{i,j}^R = \pi_3 \times F_{i-1,j}^H + \pi_2 \times F_{i,j}^H + \pi_1 \times F_{i+1,j}^H \quad (5)$$

this time with  $i, j$  only including the interior fragments, not the halo ones. Again, all  $F_{i,j}^R$  are computed in parallel, one per warp, using the tensor cores of the GPU. It is worth mentioning that the combination of the horizontal and vertical steps and the re-use of  $H$  as input for the second step allow CAT to capture the diagonal neighbors of a cell.

The entire overview of CAT, including the horizontal/vertical steps and the regions where the fragments need to be computed (dashed regions), is illustrated in Figure 4 for the case of Moore neighborhood<sup>3</sup> and fragments of  $4 \times 4$  for visual simplicity. From the Figure, matrix  $\Lambda$  is loaded into fragments  $F_{i,j}^\Lambda$  using a per-warp logic. Then, each fragment  $F_{i,j}^H$  inside the dashed region needs to be computed and is the result of computing Eq. (4). Once all threads synchronize with the first step finished, they proceed to the second step which is similar, but uses  $H$  as input instead of  $\Lambda$  and computes Eq. (5) on the dashed region of  $R$ . Once  $R$  is computed, threads can continue in the same kernel with the transition function to each cell, using a per-thread logic.

<sup>3</sup>The cost does not change when switching to the simplified Von Neumann neighborhood, as both neighborhoods employ the same amount of MMAs.

## CAT Overview

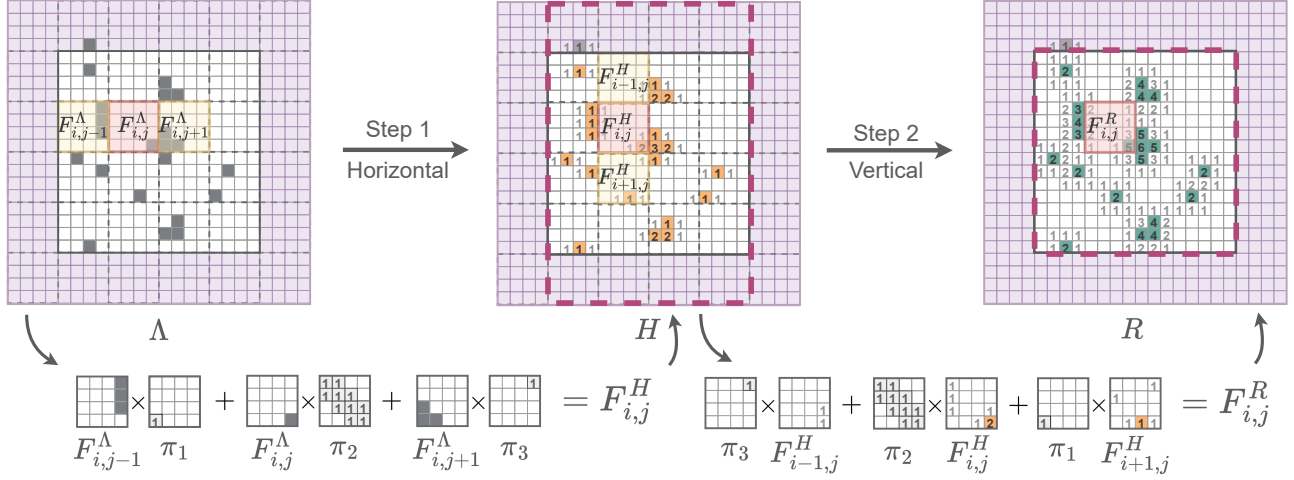


Fig. 4. Overview of CAT illustrated with a Game of Life of  $n \times n = 16 \times 16$  cells, neighborhood  $r = 1$  and periodic boundary conditions using a global halo of ghost fragments (purple). In the first step all fragments  $F_{i,j}^H$  inside the dashed region of  $H$  contain the horizontal reduction computed with three sequential MMAs between fragments  $F_{i,j-1}^\Lambda$ ,  $F_{i,j}^\Lambda$ ,  $F_{i,j+1}^\Lambda$  and  $\pi_1, \pi_2, \pi_3$ . In the second step all  $F_{i,j}^R$  inside the dashed region of  $R$  contain the full reduction computed with three more MMAs between the fragments  $\pi_3, \pi_2, \pi_1$  and  $F_{i-1,j}^H, F_{i,j}^H, F_{i+1,j}^H$ . This gives a total cost of six MMAs per fragment at any radius that fits in the fragment. For this example the fragments were shown as  $4 \times 4$  for visual clarity, but in practice the ones employed in CAT are of size  $16 \times 16$ .


Matrices  $\Lambda$  and  $R$  need to be in GPU memory, as they play the role of *in* and *out*, respectively, as in any standard GPU-based CA implementation. Matrix  $H$  does not need to be in memory, as they are actually fragments that emerge at GPU cache level.

#### A. CAT with large neighborhood radius

The main benefit of CAT is that by design it already supports simulation on a large neighborhood radius up to the fragment size. This is because instead of using an *in-halo* scheme, *i.e.*, to restrict the effective simulation area to the inside of the fragment as  $r$  increases, it uses an *out-halo* scheme where the effective area is always the entire fragment and increasing  $r$  only has an impact on the width of the band matrix which in practice means defining  $\pi_1, \pi_2, \pi_3$  according to the  $r$  value of the CA model. In other words, increasing  $r$  in CAT produces the same six MMAs per fragment, but with more effective computation than with a lower radius. At first, six MMAs per fragment may be seen as too much work for  $r = 1$ , but it is greatly compensated by the fact that the cost has already been paid for higher radius, thus in theory the execution time of CAT is unaffected when increasing the radius (in the experimental results this is confirmed) as long as  $r$  is within<sup>4</sup> the fragment size. In general, if the fragments are of size  $p \times p$ , then the maximum supported radius is  $r = f$ . In the case of CAT, it uses CUDA fragments of  $16 \times 16$  thus the maximum radius is  $r = 16$ . Figure 5 shows how  $\pi_1, \pi_2, \pi_3$  become for radius  $r = 8$  and  $r = 16$ .

<sup>4</sup>It is possible to extend CAT to support neighborhood radiuses beyond the fragment size. This is discussed in Section VI.

$r = 8 \rightarrow \text{width} = 17$



$r = 16 \rightarrow \text{width} = 33$




Fig. 5. The band fragments at two different radius values  $r = 8, r = 16$ . Here, the size of the fragment size is  $16 \times 16$ , the actual size used by CAT. With this fragment size, the maximum supported radius is  $r = 16$ .

#### B. CUDA Specific Optimizations

CAT includes three technical optimizations: 1) Fragment-level contiguous memory layout, 2) use of *shared memory* and 3) optimal tile per CUDA block.

1) *Fragment-level contiguous memory layout:* CAT's GPU kernel begins with each warp loading the data from GPU global memory into fragments of  $16 \times 16$  cells which reside at the register level. Usually, the CA data in global memory is in row-major layout, which is efficient for traditional GPU

approaches in terms of memory accesses. However, with tensor cores, this layout is less efficient because fragments are continuous 2D regions of memory and if the CA is in row-major then there will be large memory strides of size  $n$  between the rows of the fragment. This would produce a slowdown in memory bandwidth, which can diminish the benefit of using CAT significantly. To overcome this problem, CAT's memory layout linearizes fragments in memory, one after another. In this new layout, a group of  $16 \times 16 = 256$  consecutive elements corresponds to a 2D fragment. This change makes the loading of data into fragments fully coalesced. Figure 6 illustrates CAT's memory layout.




Fig. 6. The fragment-level contiguous memory layout used by CAT. At the inner level, each fragment has its own row-major layout. At the outer level, the entire CA has a row-major layout of fragments.

Tensor core memory layouts have been studied in the context of deep learning, such as in the pooling operation, where the NHWC ( $N$  images,  $H$  height,  $W$  width,  $C$  channels) layout is recommended for tensor cores [33]. In these terms, CAT uses a NHWC layout but local to each fragment, similar to recent ideas in deep learning [34] and with  $C = 1$ .

2) *Use of shared-memory for intermediate results:* CUDA's shared memory resource was first considered for the entire process of CAT, that is, at the beginning of the kernel each CUDA block of threads would load all their corresponding cells from  $\Lambda$  into a 2D *shared memory* buffer, synchronize, and then perform all the remaining steps at the shared memory level until the result is written back into global memory. However, preliminary profiling of CAT showed that moving data from global to shared, and from shared to fragment registers, turned out to be slower than just loading from global to fragment. Still, CAT does use shared memory for the tile (per block) of intermediate fragments  $F_{i,j}^H$ , which is a fast movement of data from register level to L1 Cache *shared memory* that occurs physically in the same streaming multiprocessor (SM). Additionally, in the case of the constant fragments  $\pi_1, \pi_2, \pi_3$ , given that they are the same for all warps, they are generated once in shared memory.

3) *Optimal Tile per CUDA block:* CAT works by mapping a CUDA block of threads to a tile  $Q$  of CA cells. Doing a full pass on all the tiles produces a simulation step of the CA. Here, a tile  $Q$  is defined as a rectangular region of  $w \times h$  fragments of cells. A relevant aspect of CAT is the detachment of the one-to-one mapping between the CUDA block of warps and a tile  $Q$  of data. That is, CAT allows a tile  $Q$  to have more fragments than the number of warps available in the mapped

CUDA block, and a different geometry as well. This relaxation opens the possibility to explore what values of width ( $w$ ) and height ( $h$ ) of  $Q_{w \times h}$  produces the highest performance of CAT.

Figure 7 presents a heat map where different tile shapes were explored using a large CA of size  $n \times n = 60416 \times 60416$ . From the Figure, one can note that the most efficient tile is of




Fig. 7. Heat map of the optimal shape for tile  $T^H$ . Long shaped tiles produce faster performance than square ones, specially narrow ones in width. The optimal tile shape for  $T^H$  is near  $w \times h = 1 \times 14$ .

vertical shape, near  $w \times h = 1 \times 14$ . It was unexpected that a column-shaped tile would be more efficient than one with a regular shape. The experimental results shown in Section V use the optimal shape for CAT.

### C. Pseudocode of CAT's Kernel

CAT uses the well known *ping-pong* simulation scheme where two copies of the CA are used. One of them, the  $CA_{in}$ , holds the current state of cells for reading, and  $CA_{out}$  is used to write the future state of cells. For each simulation step these buffers switch places. Algorithm 1 presents the pseudocode of CAT's GPU kernel, with its three optimizations recently described. From the pseudocode, one can note that there are three major synchronization barriers. The algorithm starts by creating/initializing shared memory data such as the tile  $Q^H$  and shared arrays for  $\pi_1, \pi_2, \pi_3$  before reaching the first barrier. After this barrier, threads retrieve some warp information such as warp ID and the total number per CUDA block. Then, each warp proceeds to compute the horizontal reduction for all of its corresponding fragments using a `for` loop that has a stride of the number of warps in a CUDA block. At each iteration, the resulting horizontal reduction is stored in tile  $Q^H$ . The second sync barrier ensures that all warps have finished the horizontal reduction and stored its results in  $Q^H$  before passing to the vertical reduction which is similar in logic. In general, for each reduction phase, three MMA operations are employed with the help of an auxiliary fragment  $F$  to accumulate the products and store the resulting fragment to the corresponding array;  $Q^H$  for horizontal and  $R$  for vertical. The third barrier ensures all reductions have

**Algorithm 1** Pseudocode of CAT's GPU Kernel

---

**Require:**  $n$ :size,  $\Lambda$ : In,  $R$ : Out,  $w, h$ :tile size,  $r$ : radius

$Q^H \leftarrow \text{DeclareSharedMemoryTile}(w+2, h+2)$

$g_i, g_j \leftarrow \text{TileGlobalOffset}(n, w, h, \text{BlockID})$

$S_{\pi_1}, S_{\pi_2}, S_{\pi_3} \leftarrow \text{GenSharedBandMats}(r)$

---

$\text{SyncThreads}()$  ▷ **Step 1: Compute H**

$W_x \leftarrow \text{getWarpID}()$

$W_b \leftarrow \text{getNumWarpsPerBlock}()$

$\pi_1, \pi_2, \pi_3 \leftarrow \text{getBandFragments}(S_{\pi_1}, S_{\pi_2}, S_{\pi_3})$

**for**  $k \leftarrow W_x$ ;  $k < w \times (h+2)$ ;  $k \leftarrow k + W_b$  **do**

$F \leftarrow \text{ZeroFragment}()$

$(i, j) \leftarrow ([k/w], k \bmod w) + (0, 1)$

$F_{i,j-1}^A, F_{i,j}^A, F_{i,j+1}^A \leftarrow \text{loadFragments}(\Lambda, g_i + i, g_j + j)$

$F \leftarrow \text{MMA}_{TC}(F_{i,j-1}^A, \pi_1, [0])$

$F \leftarrow \text{MMA}_{TC}(F_{i,j}^A, \pi_2, F)$

$F \leftarrow \text{MMA}_{TC}(F_{i,j+1}^A, \pi_3, F)$

$\text{Store}(Q^H, F, i, j)$

**end for**

---

$\text{SyncThreads}()$  ▷ **Step 2: Compute R**

**for**  $k \leftarrow W_x$ ;  $k < w \times h$ ;  $k \leftarrow k + W_b$  **do**

$F \leftarrow \text{ZeroFragment}()$

$(i, j) \leftarrow ([k/w], k \bmod w) + (1, 1)$

$F_{i-1,j}^H, F_{i,j}^H, F_{i+1,j}^H \leftarrow \text{loadFragments}(Q^H, i, j)$

$F \leftarrow \text{MMA}_{TC}(\pi_3, F_{i-1,j}^H, [0])$

$F \leftarrow \text{MMA}_{TC}(\pi_2, F_{i,j}^H, F)$

$F \leftarrow \text{MMA}_{TC}(\pi_1, F_{i+1,j}^H, F)$

$\text{Store}(R, F, g_i + i, g_j + j)$

**end for**

---

$\text{SyncThreads}()$

**for** cells assigned to current thread  $t$  **do** ▷ Per Thread

$c \leftarrow \text{getCell}(\Lambda, t)$

$q \leftarrow \text{getNeighborsReduction}(R, t)$

$c' \leftarrow \text{applyCARule}(c, q)$

$\text{Store}(R, c', t)$

**end for**

---

finished, and allows threads to compute the future state of all their corresponding cells.

**D. Cost Analysis of CAT**

Computing the cost of CAT involves adding the costs of the two main steps (Figure 4) and the application of the transition function to each cell. One way to proceed with the analysis is to obtain the parallel time of a CUDA block processing a CA tile with Algorithm 1, considering a finite number of resources (regular cores and tensor cores) residing on a GPU streaming multiprocessor (SM). With the parallel time per block computed, then one can expand it to the total number of tiles that need to be processed, considering that there is a finite number of SMs in a GPU chip. For this, we employ a PRAM-like model [35] in its CREW variant (Concurrent Read, Exclusive Write), but with finite resources and two extensions: i) two types of memory accesses; global memory accesses assisted by automatic L2 caching, with a cost of  $C$ , and cache memory (L1) with a smaller cost of  $c$ , i.e.,  $C = \alpha c$  with  $\alpha > 1$ ; ii) each CUDA MMA costs  $\tau$  which is the actual number of one-cycle physical tensor core executions for the  $p \times q \times k$  fragment.

In CAT, the task of each CUDA block is to simulate a tile of  $w \times h$  fragments from the cell matrix  $\Lambda$ . In the first step of Algorithm 1, the computation of each fragment  $F_{i,j}^H$  involves reading three fragments from  $\Lambda$  and the band

fragments  $\pi_1, \pi_2, \pi_3$  from shared memory, then perform three MMAs (see Figure 4), and write the resulting fragment into  $Q^H$ . The first three memory accesses are coalesced (because of the fragment-row memory layout optimization), costing  $C$  units each, while the reads on the band fragments cost  $c$  units each. For the three MMAs, these cost  $\tau$  units each and run one after another because one warp is in charge. Lastly, the write operation costs  $c$  as it is on shared memory. With these considerations, the time for computing a fragment  $F_{i,j}^H$  is

$$\mathcal{T}_{F_{i,j}^H} = 3C + 3\tau + 4c. \quad (6)$$

Considering that each CUDA block gets one tile assigned, there are  $w \times (h+2)$  fragments to compute for the tile  $Q^H$ . At a logical level, the number of fragments that get computed in parallel depends on the number of warps per block  $W_b$ , but at the physical level it depends on the number of tensor cores in a SM, which we denote  $Z_{sm}$ . Given that we have  $Z_{sm} \leq W_b$ , then we can just consider  $Z_{sm}$  as it is the dominant restriction. With this consideration, the parallel time for computing a tile  $Q^H$  becomes:

$$\mathcal{T}_{Q^H} = \left\lceil \frac{w \cdot (h+2)}{Z_{sm}} \right\rceil \mathcal{T}_{F_{i,j}^H} \quad (7)$$

$$= \left\lceil \frac{w \cdot (h+2)}{Z_{sm}} \right\rceil (3C + 3\tau + 4c). \quad (8)$$

For the second step, fragments  $F_{i,j}^R$  need to be computed. In this case, the warps read the three fragments from the resulting tile  $Q^H$  of step 1 that lies in shared memory, thus they cost  $c$  units each. The band fragments  $\pi_1, \pi_2, \pi_3$  must be loaded again this time as the first operand in the MMA, each one costing  $c$ . Adding these costs with the three MMAs each of cost  $\tau$  and with the store operation on  $R$  that lies on global memory, we have that the time for computing  $F_{i,j}^R$  is

$$\mathcal{T}_{F_{i,j}^R} = 6c + 3\tau + C. \quad (9)$$

In this second step, the target tile  $Q^R$  has  $w \times h$  fragments to compute for a given CUDA block. With this, the parallel time for computing tile  $Q^R$  is

$$\mathcal{T}_{Q^R} = \left\lceil \frac{w \cdot h}{Z_{sm}} \right\rceil \mathcal{T}_{F_{i,j}^R} \quad (10)$$

$$= \left\lceil \frac{w \cdot h}{Z_{sm}} \right\rceil (6c + 3\tau + C). \quad (11)$$

Then, in the last stage of Algorithm 1, each thread applies the CA's transition function  $f()$ , of  $\delta$  units, to each of its corresponding cells. For each cell in the tile, this involves reading the original cell state from  $\Lambda$  ( $C$  units), the neighborhood reduction from  $R$  ( $C$  units), then applying  $f()$  ( $\delta$  units) and lastly writing the result back again on  $R$  ( $C$  units). This gives a cost of  $(\delta + 3C)$  per cell. For the parallelism at this stage, the number of regular cores (i.e., the number of FP32 units as reference) are considered per SM, denoted  $P_{sm}$ , as it is the dominant restriction for threads in a CUDA block. With this, the parallel time for this stage of the tile is

$$\mathcal{T}_f = (\delta + 3C) \left\lceil \frac{(w \cdot h)(p \cdot q)}{P_{sm}} \right\rceil \quad (12)$$



The last cost to include is the one at the beginning of Algorithm 1, where threads cooperatively initialize the band matrices  $S_{\pi_1}, S_{\pi_2}, S_{\pi_3}$  in parallel with the regular cores  $P_{sm}$  in the SM. This gives a cost per band of

$$\mathcal{T}_{S_{\pi}} = c \left\lceil \frac{p \cdot q}{P_{sm}} \right\rceil \quad (13)$$

Adding the costs of Eqs. (7), (10), (12) and (13), we get the following total parallel time per tile:

$$\mathcal{T}_{Q_{w \times h}} = 3\mathcal{T}_{S_{\pi}} + E \cdot (\mathcal{T}_{Q^H} + \mathcal{T}_{Q^R} + \mathcal{T}_f) \quad (14)$$

where  $E \geq 1$  is a tile efficiency factor that models the behavior found in Figure 7, where some  $w \times h$  tile shapes make CAT faster, while others make it slower. Surface  $E$  was built by fitting a two-variable  $(w, h)$  fourth degree polynomial in the discrete points of the heatmap.

For a CA of  $n \times n$  cells where  $n$  is a multiple<sup>5</sup> of the tile size  $w \times h$  and the fragment size  $p \times q$ , CAT's total parallel time running on a GPU with  $\mathcal{P}$  SMs is

$$\mathcal{T}_{CAT} = \left\lceil \frac{n^2 / (p \cdot q \cdot w \cdot h)}{\mathcal{P}} \right\rceil \mathcal{T}_{Q_{w \times h}} \quad (15)$$

As a comparison, using the same cost model, the parallel running time of a reference (REF) baseline GPU method would do the following per cell; a thread reads its corresponding cell from global memory as well as its neighborhood of radius  $r$ , costing  $(1 + 2r)^2 C$  units. Then, it computes a reduction on its neighborhood (such as addition for example), which costs  $(1 + 2r)^2 - 1$  units and applies the transition function  $f()$  which costs  $\delta$  units per cell. Finally, each thread writes its resulting cell state back to global memory ( $C$  units), leading to a time per cell of

$$\mathcal{T}_{REF}^{cell} = (1 + 2r)^2 C + (1 + 2r)^2 - 1 + \delta + C \quad (16)$$

For an entire CA of  $n \times n$  cells, the parallel time of a reference (REF) baseline GPU algorithm would be

$$\mathcal{T}_{REF} = \left\lceil \frac{n^2}{\mathcal{P} \cdot P_{sm}} \right\rceil \mathcal{T}_{REF}^{cell} \quad (17)$$

$$= \left\lceil \frac{n^2}{\mathcal{P} \cdot P_{sm}} \right\rceil ((1 + 2r)^2 C + (1 + 2r)^2 - 1 + \delta + C). \quad (18)$$

When considering the speedup of  $\mathcal{T}_{CAT}$  with respect to  $\mathcal{T}_{REF}$ , we obtain

$$S_{CAT} = \frac{\mathcal{T}_{REF}}{\mathcal{T}_{CAT}} = \frac{\left\lceil \frac{n^2}{\mathcal{P} \cdot P_{sm}} \right\rceil \mathcal{T}_{REF}^{cell}}{\left\lceil \frac{n^2 / (p \cdot q \cdot w \cdot h)}{\mathcal{P}} \right\rceil \mathcal{T}_{Q_{w \times h}}}. \quad (19)$$

From Eq. (19) one can note that the number of streaming multiprocessors  $\mathcal{P}$  is not as relevant as the internal parallel times at each one of them. To further analyze the speedup,  $S_{CAT}$  must be evaluated with actual values. For this, we consider a large CA such as  $n = 2^{30}$ , a tile size of  $w \times h = 1 \times 14$  and the hardware related parameters aligned to a state of the art GPU such as NVIDIA's full<sup>6</sup> GH100 chip [36]. Table I summarizes

these parameters. The  $C$  value was defined as  $6c$ , and not orders of magnitude higher than  $c$ , because it considers the existence of the L2 cache that automatically assists the global memory on each access.

TABLE I  
CHOSEN PARAMETERS FOR COST MODEL

Parameter	Value	Description
$n \times n$	limit $n \mapsto \infty$	Representing a very large CA size.
$w \times h$	$1 \times 14$	CAT optimal tile size.
$C$	$6c$	Global + automatic L2 as a factor of $c$ .
$c$	1	Manual cache (L1) cost.
$p \times q$	$16 \times 16$	MMA Fragment size.
$\tau$	16	# of internal 1-cycle calls per MMA.
$P_{sm}$	128	Regular cores per SM.
$Z_{sm}$	4	Tensor cores per SM.
$\delta$	20	Cost of CA transition function $f()$ .

Using these parameters, Table II presents theoretical speedups of CAT with respect to the reference GPU approach REF, at radiuses  $r = 1, 4, 8, 16$  under different scenarios including hypothetical ones.

TABLE II  
CAT'S THEORETICAL SPEEDUPS FOR A CA IN THE LIMIT  $n \mapsto \infty$ .

Scenario	Parameter change	Theoretical Speedup $S_{CAT}$			
		$r = 1$	$r = 4$	$r = 8$	$r = 16$
GH100 Chip	n/a	1.20×	8.07×	27.9×	104×
More TC Units	$Z_{sm} : 4 \rightarrow 16$	1.59×	10.6×	37.1×	138×
Faster TC Units	$\tau : 16 \rightarrow 1$	1.55×	10.4×	36.1×	134×
More FP Units	$P_{sm} : 128 \rightarrow 512$	0.60×	4.06×	14.1×	52.5×
Regular Tiles	$w \times h : 1 \times 14 \rightarrow 16 \times 16$	0.17×	1.15×	3.99×	14.8×
High Cost $f()$	$\delta : 20 \rightarrow 1000$	0.79×	1.17×	2.25×	6.44×

In the first scenario, which simulates the full GH100 chip specification, the cost model shows that CAT can reach significant speedups as the CA neighborhood  $r$  increases, being up to 104× faster than REF. This GH100 scenario serves as a reference when measuring experimental speedups with the H100 GPU, which is a slightly cut-down version of the full GH100 chip. The next scenario, which is hypothetical, increases the number of TC units in a SM by a factor of four, showing a significant speedup boost of roughly  $\sim 32\%$ , assuming the tile size has enough fragments to keep all of these extra units active. In the next hypothetical scenario, making the TC units faster shows that it also has a strong effect, giving a boost of  $\sim 28\%$  to its speedup. The third hypothetical scenario explores the case of more regular cores per SM and shows that CAT's speedup is reduced and slower than REF for  $r = 1$ . The reason for this is because the REF implementation benefits more than CAT with this change. The fifth scenario explores the potential penalty if a non-optimal tile size is chosen for CAT, such as  $w \times h = 16 \times 16$ . In this case, the theoretical speedup is significantly reduced to the point of being the slowest scenario for CAT at  $r = 1$ , but it manages to provide favorable speedup for the rest of the radiuses, although highly penalized. The last scenario explores a more expensive transition function, showing that CAT's speedup is also penalized, becoming a less favorable scenario for CAT at high radius with a maximum speedup of 6.44×. The reason why the amount of work in  $f()$  affects CAT so much is that this work is done by the regular GPU cores, even in CAT, thus the tensor core work becomes a smaller fraction of the total

<sup>5</sup>This assumption is just for analysis, CAT supports any value of  $n$ .

<sup>6</sup>The H100 GPU is a slightly cut down version of the full GH100 chip.



work per cell. The scenarios covered in Table II provide useful insights on what hardware aspects and CA settings affect CAT the most. The next Section presents an experimental evaluation of CAT, comparing its performance against a reference GPU baseline (BASE) as well as with other state-of-the-art GPU techniques.

In terms of usability, this cost model only requires defining  $n$  and  $\delta$  which come from the CA in question. The  $C$  and  $c$  parameters represent the latest behavior of GPUs in the last years, therefore can be reused or redefined at will. For the  $p \times q$  parameter, this comes from the chosen fragment type and size. In our case these we chose  $16 \times 16$  FP16 fragments. The other parameters such as  $\tau$ ,  $P_{sm}$  and  $Z_{sm}$  are determined by the GPU architecture, and the optimal tile size  $w \times h$  is experimentally found, being a one-time task per GPU architecture.

#### E. Tensor Core Utilization and Impact

A lower bound on the effective FP16 tensor core TFLOPS (floating point operations per second) was measured using the optimal  $w \times h = 1 \times 14$  tile shape. The lower bound is computed as the quotient between the total number of neighborhood FP summations in one simulation step, *i.e.*,  $(n \times n) \cdot [(1 + 2r) \times (1 + 2r)]$ , divided by the average kernel time of one simulation step. For this measurement, the problem size was set at  $n = 60416$  and the GPU used was the H100 GPU, which has a peak FP16 tensor core performance of 989.5 TFLOPS for dense matrices (the sparsity feature is not used). Table III shows that at low radius the lower bound of TFLOPS has a low utilization of the tensor core (0.27% at  $r = 1$ ) but rapidly increases with the radius of the neighborhood, reaching an effective utilization of at least 33.47%. Having a tensor core utilization that is distant to the peak hardware value is not a surprise considering that CA simulation is by definition a memory bound problem.

TABLE III  
CAT'S EFFECTIVE FP16 TENSOR CORE TFLOPS FOR CA SIMULATION.

$r$	CAT (FP16)	H100 Tensor Core Utilization (Peak 989.5 TFLOPS)
1	2.73 TFLOPS	0.27%
4	24.63 TFLOPS	2.48%
8	87.90 TFLOPS	8.88%
16	331.24 TFLOPS	33.47%

We also measured just the improvement that tensor cores provide to CAT compared to not using them. For this, we compared the CAT approach with another variant that computes the MMAs with regular CUDA threads. Table IV compares the average time per simulation step of both implementations using an H100 GPU simulating a CA at different sizes (CAT running time is not affected by radius). The comparison clearly shows that the tensor cores contribute significantly to the performance. At  $n \geq 10000$  the tensor cores already provide a significant improvement in performance compared to not using them, reaching approximately an order of magnitude faster performance than the regular CUDA core version.

TABLE IV  
IMPACT OF TENSOR CORE MMAS ON CAT'S PERFORMANCE.

$n$	CAT (Tensor Cores)	CAT (CUDA cores)	Improvement by Tensor Cores
1024	0.02 ms	0.06 ms	$\sim 3.0\times$
10240	0.37 ms	3.49 ms	$\sim 9.43\times$
20480	1.42 ms	13.89 ms	$\sim 9.78\times$
30720	3.17 ms	31.13 ms	$\sim 9.82\times$
40960	5.67 ms	55.23 ms	$\sim 9.74\times$
50176	8.39 ms	82.77 ms	$\sim 9.86\times$
60416	12.05 ms	119.98 ms	$\sim 9.95\times$

#### IV. CPU VERSION OF CAT USING AMX

We also implemented CAT on CPU to work with the recent CPU Advanced Matrix Extensions (AMX) [37] in parallel. Support for the AMX instructions began just recently in 2023 with the 4th Generation Intel Xeon CPUs (Sapphire Rapids), where each CPU core has a physical AMX region that can do hardware-accelerated matrix multiplications. AMX instructions have proven to be very competitive in performance for dealing with inference AI workloads [38]. This CPU version of CAT, named CAT-CPU<sub>AMX</sub>, was implemented with OpenMP for the thread-level parallelization, and each thread was programmed to execute manual AMX instructions provided by `gcc intrinsics` header, in a very similar way to how tensor cores are programmed with the `wmma` CUDA interface. The only relevant difference is that AMX fragments are currently of size  $16 \times 64$  elements only, and CAT needs square fragments. Therefore, this CPU variant executes the CAT scheme on a super-fragment of  $64 \times 64$  elements by stacking four  $16 \times 64$  AMX fragments<sup>7</sup>. CAT-CPU<sub>AMX</sub> also includes `AVX256` instructions<sup>8</sup> for the transition function, handling groups of 8 cells at a time on each thread.

In order to have a reference point for CAT-CPU, this approach was compared to other known CPU approaches for CA simulation, such as i) classic OMP (OpenMP), ii) OMP + AVX256, and iii) OMP + AVX512. Table V shows the performance of all these CPU approaches running on a 48-core Intel(R) Xeon(R) Platinum 8488C CPU. The results show that at  $r = 4$  CAT-CPU<sub>AMX</sub> already outperforms all other approaches, and at higher radius it is even an order of magnitude faster than the second fastest one.

TABLE V  
CAT-CPU<sub>AMX</sub> VERSUS OTHER CPU APPROACHES. AVERAGE TIME (MS)  
PER SIMULATION STEP AT  $n = 60416$ .

$r$	OMP	OMP + AVX256	OMP + AVX512	CAT-CPU <sub>AMX</sub>
1	623 ms	265 ms	268 ms	<b>355 ms</b>
4	3999 ms	488 ms	374 ms	<b>354 ms</b>
8	34955 ms	3879 ms	2128 ms	<b>356 ms</b>
16	115156 ms	14586 ms	7862 ms	<b>357 ms</b>

CAT-CPU<sub>AMX</sub> is completely evaluated in the next section together with the original GPU version and the other state-of-the-art GPU methods.

<sup>7</sup>A second option was to only use  $16 \times 16$  of the  $16 \times 64$  elements, but this was not faster than the super fragment of  $64 \times 64$  cells.

<sup>8</sup>AMX + AVX256 ran faster than AMX + AVX512.

## V. EXPERIMENTAL EVALUATION

## A. Experimental Design

Experiments consist of measuring the time, speedup, power and energy efficiency of CAT in both its GPU and CPU versions, and other state of the art GPU implementations at simulating cellular automata simulation of  $n \times n$  cells with neighborhood radius between  $1 \leq r \leq 16$ . These CA are initialized with a random uniform distribution of living cells, with density  $\delta$  that varies according to the radius, and periodic boundary conditions.

1) *Chosen CA tests*: We chose instances of the Larger Than Life (LTL) family of CA [32], [14], [4], [39], which is the generalization of the well-known Game of Life (GoL) to any neighborhood radius  $r$  with its own survival/birth rules. The neighborhood type is Moore, and the rule sets have been specified to produce complex phenomena for each value of  $r$ . The standard notation for defining LTL instances of CA is

$$R_r C_c M_{s_1..s_2} B_{b_1..b_2} N_n \quad (20)$$

where

- $R_r$ : Neighborhood radius  $r$ .
- $C_c$ : Number of states ( $c$ ) per cell.
- $M_m$ : Include ( $m = 1$ ) or not ( $m = 0$ ) the center cell.
- $S_{s_1..s_2}$ : Survival range  $[s_1..s_2]$  for living cells.
- $B_{b_1..b_2}$ : Birth range  $[b_1..b_2]$  for dead cells.
- $N_n$ : Moore ( $n = M$ ) or Von Neumann ( $n = N$ ) neighborhood.

The LTL instances used for the tests are presented in Table VI, each one using a certain value of  $r$  and an initial density  $\delta$  of living cells. Some of these LTL instances are known in the literature and cited, such as Bosco's rule by Evans [40], or Bugsmovie by Griffeath [32], among others. The other CA instances in the table are custom variations made by our team, with their birth and survival ranges adapted to the specific value of  $r$ . It is worth clarifying that choosing one or another

TABLE VI  
LARGER THAN LIFE INSTANCES FOR TESTING AT DIFFERENT  $r$ .

$r$	LTL Instance	$\delta$	Name
1	$R_1 C_2 M_0 S_{2..3} B_{3..3} N_M$	0.07	Game of Life [20]
2	$R_2 C_2 M_0 S_{7..12} B_{8..11} N_M$	0.15	Starry Night (custom)
3	$R_3 C_2 M_0 S_{15..23} B_{14..17} N_M$	0.25	Boiling Gnocchi (custom)
4	$R_4 C_2 M_0 S_{40..80} B_{41..80} N_M$	0.50	Majority [32]
5	$R_5 C_2 M_0 S_{35..59} B_{34..45} N_M$	0.21	Bosco's Rule [40]
6	$R_6 C_2 M_0 S_{49..81} B_{46..65} N_M$	0.22	Radiation (custom)
7	$R_7 C_2 M_0 S_{101..201} B_{75..170} N_M$	0.29	Waffles [41]
8	$R_8 C_2 M_0 S_{163..223} B_{74..252} N_M$	0.23	Globe [41]
9	$R_9 C_2 M_0 S_{108..181} B_{100..140} N_M$	0.24	Gravity (custom)
10	$R_{10} C_2 M_0 S_{122..211} B_{123..170} N_M$	0.25	Bugsmovie [32]
11	$R_{11} C_2 M_0 S_{156..265} B_{147..205} N_M$	0.24	Broken Ships (custom)
12	$R_{12} C_2 M_0 S_{170..296} B_{170..240} N_M$	0.25	Scaled GOL [32]
13	$R_{13} C_2 M_0 S_{213..364} B_{203..283} N_M$	0.25	The Cleansing (custom)
14	$R_{14} C_2 M_0 S_{245..420} B_{234..326} N_M$	0.25	Scaled Bugsmovie [32]
15	$R_{15} C_2 M_0 S_{170..296} B_{170..240} N_M$	0.28	Pretzels [32]
16	$R_{16} C_2 M_0 S_{170..296} B_{170..300} N_M$	0.26	Tangy Ramen (Custom)

LTL instance does not affect the performance of the GPU approaches, as they are only affected by  $n$  and  $r$ .

TABLE VII  
GPU APPROACHES SELECTED FOR EVALUATION

Approach	Main idea and GPU implementation details
BASE	Baseline global memory approach. • Implemented by our team. • One cell ( <code>char</code> ) per thread in global memory. • CUDA Block Size: $B_x \times B_y = 32 \times 32$ threads. • Neighborhood radius: $r \in [1..16]$ .
SHARED	Classic Shared Memory approach [24]. • By Millan et al. [11] supporting $r \in [1..5]$ . • <i>Out-halo</i> : sh-mem of $(B_x + 2r) \times (B_y + 2r)$ cells. • One cell ( <code>char</code> ) per thread. • CUDA Block Size: $B_x \times B_y = 32 \times 32$ threads. • <b>[us]</b> Radius extended to $r \in [1..16]$ .
CAT (proposed)	Proposed Tensor Core (TC) based approach. • Implemented by our team. • Neighborhood reduction through Tensor Core MMAs. • Multiple TC fragments (FP16) per warp. • <i>Out-halo</i> : adjacent fragments for boundary cells. • Uses sh-mem for intermediate results. • CUDA Block Size: $B_x \times B_y = 16 \times 16$ threads. • Neighborhood radius: $r \in [1..16]$ .
CAT-CPU <sub>AMX</sub> (proposed)	CPU Version of CAT. • Implemented by our team. • Same idea of CAT, but with CPU AMX Instructions. • OMP (OpenMP) for doing AMX on each CPU core. • Super-fragments of $64 \times 64$ cells. • Neighborhood radius: $r \in [1..64]$ .
COARSE	Alternative shared memory approach. • Implemented by our team. • Thread Coarsening [42] and ( <code>char</code> ) cells. • <i>Out-halo</i> : sh-mem of $(80 + 2r) \times (80 + 2r)$ cells. • Neighborhood radius: $r \in [1..16]$ .
MCELL	Multi-cell + shared memory approach. • Original code by Millan et al. [11], supporting $r \in [1..5]$ . • Two adjacent cells ( <code>char</code> ) per thread for register re-use. • CUDA Block Size: $B_x \times B_y = 16 \times 16$ threads. • <b>[us]</b> <i>Out-halo</i> : sh-mem of $(2B_x + 2r) \times (B_y + 2r)$ cells. • <b>[us]</b> Improved memory access on two-cell neighborhood. • <b>[us]</b> Radius extended to $r \in [1..16]$ .
PACK	Packet coding technique. • Idea and code by Cagigas et al. [12] supporting $r = 1$ . • 64-bit words codified as eight <code>char</code> (8-bit) cells. • No use of shared memory $\rightarrow$ faster in global memory. • CUDA Block Size: $B_x \times B_y = 16 \times 16$ threads. • <b>[us]</b> Radius extended to $r \in [1..16]$ .

2) *Approaches Selected for Evaluation*: Table VII lists the GPU implementations that were selected for performance evaluation, with their source codes available at <https://github.com/temporal-hpc/CAT>.

BASE, CAT, CAT-CPU<sub>AMX</sub> and COARSE were implemented by our team, using CUDA C/C++ for the GPU methods and gcc C/C++ with intrinsics for the CAT-CPU<sub>AMX</sub> code. BASE is a GPU baseline and corresponds to a standard GPU implementations of stencil computation [5] where each thread simulates one cell reading its entire neighborhood directly from global memory. COARSE is a shared memory approach with an *out-halo* design that also includes thread coarsening [42]. It uses a large shared memory tile of  $(80 + 2r) \times (80 + 2r)$  cells per block, and each thread simulates multiple cells with a stride of the block size. CAT is the proposed method and CAT-CPU<sub>AMX</sub> is the CPU version of CAT using AMX instructions.

SHARED, MCELL and PACK correspond to related works that made their implementation available using CUDA C/C++ as well. We extended these implementations to support  $1 \leq r \leq 16$  and also optimized some of them. For SHARED, the core logic of the technique was extended to  $r \in [1..16]$  with the expected code changes towards radius generalization. For MCELL the case was less straightforward, extending it

to  $r \in [1..16]$  required changing the design from *in-halo* to *out-halo* in order to run efficient in the range of  $r$ , and also required improving the memory access such that the two-cell neighborhood is read once for each thread at any  $r$ . In the case of PACK, the extension of Cagigas et al. implementation [12] to  $r \in [1..16]$  followed their scheme but now considering that when  $8 < r \leq 16$  the horizontal neighbor cells are found in the two consecutive 64-bit words to the left/right, and not one as in the original case. It is worth mentioning that these extended implementations run as fast, or faster than before when executed using their original supported radiuses. More details on the core idea of the last three approaches can be found in Section II.

In terms of precision, CAT uses FP16 types for the neighborhood reduction, as it is currently the most suitable type offered by GPU tensor cores in which square fragments can be defined (*i.e.*, full INT32 MMAs are still not supported in square shape). Still, converting the floating point reduction into INT32 provided correct simulations throughout the range  $1 \leq r \leq 16$ . The reason why this works is because FP16 has a 10 bit mantissa, which corresponds to a precision of  $10 + 1$ , then  $2^{11}$  is the largest integer that can be precisely represented, which is sufficient to store the maximum possible amount of living neighbors with radius  $r = 16$  (*i.e.*  $33 \times 33 - 1$  cells). In the case of CAT-CPU<sub>AMX</sub> the matrix multiplication uses INT8 in the operands and INT32 in the results, matching the precision requirements. For the rest of the approaches, they employ INT32 precision for neighbor counting and *char* precision for the cell state.

3) *Performance measures*: The performance measures are *time* in milliseconds and *speedup* with respect to the BASE approach. Each measure is taken as the average of multiple realizations, varying  $n$  in the range  $n \in [1024, 60416]$ . For the lower values of  $n$  we used up to 32 realizations while for high  $n$  we used up to four. Each realization measures the average running time for simulating the given LTL instance, at a given  $n$ , for 25 time steps. These measuring settings ensured averages with a standard error of 1% or less.

The energy measurements were done by running a large size simulation of  $n \times n = 60416 \times 60416$  cells and measuring the instant GPU Power in Watts (W) as a time series for 1000 simulation steps, using the *nvml* library from Nvidia. From these GPU power time series, the average total energy per simulation step is obtained in Joules (J), and the average energy efficiency per simulation step is obtained as *Cells/J*.

4) *Testing Platform*: Table VIII lists the testing platform. As for software, we used CUDA 11.8, NVML 550.54.15 and

TABLE VIII  
TESTING PLATFORM

Property	Value
OS	Linux Ubuntu 24.04 LTS
CPU	48-core Intel Xeon Platinum 8488C
RAM	251 GB RAM
GPU	NVIDIA H100 SXM5 80GB HBM3

GCC 11.4, except for CAT-CPU<sub>AMX</sub> that used GCC 13.2.

## B. Performance Results

Figure 8 presents the time and speedup of the selected approaches. The first row shows the average time per simulation for radiuses  $r = 1, 4, 8, 16$ , while the second row shows the speedups with respect to BASE. The results show that CAT keep's its running time constant as  $r$  increases while the other approaches take more time to complete. This behavior makes CAT's speedup increase with  $r$ , starting as the second fastest at  $r = 1$  with a speedup of  $1.3\times$ , but becoming the fastest one at  $r = 4, 8, 16$  with speedups of  $9\times$ ,  $27\times$ , and  $101\times$  respectively. These results also agree with the cost model and its theoretical speedups from Table II when the full GH100 chip was assumed. For the other approaches, one can observe that their speedups cluster into three groups as  $r$  increases; the first top group with just CAT, the second group with PACK followed by MCELL, and the third group with COARSE followed by SHARED. It is worth noting that PACK is the fastest approach at  $r = 1$  with  $\sim 3\times$  of speedup, which translates to being  $2.3\times$  faster than CAT. On the other hand, at  $r = 16$  CAT is  $\sim 14\times$  faster than PACK. Another behavior to note is that COARSE performs better than SHARED as a shared memory based solution, and it avoids being slower than the BASE at  $r = 1$  as SHARED did, which is a known issue of shared memory solutions in modern GPUs [26], [27]. This difference between COARSE and SHARED has less of an impact at higher  $r$ , where both are much closer in performance. For CAT-CPU<sub>AMX</sub>, it is worth noticing that it behaves just like CAT but at a slower level of performance. Still, at  $r = 16$  it is up to  $3.4\times$  faster than BASE, and even outperforms COARSE and SHARED.

More details are given in Figure 9, which shows the performance impact from increasing the neighborhood radius  $r$ , given a large CA of  $n \times n = 60416 \times 60416$  cells. The plots show that  $r = 3$  is the crossover point where CAT surpasses PACK and becomes the fastest approach for the remaining range of  $r$ . Coincidentally, it is also the crossover zone where MCELL surpasses COARSE, and where SHARED surpasses BASE. Another behavior to note is that CAT's speedup increases with  $r$  being up to two orders of magnitude faster than BASE, whereas the other approaches converge at specific values, with PACK being the second fastest approach reaching near  $7\times$  of speedup. In the case of CAT-CPU<sub>AMX</sub>, it has a crossover point of  $r = 8$  where it is faster than BASE, and near  $r = 14$  it becomes faster than SHARED and COARSE.

## C. Energy Efficiency

Figure 10 presents the power consumption and energy efficiency of all approaches, for radiuses  $r = 1, 4, 8, 16$ . For the power consumption plots, CAT's energetic behavior is a short high-power curve that peaks near 700W which is the TDP of the H100 GPU. At  $r = 1$ , CAT's energy efficiency is in the same range as the other approaches, except for PACK, which is much more energy efficient by a great margin. This can be explained by the bit-level logic of PACK, which is highly efficient. However, as  $r$  increases, CAT's power consumption curve remains roughly the same, while the other

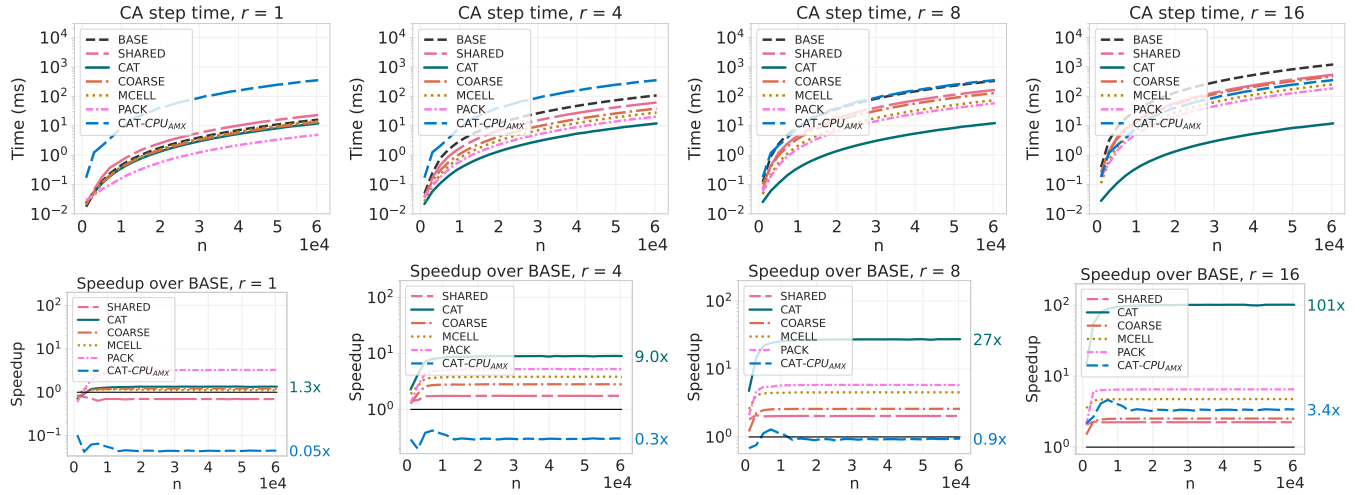


Fig. 8. Time and Speedup for all approaches at different values of neighborhood  $r$ .




Fig. 9. Impact of the neighborhood radius  $r$  in the time and speedup of all approaches for a large CA of  $n \times n = 60416 \times 60416$  cells.

approaches sustain their consumption for longer, thus using more energy. This difference is most noticeable at  $r = 16$ , where CAT can be up to  $6.45\times$  more energy efficient than the second best, and an order of magnitude more efficient than the rest. For the other approaches, their power curves tend to transform from pulse-like shapes to plateau-like ones. In the case of CAT-CPU<sub>AMX</sub>, its behavior is a low-power curve with high total energy but constant with  $r$ . This makes it a competitive approach at high radius such as  $r = 16$ .

Figure 11 shows more in detail how increasing  $r$  affects the total energy used and energy efficiency of each approach. The total energy of CAT remains almost unchanged throughout the range of  $r$ , while the other approaches use more energy with higher  $r$ . In terms of energy efficiency (cells per Joule),  $r = [4..5]$  is the crossover zone where CAT's surpass PACK and become the most energy efficient approach. Below this value, PACK is the most energy efficient one. As for the other approaches, in general, they are more energy efficient than BASE, although their differences are smaller at high  $r$ . For CAT-CPU<sub>AMX</sub>, it shows a constant energy consumption just as CAT for GPU, but at one order of magnitude higher. For this CPU version two crossover points can be identified along  $r$ , one at  $r = 6$ , where it becomes more energy efficient than BASE, and another near  $r = 14$ , where it becomes more energy efficient than SHARED and COARSE.

#### D. Scaling with GPU Generations

The first generation of tensor cores was introduced back in 2017 with NVIDIA's V100 GPU, which had the Volta architecture. Since then, different tensor core generations were released, with improved performance and features. The second generation of tensor cores was introduced in 2018 with the Turing architecture, focused on videogames (e.g., the TITAN RTX GPU) with no data-center variant. The third generation was released in 2020 with the Ampere architecture (e.g., the A100 GPU), and the fourth generation in 2022 with the Hopper architecture (e.g., the H100 GPU). Many other aspects also got improved on each generational jump, such as the number of traditional cores (FP32/INT32 units), memory capacity, memory bandwidth, among others. These improvements make both classical and tensor core based GPU implementations automatically scale their performance by just using newer hardware. The scaling factors provided by the last GPU architectures can give key insights on what performance future GPU architectures could bring to CAT in comparison to the other approaches.

Figure 12 presents the scaling factors of CAT and all other approaches, relative to their performance on a Volta V100 GPU. Overall, CAT shows the highest scaling factors across GPU architectures, specially when switching from the A100 to the H100 where it achieves a scaling of near  $3.3\times$  in all cases. This jump in performance is in agreement with the 2-year technological improvement that the H100 GPU provides in comparison to the A100. A particular behavior is noted at  $r = 1$  where SHARED exhibits the highest scaling when jumping from the V100 to the A100. But in general, the transition from Ampere to Hopper is the one that provides the highest jump, and tensor core performance has been scaling at a higher rate than regular GPU compute. It is very likely that this scaling trend could continue given the importance of AI these days, favoring CAT with a scaling rate that is higher than traditional GPU approaches.



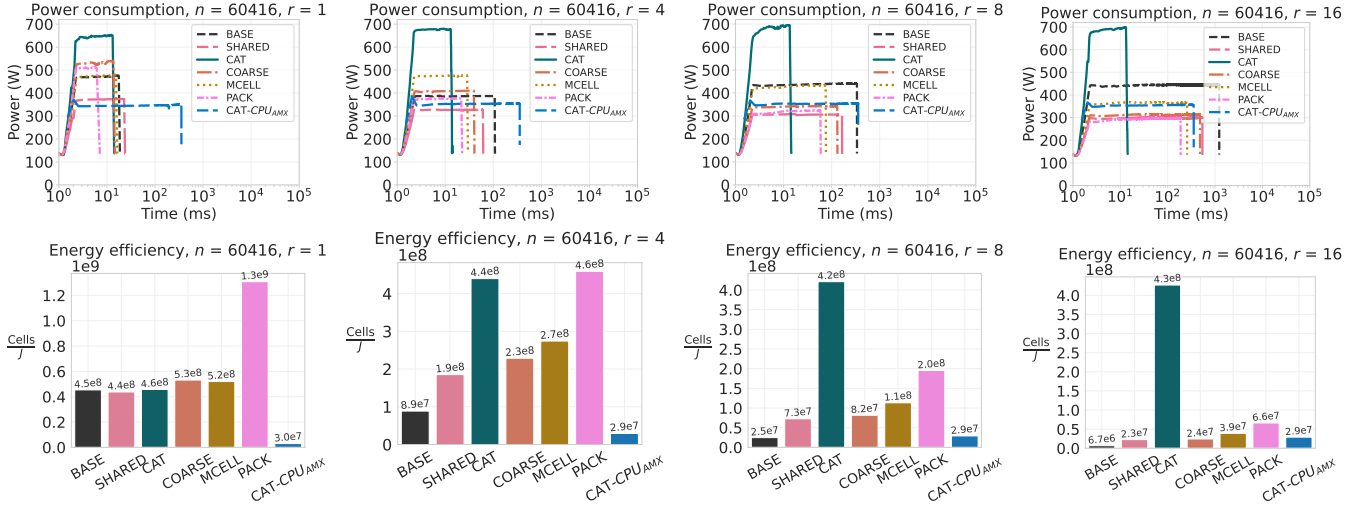


Fig. 10. Power time series and energy efficiency for all approaches at different values of neighborhood  $r$ .




Fig. 11. Impact of the neighborhood radius  $r$  in the total energy and energy efficiency for a large CA of  $n \times n = 60416 \times 60416$  cells.

## VI. DISCUSSION AND CONCLUSIONS

This work presented CAT (Cellular Automata on Tensor Cores), a GPU approach that uses tensor cores to accelerate the simulation of cellular automata (CA) at large neighborhood radius. Because of its design based on computing the neighborhood reduction via Matrix Multiply Accumulate (MMA) operations, CAT's performance is unaffected by the increase in neighborhood radius in the range  $1 \leq r \leq 16$ , while other state-of-the-art approaches do increase their work (hence time) as  $r$  increases. CAT can be employed on any CA model where the cell transition function acts over a weighted summation of its neighbors, such as the life-like CA models.

Experiments using the Larger Than Life (LTL) family of Cellular Automata (CA) as case studies showed that CAT is the fastest approach of all in the range  $3 \leq r \leq 16$ ; reaching up to  $101\times$  of speedup over a traditional baseline approach and up to  $\sim 14\times$  faster than the fastest state-of-the-art GPU method (PACK). For the low range  $1 \leq r \leq 2$ , CAT is still not the fastest approach, but it is still competitive, being only behind the GPU packet coding approach (PACK) which is efficient at low neighborhood radius. In terms of energy usage and efficiency, CAT is the most energy-efficient approach in the range  $5 \leq r \leq 16$ ; up to  $6.45\times$  more efficient than the second most efficient approach (PACK), and up to an order of magnitude more efficient than the rest.

In terms of performance scaling across GPU architectures, the last two tensor core generations (Ampere and Hopper) have provided significant scaling factors to CAT, which are much higher than the ones observed for the other approaches. Assuming that least part of this trend continues in the future because of AI, supported by the fact that AI is driving the evolution of GPUs towards more tensor core performance, then CAT becomes a promising approach for upcoming GPU architectures.

We also implemented a CPU version of CAT, named CAT-CPU<sub>AMX</sub>, which uses recently introduced AMX instructions that accelerate matrix multiplication on the CPU in a very similar way as how GPU tensor cores work. The results showed that although CAT-CPU<sub>AMX</sub> is not as fast as its GPU version, it is still highly competitive and can outperform some GPU approaches in a large neighborhood radius. Furthermore, the performance of AMX instructions is progressively improving on each CPU generation, thus it will become an even more attractive resource in the future years.

As for future work, CAT can be further improved and extended in several aspects. For instance, caching the tiles of  $\Lambda$  into *shared memory* should make CAT run even faster; a preliminary version was attempted but reported slower performance than the actual version of CAT. Further research and experimentation is required to reach a proper caching scheme of the  $\Lambda$  tile. Another technical improvement for CAT is to support more neighborhood types; currently CAT's design supports *Moore* and *Simplified Von Neumann* neighborhoods. A third neighborhood that is often required is the original *Von Neumann* neighborhood, which has a diamond shape. Supporting this neighborhood would require redefining the nonzero entries of the band matrices to capture the diamond-shaped neighborhood, as well as verifying if the transition from step 1 to step 2 of MMAs should accumulate or multiply the previous result. A third technical improvement is to take advantage of the low precision types currently supported by tensor cores, such as INT8 which is faster than the FP16 types currently used, or the experimental types BIT, INT4, which are even

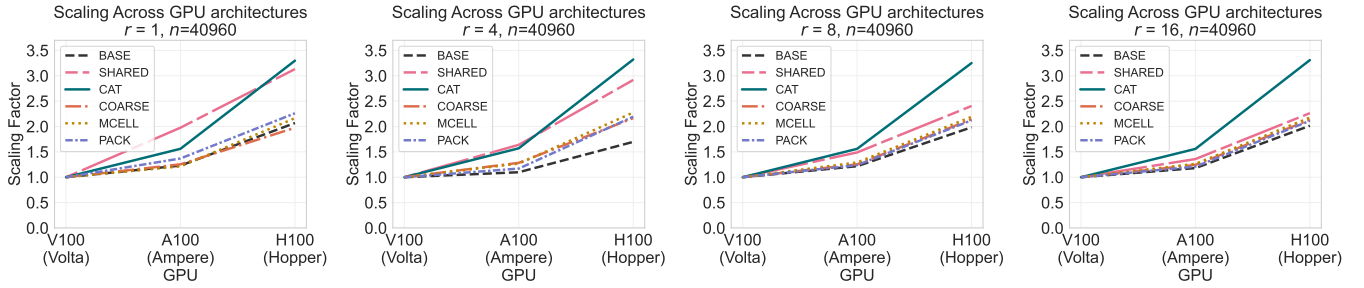


Fig. 12. Scaling factors of CAT and all other approaches with respect to their performance on the Volta V100.

faster. Some of these types are only supported by rectangular shaped fragments, but the square shaped requirement can still be accomplished by stacking several of these fragments until a major square region is built again. Currently, what penalizes this last idea is the need of converting all  $\text{INT}_{32}$  results back again into low precision types in between the two major steps of CAT. Finding a faster way of re-using the resulting  $\text{INT}_{32}$  fragment back as a product operand for another MMA could indeed accelerate CAT significantly.

CAT's neighborhood radius range can be extended beyond  $1 \leq r \leq 16$ ; it only requires to expand the main idea presented in Figure 4 to consider the next closest fragments at each reduction direction, i.e.,  $F_{i \pm 2, j \pm 2}$ . This expansion implies having six band fragments  $\pi_1, \pi_2, \dots, \pi_6$  instead of three, as well as to widen the global halo in one more fragment. All of these changes would make CAT support neighborhood radiuses in the range  $1 \leq r \leq 32$ . In general, successive applications of this expansion would make CAT increase its supported radius range by +16. Each new expansion would put CAT's running time at a higher plateau of running time, but would remain constant for that new supported range. The challenge however is that  $\text{FP}_{16}$  may not be precise enough to store the maximum amount of neighbors for such large radiuses. Possible solutions could be to normalize the range of values, which would push the limitation to appear in a larger  $r$ , or a more fundamental solution is that future GPUs could support full  $\text{INT}_{32}$  MMAs.

Lastly, exploring the 3D extension of CAT is a natural extension to continue research. The challenge would be to find ways to represent 3D reductions with tensor cores that operate on 2D MMAs. Reusing the main idea of CAT, now by layers is feasible, but other new approaches could also be explored and tested, especially given that generic arithmetic reductions are known to run efficiently on tensor cores [43], opening the possibility of reducing entire volumes of cells with a few MMAs.

To conclude, this work has shown that tensor cores can indeed accelerate non-AI applications such as the simulation of cellular automata, and not only GPU tensor cores, but also the recently introduced CPU AMX instructions which for practical purposes can be thought of as the CPU tensor cores. CAT, both in GPU or CPU, can be of great interest for research teams that need to study complex phenomena that emerge from CA with a large neighborhood radius.

## ACKNOWLEDGMENT

This work was supported by the ANID FONDECYT grants #1221357, #1241596, the Temporal research lab, and the Patagón Supercomputer of Austral University of Chile (FONDEQUIP EQM180042). The authors thank Roberto Melita, who brought rich ideas and discussions in the early stages of this investigation.

## REFERENCES

- [1] E. F. Codd, *Cellular automata*. Academic press, 2014.
- [2] A. Adamatzky, *Game of life cellular automata*. Springer, 2010, vol. 1.
- [3] C. G. Langton, "Studying artificial life with cellular automata," *Physica D: Nonlinear Phenomena*, vol. 22, no. 1-3, pp. 120–149, 1986.
- [4] K. M. Evans, "Larger than life: Digital creatures in a family of two-dimensional cellular automata," *Discrete Mathematics & Theoretical Computer Science*, no. Proceedings, 2001.
- [5] D. D'Ambrosio, G. Filippone, R. Rongo, W. Spataro, and G. A. Trunfio, "Cellular automata and gpgpu: an application to lava flow modeling," *International Journal of Grid and High Performance Computing (IJGHPC)*, vol. 4, no. 3, pp. 30–47, 2012.
- [6] P. Hogeweg, "Cellular automata as a paradigm for ecological modeling," *Applied mathematics and computation*, vol. 27, no. 1, pp. 81–100, 1988.
- [7] M. Batty, H. Couclelis, and M. Eichen, "Urban systems as cellular automata," pp. 159–164, 1997.
- [8] M. J. Gibson, E. C. Keedwell, and D. A. Savić, "An investigation of the efficient implementation of cellular automata on multi-core cpu and gpu hardware," *Journal of Parallel and Distributed Computing*, vol. 77, pp. 11–25, 2015.
- [9] J. Holewinski, L.-N. Pouchet, and P. Sadayappan, "High-performance code generation for stencil computations on gpu architectures," in *Proceedings of the 26th ACM international conference on Supercomputing*, 2012, pp. 311–320.
- [10] Y. Zhang, J. Zhou, Y. Yin, X. Shen, and X. Ji, "Multi-gpu implementation of a cellular automaton model for dendritic growth of binary alloy," *Journal of Materials Research and Technology*, vol. 14, pp. 1862–1872, 2021.
- [11] E. N. Millán, N. Wolovick, M. F. Piccoli, C. G. Garino, and E. M. Bringa, "Performance analysis and comparison of cellular automata gpu implementations," *Cluster Computing*, vol. 20, no. 3, pp. 2763–2777, 2017.
- [12] D. Cagigas-Muñiz, F. Diaz-del Rio, J. L. Sevillano-Ramos, and J.-L. Guisado-Lizar, "Efficient simulation execution of cellular automata on gpu," *Simulation Modelling Practice and Theory*, vol. 118, p. 102519, 2022.
- [13] T. Fujita, D. Nishikori, K. Nakano, and Y. Ito, "Efficient gpu implementations for the conway's game of life," in *2015 Third International Symposium on Computing and Networking (CANDAR)*. IEEE, 2015, pp. 11–20.
- [14] K. M. Evans, *Larger than Life: It's so nonlinear*. The University of Wisconsin-Madison, 1996.
- [15] S. Markidis, S. W. Der Chien, E. Laure, I. B. Peng, and J. S. Vetter, "Nvidia tensor core programmability, performance & precision," in *2018 IEEE international parallel and distributed processing symposium workshops (IPDPSW)*. IEEE, 2018, pp. 522–531.

- [16] D. Narayanan, M. Shoenybi, J. Casper, P. LeGresley, M. Patwary, V. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro *et al.*, "Efficient large-scale language model training on gpu clusters using megatron-lm," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021, pp. 1–15.
- [17] X. Wang, Y. Wei, Y. Xiong, G. Huang, X. Qian, Y. Ding, M. Wang, and L. Li, "Lightseq2: Accelerated training for transformer-based models on gpus," in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2022, pp. 1–14.
- [18] S. Rybacki, J. Himmelsbach, and A. M. Uhrmacher, "Experiments with single core, multi-core, and gpu based computation of cellular automata," in *2009 first international conference on advances in system simulation*. IEEE, 2009, pp. 62–67.
- [19] G. Oxman, S. Weiss, and Y. Be'ery, "Computational methods for conway's game of life cellular automaton," *Journal of Computational Science*, vol. 5, no. 1, pp. 24–31, 2014.
- [20] M. Gardner, "The fantastic combinations of john conway's new solitaire game 'life,'" *Sc. Am.*, vol. 223, pp. 20–123, 1970.
- [21] L. Zhang, M. Wahib, P. Chen, J. Meng, X. Wang, T. Endo, and S. Matsuoka, "Revisiting temporal blocking stencil optimizations," in *Proceedings of the 37th International Conference on Supercomputing*, 2023, pp. 251–263.
- [22] J. Balasalle, M. A. Lopez, and M. J. Rutherford, "Optimizing memory access patterns for cellular automata on gpus," in *GPU Computing Gems Jade Edition*. Elsevier, 2012, pp. 67–75.
- [23] J. Tran, D. Jordan, and D. Luebke, "New challenges for cellular automata simulation on the gpu," *SIGGRAPH, Los Angeles. ACM. Poster*, 2004.
- [24] P. Topa and P. Mlcek, "Using shared memory as a cache in high performance cellular automata water flow simulations," *Computer Science*, vol. 14, no. 3, pp. 385–385, 2013.
- [25] P. Topa, "Cellular automata model tuned for efficient computation on gpu with global memory cache," in *2014 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*. IEEE, 2014, pp. 380–383.
- [26] A. Schäfer and D. Fey, "High performance stencil code algorithms for gpgpus," *Procedia Computer Science*, vol. 4, pp. 2027–2036, 2011.
- [27] P. Renc, T. Pećak, A. De Rango, W. Spataro, G. Mendicino, and J. Waş, "Towards efficient gpgpu cellular automata model implementation using persistent active cells," *Journal of Computational Science*, vol. 59, p. 101538, 2022.
- [28] K. Matsumura, H. R. Zohouri, M. Wahib, T. Endo, and S. Matsuoka, "An5d: automated stencil framework for high-degree temporal blocking on gpus," in *Proceedings of the 18th ACM/IEEE International Symposium on Code Generation and Optimization*, 2020, pp. 199–211.
- [29] H. Zhuang, X. Liu, X. Liang, Y. Yan, J. He, Y. Cai, C. Wu, X. Zhang, and H. Zhang, "Tensor-ca: A high-performance cellular automata model for land use simulation based on vectorization and gpu," *Transactions in GIS*, vol. 26, no. 2, pp. 755–778, 2022.
- [30] X. Liu, Y. Liu, H. Yang, J. Liao, M. Li, Z. Luan, and D. Qian, "Toward accelerated stencil computation by adapting tensor core unit on gpu," in *Proceedings of the 36th ACM International Conference on Supercomputing*, 2022, pp. 1–12.
- [31] Y. Chen, K. Li, Y. Wang, D. Bai, L. Wang, L. Ma, L. Yuan, Y. Zhang, T. Cao, and M. Yang, "Convstencil: Transform stencil computation to matrix multiplication on tensor cores," in *Proceedings of the 29th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 333–347.
- [32] D. Griffiths, "Self-organization of random cellular automata: four snapshots," *Probability and phase transition*, pp. 49–67, 1994.
- [33] N. Corporation, "Nvidia deep learning performance," <https://docs.nvidia.com/deeplearning/performance/dl-performance-convolutional/index.html>, accessed: 2024-11-9.
- [34] J. Choi, H. Kwon, W. Lee, J. Choi, and J. Lim, "Learning from distinctive candidates to optimize reduced-precision convolution program on tensor cores," *arXiv preprint arXiv:2202.06819*, 2022.
- [35] R. M. Karp, *A survey of parallel algorithms for shared-memory machines*. University of California at Berkeley, 1988.
- [36] NVIDIA, "Nvidia h100 tensor core gpu architecture," NVIDIA, May, 2024. [Online]. Available: <https://resources.nvidia.com/en-us-tensor-core/gtc22-whitepaper-hopper>
- [37] I. Corporation, "Accelerate artificial intelligence workloads with intel® advanced matrix extensions," <https://www.intel.com/content/www/us/en/content-details/785250>, accessed: 2024-11-2.
- [38] H. Kim, G. Ye, N. Wang, A. Yazdanbakhsh, and N. S. Kim, "Exploiting intel® advanced matrix extensions (amx) for large language model inference," *IEEE Computer Architecture Letters*, 2024.
- [39] C. T. Bekaroglu, "Analyzing dynamics of larger than life: Impacts of rule parameters on the evolution of a bug's geometry," Ph.D. dissertation, CALIFORNIA STATE UNIVERSITY, NORTHRIDGE, 2023.
- [40] K. M. Evans, "Is bosco's rule universal?" in *Machines, Computations, and Universality: 4th International Conference, MCU 2004, Saint Petersburg, Russia, September 21-24, 2004, Revised Selected Papers 4*. Springer, 2005, pp. 188–199.
- [41] M. Wojtowicz, "Larger than life family, cellular automata rules lexicon," [http://www.mirekw.com/ca/rullex\\_lgtl.html](http://www.mirekw.com/ca/rullex_lgtl.html), accessed: 2024-05-14.
- [42] A. Magni, C. Dubach, and M. F. P. O'Boyle, "A large-scale cross-architecture evaluation of thread-coarsening," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC '13. New York, NY, USA: Association for Computing Machinery, 2013.
- [43] C. A. Navarro, R. J. Barrientos, J. A. Riquelme, and R. Vega, "Gpu tensor cores for fast arithmetic reductions," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 1, pp. 72–84, 2020.



**Cristóbal A. Navarro** has a Ph.D. degree in computer science from the University of Chile (2015). Currently, he is an associate professor at the Universidad Austral de Chile and leads the *Temporal* research lab as well as the Patagón Supercomputer project. Today, his research interests include GPU computing, computer graphics, and computational physics.



**Felipe A. Quezada** has Bsc and Msc degrees in computer science with a focus on HPC and data science, both at Universidad Austral de Chile. He is currently working as a researcher in the Temporal Group led by Dr. Cristóbal A. Navarro and as an HPC engineer at the Patagón Supercomputer project. His interests are mainly focused on HPC, machine learning, and computer graphics.



**Enzo Meneses** holds Bsc and Msc degrees in computer science at the Universidad Austral de Chile. Currently he is a researcher of the Temporal group led by Dr. Cristóbal A. Navarro, where new GPU techniques are developed. His interests are mainly parallel algorithms and data structures, and lately the use of RT Cores beyond ray tracing.



**Héctor Ferrada** received his Ph.D. degree in Computer Science from the University of Chile in 2016, focusing on his research in the design and analysis of algorithms for compact data structures. In 2016–2017, he conducted postdoctoral research in the Genome Scaling Algorithms Group at the University of Helsinki, Finland, in collaboration with Dr. Veli Mäkinen. Currently, he mainly teaches courses related to his research interests in algorithms and data structures.



**Nancy Hitschfeld** received a PhD in Applied Sciences from the ETH-Zurich in 1993. Currently, she works as full professor at the Computer Science Department, University of Chile. Her research interests include geometric modeling, polygon and polyhedral meshes, and parallel algorithms, for problem solving in computational science and engineering applications.