We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,900
Open access books available

185,000
International authors and editors

200M
Downloads

Our authors are among the

154
Countries delivered to

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Accelerating 3D Cellular Automata Computation with GP-GPU in the Context of Integrative Biology

Jonathan Caux[1,2,3,4], Pridi Siregar[4] and David Hill[1,2,3]
*[1]Clermont Université, Université Blaise Pascal, LIMOS,
BP-10448, F-63000 Clermont-Ferrand
[2]CNRS, UMR 6158, LIMOS, ISIMA, Campus des Cézeaux, BP-10125, F-63173 Aubière
[3]ISIMA, Computer Science & Modelling Institute,BP10125, F-63177
[4]Integrative Biocomputing, Nouvelles Structures
- Place du Granier, 35135 Chantepie
France*

1. Introduction

Integrative biology has a large number of solutions available to simulate multi-scale models: Reaction-Diffusion Equations (Lu et al., 2001), Dissipative particle dynamics (Pivkin et al., 2009) or Finite element method (Sander et al., 2009). But cellular automata can also be used (von Neumann, 1966), this has been shown for instance to simulate a part of an organ behavior (Alarcón et al., 2004) or the depolarization and repolarization of the heart muscle cells which can be seen as an excitable medium (Siregar et al., 1998). In this chapter, we consider 3D cellular automata, consisting in a cube of cells where each cell can be in a specific state. A finite number of states are considered and a cell state evolves over time according to its current state, but also according to the states of the other cells in a defined neighborhood. When we simulate the evolution of an automaton, the future state of each cell is computed simultaneously according to the state of the cube cells at the current simulation time. With basic rules, it is possible to create a simplified model for natural behavior. The most common example is the Game of Life (Gardner, 1970). This model, created in 1970 by John Conway, allows simulating bacteria reproduction on a 2D toric space. In this model, a cell can only have two states: dead or alive, and evolves according to the number of cells alive in the contacting neighborhood (named Moore's neighborhood). But cellular automata can't be reduced to Conway's Game of Life; the interested reader can consult (Wolfram, 2002). For instance, electronic circuits can be simulated with the WireWorld cellular automata (Dewdney, 1990), and in biology, excitable medium, among others models, can also be simulated with cellular automata (Ermentrout & Edelstein-Keshet, 1993).

Even if cellular automata are mostly very simple, large 3D models with an important number of iterations require a very long computing time. Moreover, in integrative biology, computation of cellular automata is often only a part of a more complex problem. Working on larger cubes can therefore improve the result sensitivity, and performing this result faster can speed up the global process.

Our idea was to test cellular automaton implementation on GP-GPU (General Purpose Graphic Processing Unit)[1] (Luebke et al., 2004). This kind of processing unit was first designed to process graphics on a computer, but while the classical CPU computation performances evolution recently began to slow down, the GP-GPU continued to provide very significant speed-up. Ten years ago, developers of high performance computing applications started to port scientific software from CPU to GP-GPU to make the most of it (Trendall & Stewart, 2000). After the initial success, GPU manufacturers started to work on friendlier API for general purpose computation and we are now able to develop directly in languages which are close variants of the C language[2,3]. In the same way, other hardware accelerators like FPGA have been considered to speed up cellular automata computations (Woundeberg, 2006).

The main difficulty lies in the memory manipulation since GP-GPU have various levels of memory with different performances. The GP-GPU global memory which can be accessed by any thread at any time has very important access latency. The shared memory available for each streaming multiprocessor inside a GP-GPU does not have such latency, but this memory is only shared by threads running on the same multiprocessor. In addition, the number of concurrent threads in a streaming multiprocessor is limited. Moreover, memory transfer between the host computer and the GP-GPU device can severely imapir the global speed-up if the computation time is not significant enough in comparison with the data transfer time.

In this chapter, we propose two cellular automata implementations on GP-GPU with different memory usage (Part 3). To understand the benefits of GP-GPU, we compare the results obtained with an Nvidia Tesla C1060 board to a sequential implementation on 2 kinds of CPU (Xeon Core 2 and Nehalem) (Part 4).

2. Accelerate cellular automata computation

2.1 3D Cellular automata

A simple heart model can be created using a cellular automaton. After selecting an appropriate neighborhood (a Moore cubic-shaped neighborhood with a range of one for example), we can simulate the propagation of an electric stimulation with 3 states:

A cell is in an "activated" state over a predefined time interval. In this state, this cell stimulates its neighborhood cells which are in an "inactive" state. This state corresponds to cell depolarization.

A cell is in a "refractory" state for a certain predefined time. It can't be stimulated by an "activated" cell. This state corresponds to cell repolarization.

The rest of the time, a cell is in an "inactive" state.

But this heart model is far from perfect. A critical point is that the number of cells used will directly affect the model accuracy. If the whole heart is represented by a 10^3 cells cube, the accuracy will be low compared to the use of a cube with 100^3 cells. To compare both

---

[1]Four-Dimensional Cellular Automata Acceleration Using GPGPU:

http://eric_rollins.home.mindspring.com/gpgpu/index.html. Accession date: 20/01/2010.

[2]What is CUDA ? - Official CUDA web site: http://www.nvidia.co.uk/object/cuda_what_is_uk.html. Accession date: 20/01/2010.

[3]OpenCL Overview - Official OpenCL web site from Khronos : http://www.khronos.org/opencl/. Accession date: 20/01/2010.

possibilities, a ventricle will be represented by a cube of 3^3 cells in the first case whereas in the second case we can use 30^3 cells (27 cells compared to 27000). In the second case the finer grain-size of the 3D ventricle model will lead to a better approximation a real one. Getting a better accuracy cannot be reached only by a growth in cube size. Each time each dimension grows by a factor of 10, the total number of cells grows by a factor of 1000 (1000 cells compared to 1.000.000 cells in the second case with a 100^3 cube). Having more cells implies more computations and this significantly slows down the simulation performances. In addition, even if we assume that a 100^3 cells cube is a good size to represent a heart and that we have enough computing power, it may be interesting to take into account other features in order to have a more realistic model. However, these improvements will increase the computation load. For example, the state corresponding to the cell repolarization ("refractory") can be subdivided into two parts: a first one called the "absolute refractory state", where the cell can't be reactivated by an incoming impulse and a second, called the "relative refractory state" where it can. This small modification will only slightly increase computing load but more complex modifications will likely have wider impact on computing time.

One way or another, if the modeling technique retained is based on cellular automata, increasing the computation speed could be a good way to improve the model representation without damaging the global simulation speed and therefore achieve better accuracy.

2.2 GP-GPU utilization

In order to accelerate the computing time of 3D cellular automata simulation, we have tried a parallelization using a GP-GPU (General Purpose Graphical Processing Unit). GP-GPU is particularly powerful when used to process the same instruction set on an important number of independent data. This approach can be useful for cellular automata evolution if we consider the splitting of the cellular space in sub-cubes. However, neighborhood interactions are going to be a weighty problem for the program. To compute an iteration on a sub-cube, we need a fast access to cells' neighborhood, otherwise the new state of border cells can't be computed. In the next figure, the dark grey sub-cube will need the current state of every black cells present in other sub-cubes to compute the next state (a non-toric cube representation with a Moore cubic-shaped neighborhood with a range of one is used in this figure).
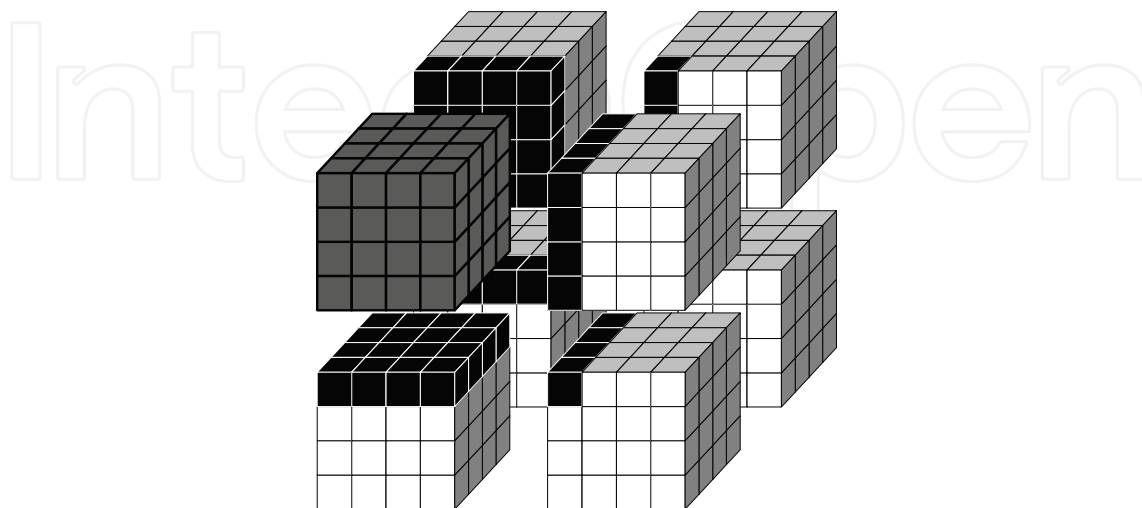

Fig. 1. Example of Moore cubic-shaped neighborhood in red with a range of one.

2.3 Problems induced by the current GP-GPU architectures

To correctly understand the consequences while programming with GP-GPU, it is important to realize that GP-GPU design is quite different from current CPU architectures. While a CPU possesses few cores, each of them allowing the execution of one thread at a time, a GP-GPU possesses a small number of streaming multiprocessors, each of them allowing the parallel execution of numerous instructions in a SIMD approach (Single Instruction Multiple Data). For instance an Nvidia Tesla 10 (see Figure 2), will have 240 vector cores split in 30 streaming multiprocessors (SM) with 8 thread processors (SP Thread Processors) each[4]. Each streaming multiprocessors can run a full warp of 32 threads with the same control flow for different data in two GPU cycle (each SP compute 2 identical operations per GPU cycle because the SM cycle is twice as fast); and since a SM can schedule up to 32 warp at a time, it leads to potentially 1024 threads running concurrently on each of the thirty streaming multiprocessor (potentially 32768 threads running concurrently on the entire card).
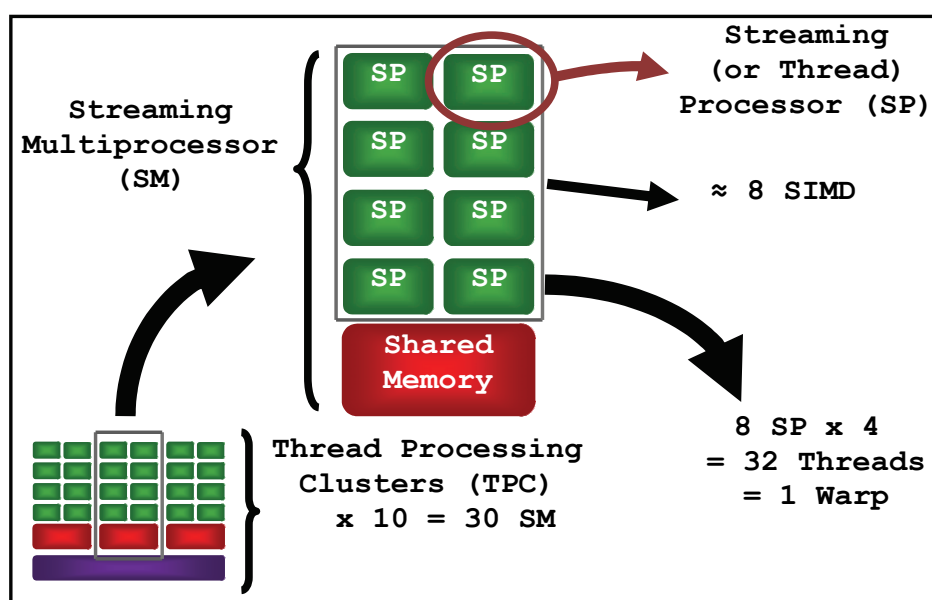

Fig. 2. Simplified architecture of an Nvidia Tesla 10.

With CUDA, all the threads needed to execute a kernel must be grouped in blocks and all these blocks must have the same, limited number of threads. All the threads of a block are executed simultaneously on the same multiprocessors and therefore can make use of its shared memory. To get the best results from GP-GPUs, we have to place data in shared memory. This memory is fast and managed by a streaming multiprocessor. The latency for accessing global memory is very high and we have to limit its access. Thus, dependant threads needing fast communications have to use the shared memory within the same streaming multiprocessor (SM).

When mapping cellular automata to a GP-GPU architecture we will not be able to place all the threads on a single SM, and consequently we will have to communicate using more

---

[4]CUDA Programming Guide v2.3 - NVIDIA_CUDA_Programming_Guide_2.3.pdf:
http://developer.download.nvidia.com/compute/cuda/2_3/toolkit/docs/NVIDIA_CUDA_Programming_Guide_2.3.pdf. Accession date: 20/01/2010.

global memory accesses, implying a significant loss in performance. The difficulty lies in successfully managing memory copy from global to shared memory. Another important consideration is the memory transfer between the host and the GP-GPU which can be very significant if the following ratio: computation duration divided by the transfer duration, is too small.

3. GP-GPU Implementation

3.1 First implementation

The first implementation we tested doesn't use shared memory. Its goal is to use the GP-GPU computation capability to see if we can already obtain a speed-up of the global computation with a basic approach.

To utilize this GP-GPU computation capacity, an elementary thread is used to compute the new state of each cell. The algorithm is then kept simple and an excerpt of the code is presented below.

```
On CPU:
  // Initialize the cell grid.
  int    * matrice, * matriceGPU;

  initGrid( &matrice, size);
  cudaMalloc( ( void ** ) &matriceGPU,
              size * sizeof( int ) );

  // Copy the cell grid from the CPU to the GPU.
  cudaMemcpy(
              matriceGPU,
              matrice,
              size * sizeof( int ),
              cudaMemcpyHostToDevice
            );

  // For each iteration needed:
  for ( i = 0; i < nbIte; ++i )
     // Run the GPU kernel to compute iteration with
     // one thread for each cell.
     nextGeneration<<< nbBlock, nbThread >>>( … );

  //Copy back the cell grid.
  cudaMemcpy(
              matrice,
              matriceGPU,
              size * sizeof( int ),
              cudaMemcpyHostToDevice
            );

On GPU kernel:
  // Compute the global thread index.
  threadIdx = threadIdx.x + blockIdx.x * blockDim.x;

  // Compute the new state of the cell corresponding
  // at this index in function of its neighborhood states.
  […]
```

Code 1. Partial code of the basic algorithm.

In this version, we first copy the cell cube on the GP-GPU. Then, we run the kernel using a thread for each cube cell. Each thread can compute the new state of the cell it handles depending on the neighborhood state. Consequently, at the end of the last thread execution, all the cells are updated for an iteration. When all the iterations have been computed, we transfer the resulting cube back onto the CPU.

This very simple algorithm has two major limiting operations. The first one is the copy of the initial cube from the CPU to the GP-GPU and the resulting cube from GPU to CPU. The second limit of this basic algorithm is almost masked in the previous code: it deals with memory access to the cells and their neighborhood.

For the first point, copying time costs will be very difficult to improve. A solution could be to take advantage of asynchronous memory transfer by dividing computation in multiple parts for example. But this solution would significantly affect the algorithm simplicity and reusability, and moreover it will not allow a significant global speed-up in most cases, particularly when we have a large number of iterations to compute. Furthermore, the greater the number of iterations, the less the copy time costs will be significant in comparison with the computing cost for all iterations.

The second point is related to global memory access. When we perform the Game of Life in three dimensions, it is necessary to know the neighborhood of each cell. This leads to a maximum of 27 cell states if we consider the current cell state. In this first version, all memory access is done in global memory with a very high access latency (between 400 and 600 clock cycles on a Tesla T10 for instance).

If we consider the handling of processing 2 cells at a time, we can notice that 2 neighboring cells share 18 cells in their common neighborhood. Thus we can obtain all their neighborhood with only 36 global memory access instead of 54. And results are even better if we increase the block size to small cubes: a 3^3 cube (27 cells) will need only 125 global memory access (the block including neighboring cells is a 5^3 cube) instead of accessing 729 cells if we consider the handling on a single cell basis (27 x 27). The following implementation tries to take advantage of this possibility.

3.2 Second implementation

To limit at minimum the global memory access, the code hereafter makes use of shared memory between threads of a same CUDA software block. Since the size of a software block is limited, it is impossible to share all the cell information between all the threads needed to handle a large cube. Each software block works on a part of the cube and each thread inside this block works on the generation of the new state of a single cell. Before computing an iteration, data required for a CUDA software block must be loaded from global memory into shared memory. This is done in every software block. When the uploading is done, a synchronization of all the threads of a same block is necessary to ensure that all the data is correctly uploaded in shared memory. After this synchronization, the computing of the new state can be done in function of neighborhood states read in shared memory instead of global memory. The new state obtained for the local cube is written in global memory otherwise it would be lost after the end of the block computation.

In this solution, the aim is the reduction of the global memory access but nothing has been done for the memory transfer between CPU and GP-GPU.

With this approach we still have global memory access to load the shared memory and to update the global memory after an iteration. Even if there is a reduction in terms of global memory access, this approach results in an increase in the total number of memory access

(global and shared) but most of them are shared memory with a final increase in performance.

```
On CPU:
    […]
    int  neighborSize = […];
    int  sharedSize   = ( size / nbThread + neighborSize ) *
                        sizeof ( int );
    […]
    // Run the GPU kernel to compute iteration with
    // one thread for each cell.
    nextGeneration<<< nbBlock, nbThread, sharedSize >>>(…);
    […]

On GPU kernel:
    // Upload from global memory to shared memory
    int   sharedMatrice[];
    sharedMatrice[ threadIdx.x ]   =   matriceGPU[ … ];

    __syncthreads();

    // Compute the new state of the cell corresponding at
    // this index in function of its neighborhood states
    // present in shared memory.
    […]
```

Code 2. Pseudo-code of the shared algorithm.

To maximize the number of cells shared in a same block, the largest block possible has been used. Each block contains 512 threads, allowing the processing of an 8^3 size cube. To limit each thread access to the global memory, the state of each cube cell is loaded in shared memory. This choice implies that all threads handling the cells at the border of the 8^3 cube are not used to compute the next state. To obtain the next iteration state, only the threads processing the inner cube of 6^3 cells are considered.

4. Results

The two algorithms previously presented have been tested on personal computers, the main one using two Intel Nehalem CPUs (at 2.53 GHz) with an NVidia Tesla C1060 board. To compare the benefits of hybrid computing, the CPU only algorithm computes for each cell the new state with a single core (a Core 2 Xeon and a Nehalem Xeon have been considered).

4.1 Memory transfer costs

Figure 3 confirms what has been said previously when explaining the basic algorithm: transfer costs between CPU and GPU are not significant anymore when we compute many iterations. In fact, for a constant cube size, the global simulation duration should be linear with the number of iterations; to this global computing time we naturally add two memory transfer costs for copying the initial cube from the CPU memory to the GPU global memory and for copying back the final result. In Figure 3, with a small 10^3 cube, the impact of this transfer time is important if we have a small number of iterations and becomes insignificant after 400 iterations.
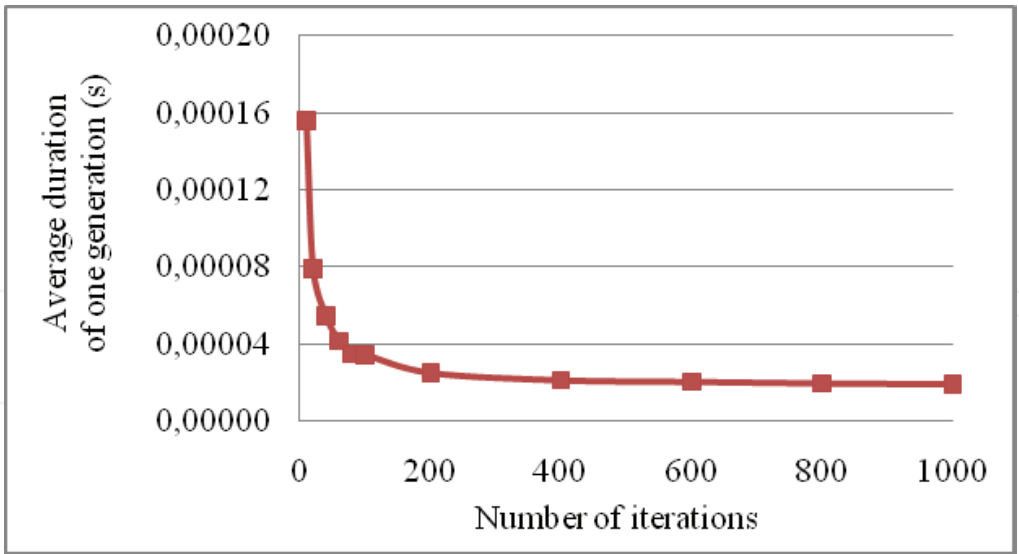
Fig. 3. Average duration (in seconds) of one generation for a cube with 10^3 cells.

The next figure shows the obvious importance of the cube size. For big cubes the transfer time measured became rapidly insignificant even with a small number of iterations. Now that the memory transfer problem has been analyzed, it is possible to start comparing CPU and GPU results.
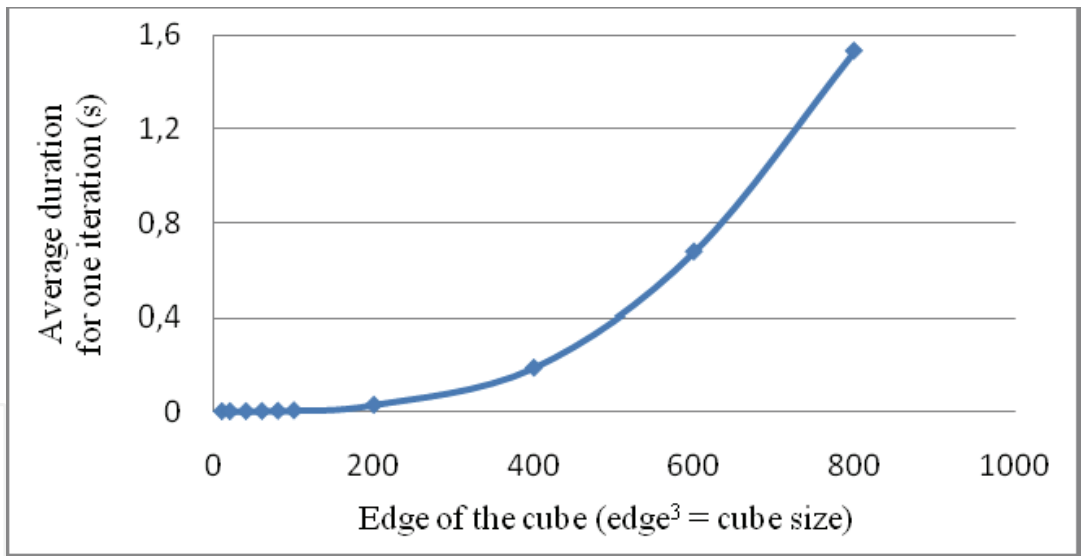

Fig. 4. Average time for one iteration (100 replicates) for the following 3 tasks: allocation of 2 cubes (cube size), copy of a cube from CPU to GPU and copy of the final resulting cube (from GPU to CPU)

4.2 Basic GPU implementation results
The first results deal with the basic implementation which uses only global memory.
Figure 5 shows a comparison between the pure CPU algorithm (Core2 Q9300 at 2.5GHz) and its hybrid implementation with a GP-GPU for a 20^3 cube (Tesla T10 – C1060). We see that the best results are obtained when a significant number of iterations. The performance improvement stops after reaching a threshold.

In Figure 5, we see that the sequential implementation using the CPU is always slower for a 20^3 cube when compared to the GP-GPU implementation.
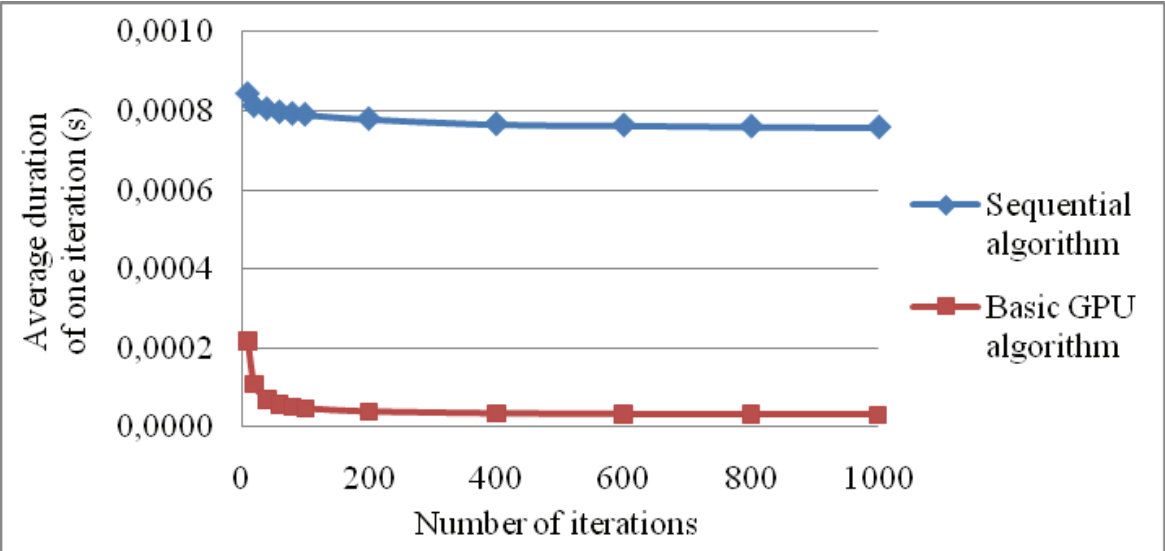
Fig. 5. Average duration of an iteration depending on the number of iterations for a 20^3 cube.

In Figure 6, we notice that when using a GP-GPU, the computing time of an iteration grows slowly with the cube size, whereas the computing time for an iteration according to the cube size shows a polynomial growth if we use the sequential CPU algorithm.
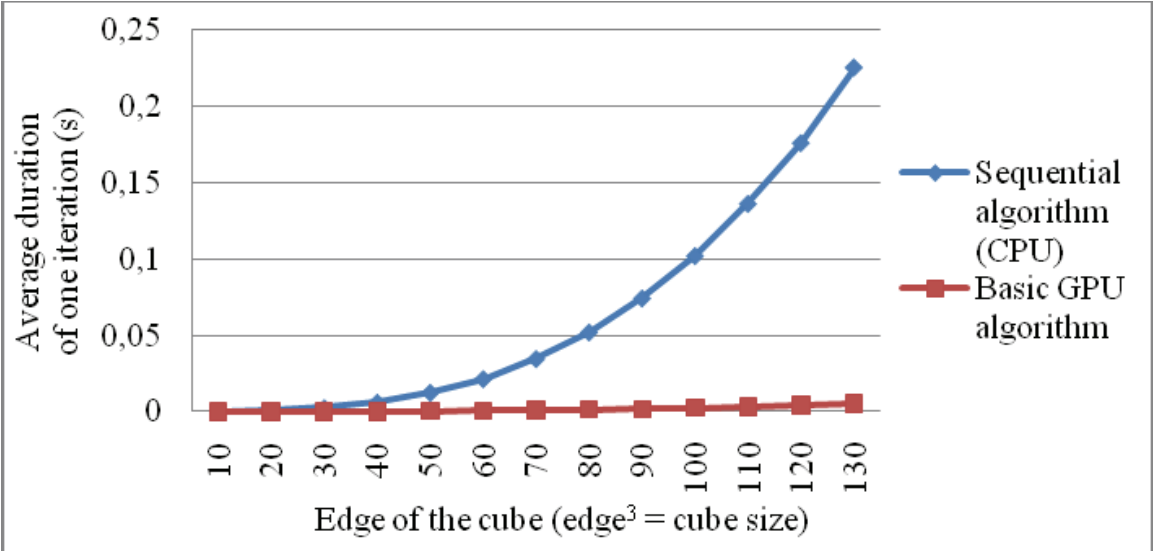
Fig. 6. Average duration of one generation for 1000 iterations depending on the cube size.

Figure 7 presents the speed-up for different cube size depending on the number of iterations. For each cube size we have an upper limit; we also note that the speed-up first increases with cube size and stays approximately at the same level when the cube size reaches 80^3 and over.

This last point can be seen more clearly on Figure 8 which represents the speed-up for 1000 iterations depending on the cube size.
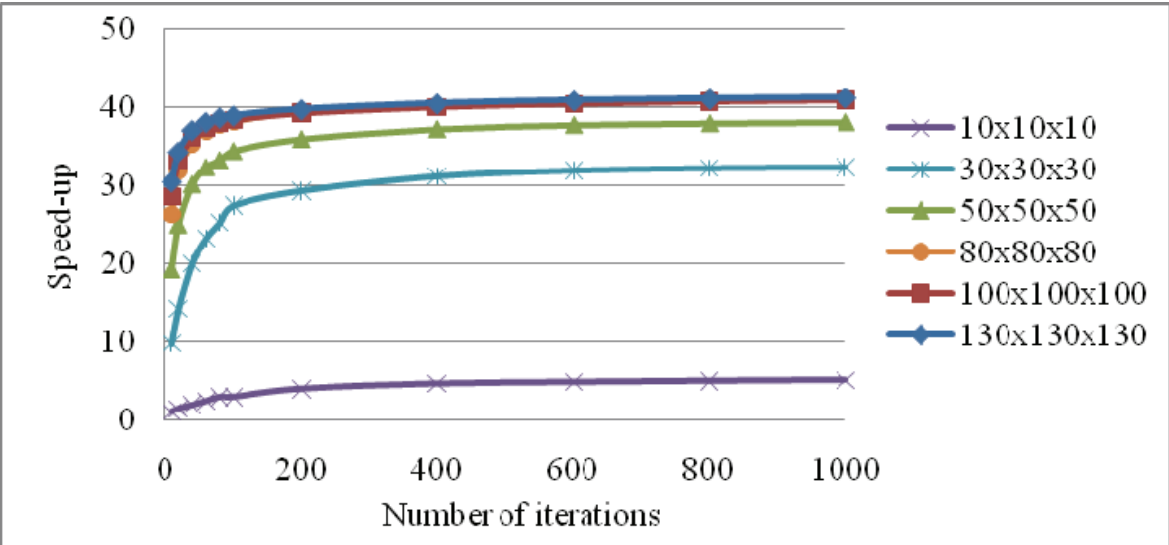
Fig. 7. Basic GPU implementation speed-up compared to a sequential CPU implementation for different cube sizes depending on the number of iteration.
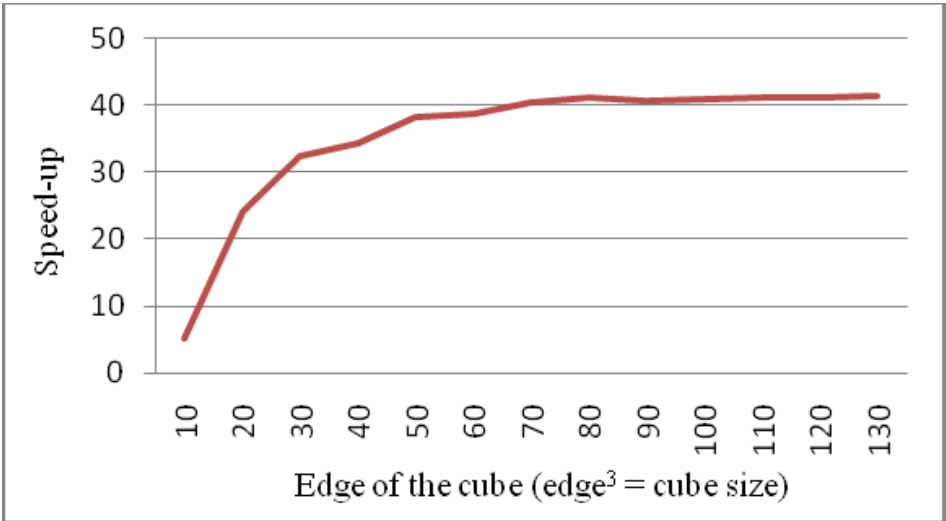
Fig. 8. Speed-up evolution for 1000 iterations depending on the cube sizes (50 measures for each point).

Now that the speed-up trend has been identified, we can discuss speed-up results. If we ignore cube sizes smaller than 80^3, the average speed-up for 1000 iterations is 41x when compared to a single Q9300.

4.3 GPU implementation with shared memory results

The next implementation takes advantage of shared memory between threads of the same streaming multiprocessor to minimize global memory access.

Figure 9 shows the performance increase compared to the first implementation (using a 100^3 cube). It confirms that the two trends are really similar and also that the shared memory algorithm is better than the previous one.

Figure 10 compares the speed-up for 1000 iterations with both implementations. Speed-ups are very similar for small cubes but as expected the shared memory algorithm has better

speed-up for larger cubes (reaching 60x compared to 40x). With this algorithm the upper limit is reached later than with the previous algorithm (see Figure 11).

These results can be generalized for larger cubes. Figure 12 shows speed-up evolution for 1000 iterations comparing the CPU algorithm for two different processors: the previous Core2 Q9300 but also a Nehalem 5500. The first point we can notice is that the speed-up did not slow-down with the increase in cube size (the 800^3 cube size limit will be discussed in conclusion). The second point is that speed-up obtained with the GPU is far better when we compare it with a Q9300 Xeon processor (average speed-up of 62.8x) than in comparison with a Nehalem 5500 processor (average speed-up of 20.5x). This difference shows an impressive advance in CPU performances, at almost the same frequency, particularly if we consider that the Xeon Q9300 was still considered as a very good Intel processor at the beginning of year 2009.
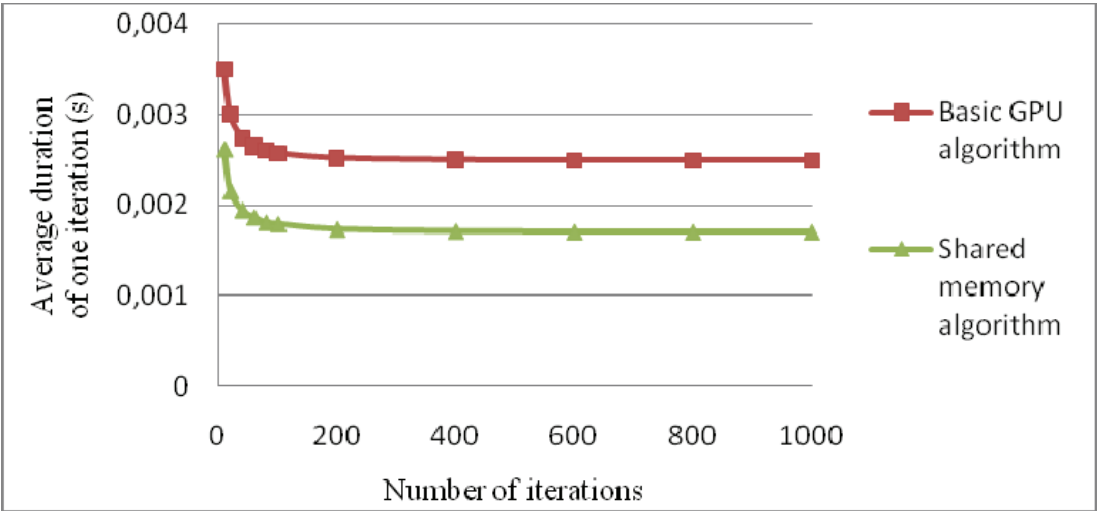
Fig. 9. Average duration of an iteration depending on the number of iterations for a 100^3 cube (50 replicates).
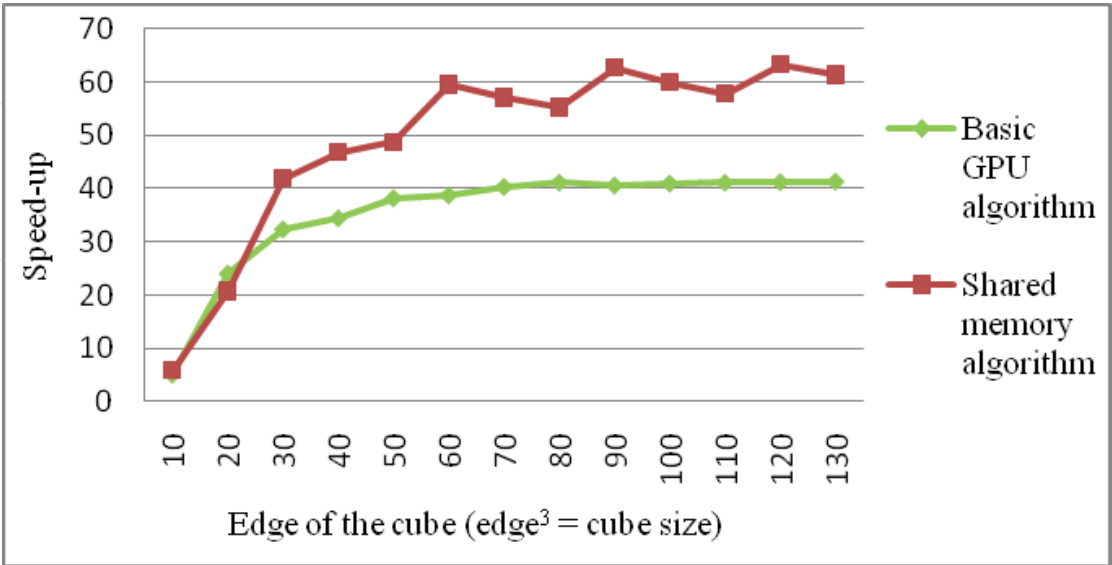
Fig. 10. Comparison of speed-up evolution for 1000 iterations depending on the cube sizes (50 measures for each point).
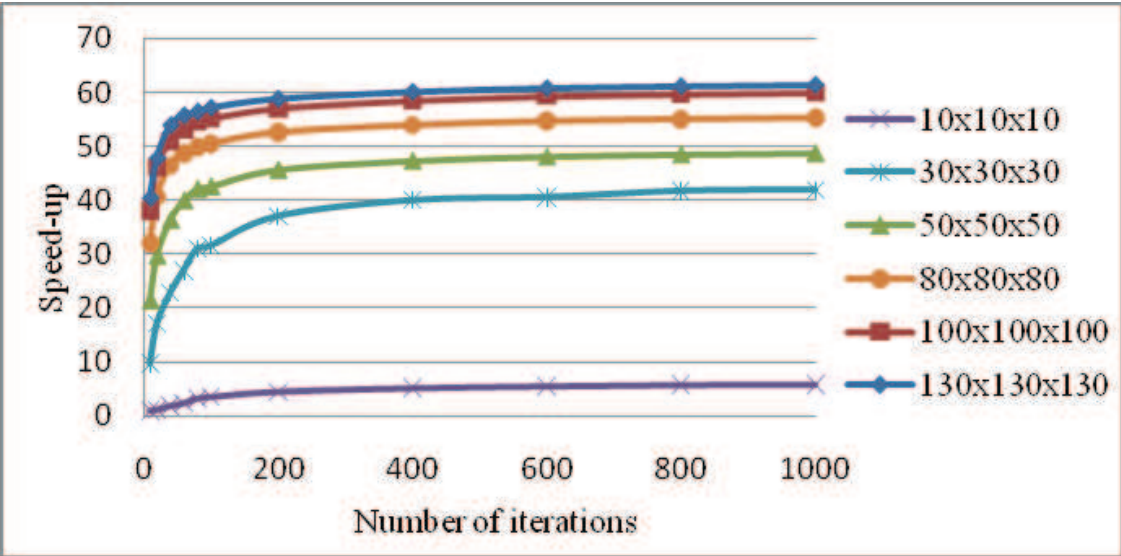
Fig. 11. Shared memory GPU implementation speed-up for different cube sizes depending on the number of iterations.
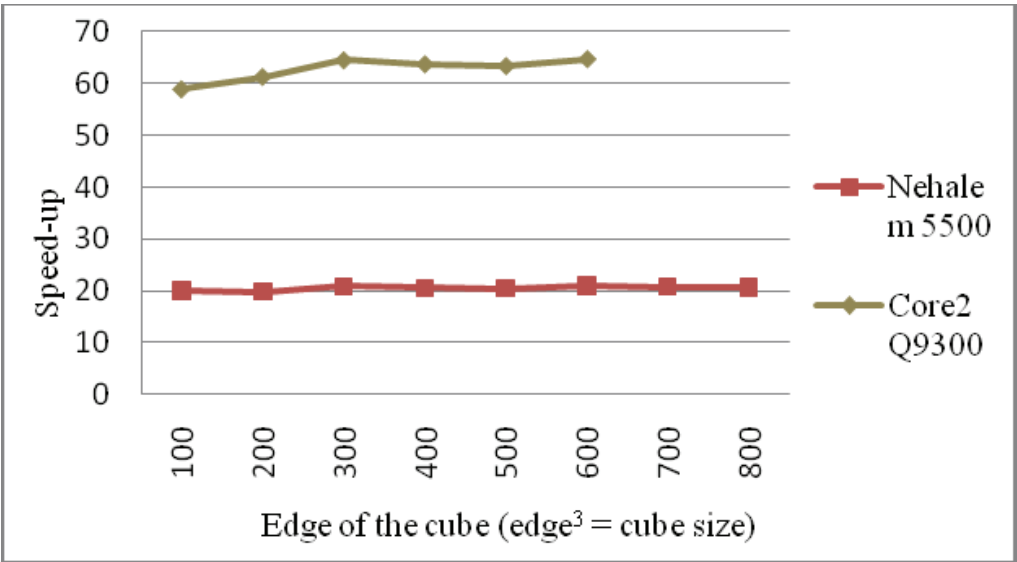

Fig. 12. Speed-up evolution on large cube for 1000 iterations depending on the cube sizes.

5. Future research

In the case of 3D cellular automata, whatever the implementation, basic or more sophisticated, the use of hybrid computing shows interesting speed up. But improving this solution is still possible. Here are some possibilities which it could be interesting to explore. First, it is possible to have all the threads of a block to work on computing the next status of a cell. Currently, all the threads corresponding to the block border are only used to transfer neighborhood status from global to shared memory. To make them work, it is necessary to make them load multiple cell status from global to shared memory. This solution, which has not been tested yet, does not give any guarantee of speed-up and it probably needs to be tested on different automata configurations.

It would also be interesting to see if using another block size can improve the performance. At the moment a block of 512 threads is used to represent 8^3 cubes. But it could be interesting to test blocks of 256 threads to represent 6^3 cubes.

The second implementation was proposed to improve the memory access of the first version (which used only global memory). The fastest implementation uses mainly shared memory but it still needs to transfer data from the global memory to the shared memory and we think that this transfer could be done faster. In my implementation, each block of threads corresponds to a block of the cube cells. In the example of Figure 13, the 125 threads of a block access 125 cube cells in global memory (the part framed in orange). When each thread accesses the state of a cell, it's clear that the global memory index will not be continuously accessed and with the CUDA architecture, we know that memory access can be done more efficiently by reading continuous data.

A solution could be to reorganize the cube data to allow continuous global memory access for all the threads in a same block. This better memory access, implies that more computation would have to be done on each thread to compute the correct index in the cube but it would be interesting to see if it improves the performance. However this kind of dedicated implementation becomes quite obfuscated and difficult to maintain.
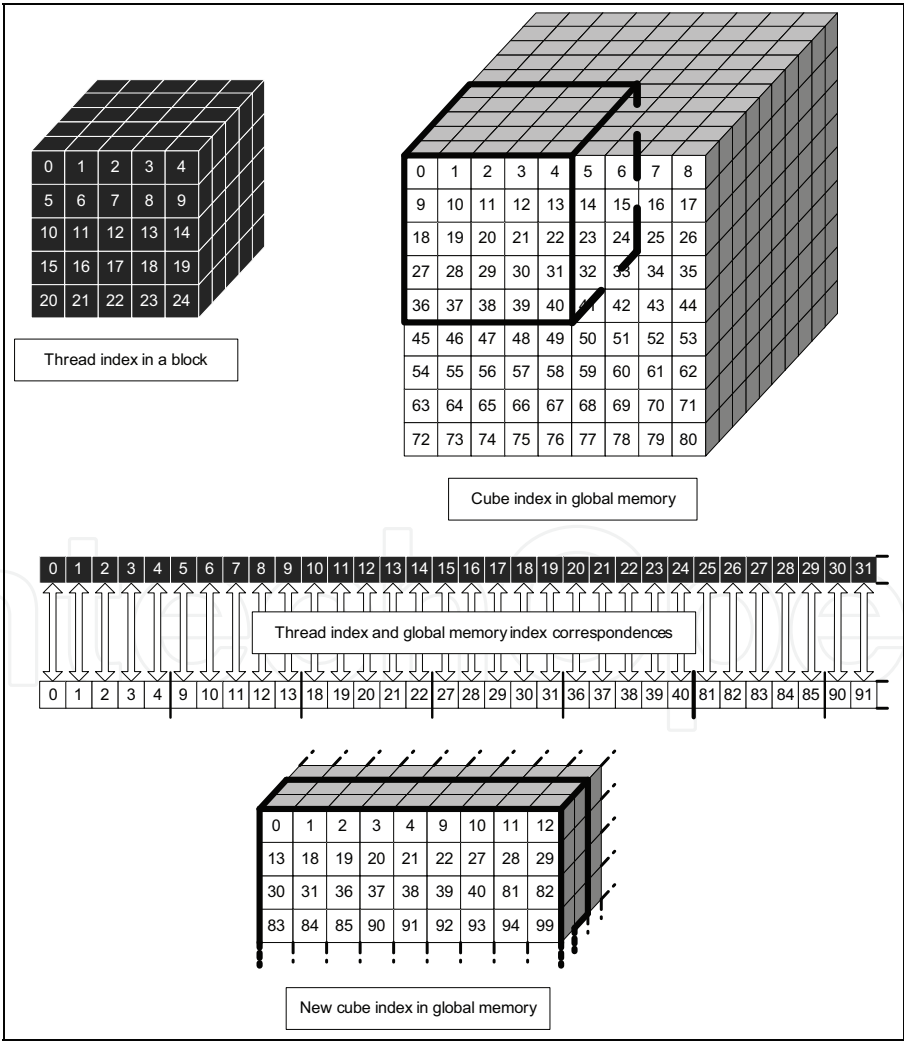

Fig. 13. Improvement of data storage to increase global memory access.

Another solution to improve the results could be to create specific generation algorithm for each cellular automata algorithm by taking advantage of their particularities. For example, in the case of excitable medium, there are generally some important cube parts where no activity is present. It could be possible to locate these parts and therefore avoid doing any computation because these parts won't evolve. But there is an important problem with this solution. It's logically the development cost. If many cellular automata algorithms are used, developing for each one a specific generation algorithm which takes advantage of its particularities could be an enormous task. And even if the number of cellular automata algorithms is limited, specific processing must be fast enough avoid imparing the global computation time. In the previous example, if searching for a large empty cube part takes a longer time than computing a new generation on this part, performance increase can't be guaranteed.

6. Conclusion

We have implemented two cellular automata algorithms on a Tesla C1060 board. The basic principle is the same for both algorithms, a GPU thread is used to compute the new state of a cell. But to compute a new cell state, a thread has to read the states of the 27 neighbouring cells (if we consider a classic Moore neighborhood extended to a 3D automaton). The total number of memory accesses can incur a significant access time when relying on the GP-GPU global memory. A second implementation has been created to minimize the number of global memory accesses by using the fast shared memory present on a GP-GPU streaming multiprocessor. The states of all the cells needed by all the threads running on a same multiprocessor are transferred only once from global memory to shared memory. This is sufficient to compute a significant set of cells in parallel with efficient shared memory accesses.

We have tested both approaches and the results show that we have an interesting speed-up for any cube size when enough iterations are computed. On small cubes, when a little number of iterations is considered, the transfer costs between the host and the GPU make the whole process significantly slower. When large cubes are considered (100^3 and above), the average speed-up for the first implementation compared to a sequential algorithm running on a Xeon Core 2 CPU at 2.5 GHz is around 41x (13x when compared to a Nehalem Xeon at 2.53 GHz). If shared memory is used (second implementation), the speed-up grows to 65x when compared to a Xeon Core 2 CPU (and 20x a Nehalem Xeon).

The results observed show that even with the best CPUs such as Intel Nehalem, we obtain significant speed-ups. However, we have to consider the following limiting points. First, the CPU algorithm used for comparison with the GP-GPU algorithm is a sequential one. Now, most of the computers have multiple cores and taking the best of all of them will significantly increase the regular host performance. This point is not so important in integrative biology because cellular automata are only a part of the whole computation and the computation power of host cores will be used for other computations which have lower speed-ups while computing on the GP-GPU.

The second point is the cube size. We have to take into account the current memory limit for Tesla GP-GPUs (4 GB). To compute a new iteration on a GP-GPU, we use two cubes, one handling the cells current statuses and the other to write the new statuses. A Tesla C1060 has 4GB of global memory, hence the size of the cube can't be too important. Two cubes with 900^3 cells can't be processed on this kind of GPU because it would need more than

5.8GB if 4 bytes are used to store a cell status (1000^3 can be considered if we store cell states on a single byte).

The next GP-GPU architecture (Fermi T20) is very promising, it improves significantly memory access performances with a configurable cache and ECC memory, and it will also be able to address up to 1TB of memory. With this kind of architecture, a cube with 5000^3 cells could be handled on a single GP-GPU board, enabling a fine multi-scale modeling for simulated organs. Grids and clusters of GP-GPUs will soon be considered to simulate interactions between organs (Fan et al., 2004).

7. References

Lu, S.; Michailova, A. P.; Sauverman, J. J.; Cheng, Y., Yu, Z.; Kaiser, T. H.; Li, W. W.; Bank, R. I.; Holst, M. J.; McCammon, J. A.; Hayashi, T.; Hoshijima, M.; Arzberger, P. & McCulloch, A. D. (2009). *Multiscale Modeling in Rodent Ventricular Myocytes*. IEEE Engineering in Medicine and Biology Magazine. March-April 2009; Volume 28 - Number 2; pp. 46-57.

Pivkin, I. V.; Richardson, P. D. & Karniadakis, G. E. (2009). *Effect of Red Blood Cells on Platelet Aggregation*. IEEE Engineering in Medicine and Biology Magazine. March-April 2009; Volume 28 - Number 2; pp. 32-37.

Sander, E. A.; Strylianopoulos, T.; Tranquillo R. T. & Barocas V. H. (2009). *Image-Bases Biomechanics of Collagen-Based Tissus Equivalents*. IEEE Engineering in Medicine and Biology Magazine. May-June 2009; Volume 28 - Number 3; pp. 10-18.

Von Neumann, J. & Burks, A. W. (1966). *The Theory of Self-reproducing Automata*. Univ. of Illinois Press, Urbana, Illinois.

Alarcón, T.; Byrne; H. M. & Maini, P. K. (2004). *Towards whole-organ modelling of tumour growth*. Progress in Biophysics and Molecular Biology 2004, 85; pp. 451-472.

Siregar, P.; Sinteff, J. P.; Julen, N. & Le Beux, P. (1998). *An Interactive 3D Anisotropic Cellular Automata Model of the Heart*. Computers and Biomedical Research, Volume 31, Issue 5, October 1998; pp. 323-347.

Martin Gardner, (October 1970), "Mathematical Games: The fantastic combinations of John Conway's new solitaire game "Life"", Scientific American 223: pp. 120–123.

Wolfram Stephen, A new kind of Science, Wolfram Media, Inc, ISBN 1-57955-008-8, 2002.

A. K. Dewdney. The cellular automata programs that create wireworld, rugworld and other diversions, Computer Recreations, *Scientific American,* Jan: pp. 146-149. 1990.

Ermentrout, G. B. & Edelstein-Keshet, L. (1993). *Cellular automata approaches to biological modelling*. Journal of Theoretival Biology 160, pp. 97–133.

Luebke, D. L.; Harris M.; Krüger J.; Purcell T.; Govindaraju N.; Buck I.; Wooley C. & Lefohn A. *GPGPU : General purpose computation on graphics hardware*. (2004). In SIGGRAPH'04, Proceedings of the SIGRAPH 2004 conference, course notes, New York, NY, USA, ACM Press (2004); p. 33.

Trendall, C. & Stewart A. J. (2000). *General calculations using graphics hardware with application to interactive caustics*. In Rendering Techniques '00 (Proc. Eurographics Rendering Workshop), pp. 287-298. Springer, June 2000.

Woundeberg M. *Using FPGAs to Speed Up Cellular Automata Computations*. (2006). Master's thesis, University of Amsterdam.

Fan, Z.; Qiu, F.; Kaufman, A. & Yoakum-Stover, S. *GPU cluster for high performance computing*. (2004). In SC '04: Proceedings of the 2004 ACM/IEEE conference on Supercomputing, page 47, Washington, DC, USA, 2004. IEEE Computer Society.
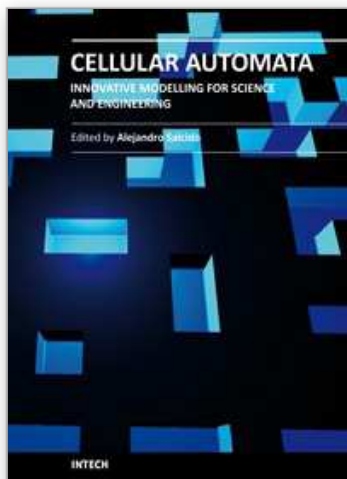
Modelling and simulation are disciplines of major importance for science and engineering. There is no science without models, and simulation has nowadays become a very useful tool, sometimes unavoidable, for development of both science and engineering. The main attractive feature of cellular automata is that, in spite of their conceptual simplicity which allows an easiness of implementation for computer simulation, as a detailed and complete mathematical analysis in principle, they are able to exhibit a wide variety of amazingly complex behaviour. This feature of cellular automata has attracted the researchers' attention from a wide variety of divergent fields of the exact disciplines of science and engineering, but also of the social sciences, and sometimes beyond. The collective complex behaviour of numerous systems, which emerge from the interaction of a multitude of simple individuals, is being conveniently modelled and simulated with cellular automata for very different purposes. In this book, a number of innovative applications of cellular automata models in the fields of Quantum Computing, Materials Science, Cryptography and Coding, and Robotics and Image Processing are presented.

INTECH
open science | open minds