

Cellular Automata as a Mesoscopic Approach to Model and Simulate Complex Systems

P.M.A. Sloot and A.G. Hoekstra

Section Computational Science, Faculty of Science, University of Amsterdam,
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
{sloot, alfons}@science.uva.nl

Abstract. We discuss the cellular automata approach and its extensions, the lattice Boltzmann and multiparticle methods. The potential of these techniques is demonstrated in the case of modeling complex systems. In particular, we consider simple applications taken from various scientific domains. We discuss our distributed particle simulation of flow, based on a parallel lattice Boltzmann Method. Efficient parallel execution is possible, provided that (dynamic) load balancing techniques are applied. Next, we present a number of case studies of flow in complex geometry, i.e. flow in porous media and in static mixer reactors.

1. The Cellular Automata Approach

A natural way to describe a physical, chemical or biological system is to propose a model of what we think is happening. During this process we try to keep only the ingredients we believe to be essential whilst still capturing the behavior we are interested in. Using an appropriate mathematical machinery, such a model can be expressed in terms a set of equations whose solution gives the desired answers on the system. The description in terms of equations is very powerful and corresponds to a rather high level of abstraction. For a long time, this methodology has been the only tractable way for scientists to address a problem.

Another approach, which has been made possible by the advent of fast computers, is to stay at the level of the model and its basic components. The idea is that all the information is already contained in the model and that a computer simulation will be able to answer any possible question on the system by just running the model for some time. Thus, there is no need to use a complicated mathematical tool to obtain a high level of description. We just need to express the model in a way which is suitable to an effective computer implementation. Cellular automata constitute a paradigm in which simple models of complex phenomena can be easily formulated. In particular, cellular automata models illustrate the fact that a complex behavior emerges out of many simply interacting components through a collective effect.

The degree of reality of the model depends on the level of description we expect. When we are interested in the global or macroscopic properties of a system the microscopic details of a system are often irrelevant. On the other hand, symmetries and conservation laws are usually the essential ingredients of such mesoscopic

description. It is therefore a clear advantage to invent a much simpler microscopic reality, which is more appropriate to our computational means of investigation.

A cellular automata model can be seen as a fictitious universe which has its own microscopic reality but, nevertheless, has the same macroscopic behavior as the real system we are interested in. The examples in the next section will illustrate this statement.

Cellular automata (CA) are an idealization of a physical system in which space and time are discrete. In addition, the physical quantities (or state of the automaton) take only a finite set of values. Since the time it has been invented by von Neumann in the late 1940s, the cellular automata approach has been applied to a large range of scientific problems [1, 2].

The original motivation of von Neumann was to extract the abstract mechanisms leading to self-reproduction of biological organisms. [2]

Following the suggestions of S. Ulam, von Neumann addressed this question in the framework of a fully discrete universe made up of cells. Each cell is characterized by an internal state, which typically consists of a finite number of information bits. Von Neumann suggested that this system of cells evolves, in discrete time steps. The rule, determining the evolution of this system is the same for all cells and is a function of the states of the neighboring cells. Similarly to what happens in any biological system, the activity of the cells takes place simultaneously. However, the same clock drives the evolution of each cell and the updating of the internal state of each cell occurs synchronously. Such fully discrete dynamical systems (cellular space) as invented by von Neumann are now referred to as a *cellular automata* (CA).

After the work of von Neumann, other authors have followed the same line of research and nowadays the problem is still of interest [3] and has led to interesting developments for new computer architectures and algorithms. [4, 5]

A very important feature of CAs is that they provide simple models of complex systems. They exemplify the fact that a collective behavior can emerge out of the sum of many, simply interacting, components. Even if the basic and local interactions are perfectly known, it is possible that the global behavior obeys new laws that are not obviously extrapolated from the individual properties, as if the whole is more than the sum of all the parts. This properties makes cellular automata a very interesting approach to model physical systems and in particular to simulate complex and nonequilibrium phenomena. [6]

The studies undertaken by S. Wolfram in the 1980s [7] clearly establishes that a CA (the famous Wolfram's rules) may exhibits many of the behaviors encountered in continuous dynamical systems, yet in a much simpler mathematical framework. A further step is to recognize that CAs are not only behaving similarly to some dynamical processes, they can also represent an actual model of a given physical system, leading to macroscopic predictions that could be checked experimentally. This fact follows from statistical mechanics which tells us that the macroscopic behavior of many systems is quite disconnected from its microscopic reality and that only symmetries and conservation laws survives to the change of observation level: it is well known that the flows of a fluid, a gas or even a granular media are very similar at a macroscopic scale, in spite of their different microscopic nature.

Although John von Neumann introduced the cellular automata theory several decades ago, only in recent years it became significant as a method for modeling and simulation of complex systems. This occurred due to the implementation of cellular automata on massively parallel computers. Based on the inherent parallelism of

cellular automata, these new architectures made possible the design and development of high-performance software environments. These environments exploit the inherent parallelism of the CA model for efficient simulation of complex systems modeled by a large number of simple elements with local interactions. By means of these environments, cellular automata have been used recently to solve complex problems in many fields of science, engineering, computer science, and economy. In particular parallel cellular automata models are successfully used in fluid dynamics, molecular dynamics, biology, genetics, chemistry, road traffic flow, cryptography, image processing, environmental modeling, and finance. [3]

Interesting examples are the Lattice Gas Automata for fluid flow, and the derived Lattice Boltzmann models. The sequel of the paper will introduce these mesoscopic models and demonstrate the potential of Cellular Automata as a Mesoscopic Approach to Model and Simulate Complex Systems.

2. Particle Simulation of Flow

Some 12 years ago theoretical physicists showed that an highly idealized model of a gas, consisting of particles that have a very limited set of velocities and that are confined to a lattice behaves, on average, as an incompressible fluid [8]. This Lattice Gas Automaton (LGA) is a particle model for fluid flow and is a true CA. In this section we shortly introduce LGA and show how it relates to CA. Detailed theory behind LGA is introduced in two recent books. [9, 10]

Consider a hexagonal lattice, as in Fig. 1. Particles live on the links of the lattice, and they can move from one node to another. The dynamics of the particles is such that they all move from one node to another. Next, if particles meet on a node, they collide and change direction (see Fig. 1). The collisions are such that they obey the physical constraints of conservation of mass, momentum, and energy.




Fig. 1. Lattice and particle update mechanism for a LGA. A dot denotes a particle and the arrow its moving direction. From left to right an initial condition, streaming, and collision of particles are shown.

We can formally define a CA rule for such an LGA as follows. Suppose that the state of a cell is determined by b_m surrounding cells. Usually, only the nearest and next-nearest neighbors are considered. For example, on a square lattice with only nearest neighbor interactions $b_m = 4$, if next-nearest neighbors are also included $b_m = 8$, and on a hexagonal lattice with nearest neighbor interactions $b_m = 6$. Furthermore,

suppose that the state of the cell is a vector \mathbf{n} of $b = b_m$ bits. Each element of the state vector is associated with a direction on the CA lattice. For example, in the case of a square grid with only nearest neighbor interactions we may associate the first element of the state vector with the *north* direction, the second with *east*, the third with *south*, and the fourth with *west*. With these definitions we construct the following CA rule (called the LGA rule), which consists of two sequential steps:

1. Each bit in the state vector is moved in its associated direction (so in the example, the bit in element 1 is moved to the neighboring cell in the north) and placed in the state vector of the associated neighboring cell, in the same position (so, the bit in element 1 is moved to element 1 of the state vector in the cell in the north direction). In this step each cell is in fact moving bits from its state vector in all directions, and at the same time is receiving bits from all directions, which are stored into the state vector.
2. Following some deterministic or stochastic procedure, the bits in the state vector are reshuffled. For instance, the state vector (1,0,1,0) is changed to (0,1,0,1).

Maybe very surprisingly, if we assign physical quantities to this CA, enforce physical conservation laws on the bit reshuffling rule of step 2, and use methods from theoretical physics to study the dynamics, we are in fact able to analyze the CA in terms of its average behavior. The average state vector of a cell and the average flow of bits between cells can be calculated. Even better, it turns out, again within the correct physical picture that this CA behaves like a real fluid (such as water) and therefore can be used as a model for hydrodynamics. Furthermore, as the LGA rule is intrinsically local (only nearest and next nearest neighbor interactions) we constructed an inherently parallel particle model for fluid flow.

Associate the bits in the state vector with *particles*; a one-bit codes for the presence of a particle, and a zero bit codes for the absence of a particle. Assume that all particles are equal and have a mass of 1. Step 1 in the LGA-CA is now interpreted as streaming of particles from one cell to another. If we also introduce a length scale, i.e. a *distance* between the cells (usually the distance between nearest neighbors cells is taken as 1) and a time scale, i.e. a *duration* of the streaming (i.e. step 1 in the LGA-CA rule, usually a time step of 1 is assumed), then we are able to define a *velocity* \mathbf{c}_i for each particle in direction i (i.e. the direction associated with the i -th element of the state vector \mathbf{n}). Step 1 of the LGA-CA is the streaming of particles with velocity \mathbf{c}_i from one cell to a neighboring cell. Now we may imagine, as the particles meet in a cell that they collide. In this collision the velocity of the particles (i.e. both absolute speed and direction) can be changed. The reshuffling of bits in step 2 of the LGA-CA rule can be interpreted as a collision of particles. In a real physical collision, mass, momentum, and energy are conserved. Therefore, if we formulate the reshuffling such that these three conservation laws are obeyed, we have constructed a true Lattice Gas Automaton, i.e. a gas of particles that can have a small set of discrete velocities \mathbf{c}_r , moving in lock-step over the links of a lattice (space is discretized) and that all collide with other particles arriving at a lattice point at the same time. In the collisions, particles may be sent in other directions, in such a way that the total mass and momentum in a lattice point is conserved.

We can now define for each cell of the LGA-CA a density ρ and momentum $\rho\mathbf{u}$, with \mathbf{u} the velocity of the gas:

$$\rho = \sum_{i=1}^b N_i, \quad \rho \mathbf{u} = \sum_{i=1}^b \mathbf{c}_i N_i, \quad (1)$$

where $N_i = \langle \mathbf{n}_i \rangle$, i.e. a statistical average of the Boolean variables; N_i should be interpreted as a particle density.




Fig. 2. LGA simulation of flow around a cylinder. The arrows are the flow velocities, the length is proportional to the absolute velocity. The simulations were done with FHP-III, on a 32×64 lattice, the cylinder has a diameter of 8 lattice spacings, only a 32×32 portion of the lattice is shown; periodic boundary conditions in all directions are assumed. The left figure shows the result of a single iteration of the LGA, the right figure shows the velocities after averaging over 1000 LGA iterations.

If we let the LGA evaluate and calculate the density and momentum as defined in eqn. (1), these quantities behave just like a real fluid. In Fig. 2 we show an example of an LGA simulation of flow around a cylinder. In left figure we show the results of a single iteration of the LGA, so in fact we have assumed that $N_i = \mathbf{n}_i$. Clearly, the resulting flow field is very noisy. In order to arrive at smooth flow lines one should calculate $N_i = \langle \mathbf{n}_i \rangle$. Because the flow is static, we calculate N_i by averaging the Boolean variables \mathbf{n}_i over a large number of LGA iterations. The resulting flow velocities are shown in the left panel of Fig. 2.

Immediately after the discovery of LGA as a model for hydrodynamics, it was criticized on three points; noisy dynamics, lack of Galilean invariance, and exponential complexity of the collision operator. The noisy dynamics is clearly illustrated in Fig. 2. The lack of Galilean invariance is a somewhat technical matter which results in small differences between the equation for conservation of momentum for LGA and real Navier-Stokes equations, for details see e.g. [10]. Finally, adding more velocities in an LGA leads to increasingly more complex collision operators, exponentially in the number of particles. Therefore, another model, the Lattice Boltzmann Method (LBM), was introduced. This method is reviewed in detail in [11].

The basic idea in LBM is that one should not model the individual particles \mathbf{n}_i , but immediately the particle densities N_i . This means that particle densities are streamed

from cell to cell, and particle densities collide. This immediately solves the problem of noisy dynamics. However, in a strict sense we no longer have a CA with a Boolean state vector. However, we can view LBM as a generalized CA. It is easy to make LBM Galilean invariant, thus solving the second problem of LGA. Finally, a very simple collision operator is introduced. This so-called BGK collision operator models the collisions as a single-time relaxation towards equilibrium. This L-BGK method is also developed for many other lattices, e.g. in two or three-dimensional cubic lattices with nearest and next nearest neighbor interactions. The LBM and especially the L-BGK has found widespread use in simulations of fluid flow.

3. Parallel Lattice Gas and Lattice Boltzmann Simulations

The local nature of the LGA and LBM interactions allows a very straightforward realization of parallelism. A geometric decomposition of the lattice with only local message passing between the boundaries of the different domains is sufficient to realize an efficient parallel simulation. For instance, we have developed a generic 2-dimensional LGA implementation that is suitable for any multi-species (thermal) LGA [12]. Parallelism was introduced by means of a 1-dimensional, i.e. strip-wise, decomposition of the lattice. As long as the grid dimension compared to the number of processors is large enough, this approach results in a very efficient parallel execution.




Fig. 3. The processor load, on 16 processors, for slice and ORB decomposition of the elutriator chamber benchmark (see [13]). On the x-axis are the processor ID's and on the y-axis the execution time of each processor.

This LGA system is mainly used for simulations in simple rectangular domains without internal obstacles. However, in a more general case, where the boundaries of the grid have other forms and internal structure (i.e. solid parts where no particles will flow) the simple strip-wise decomposition results in severe load imbalance. In this case, as was shown in [13], a more advanced decomposition scheme, the Orthogonal Recursive Bisection (ORB) method [14], still leads to highly efficient parallel LBM simulations. ORB restores the load balancing again, however at the price of a somewhat more complicated communication pattern between processors. In Fig. 3 we

show, for a representative 2D benchmark, the processor load. In a simple slice decomposition, the load on each processor differs a lot. For the ORB the load is almost balanced. This results for this benchmark, in reductions of execution times as large as 40%. For details we refer to [13].

For flow problems with a static geometry the ORB decomposition seems to be most appropriate. However, we are also interested in flow simulations where the geometry dynamically changes during the simulation. Here we refer to e.g. free flow around growing biological objects [15] or bounded flows with dynamically changing boundaries (e.g. blood flow in the heart). In this case an initially well-balanced parallel computation may become highly unbalanced. This may be overcome by a redundant scattered decomposition or by dynamic load balancing. Here we report some preliminary results of the first approach, the latter will be published elsewhere [16].

We consider growth of an aggregate in a three dimensional box. We must simulate flow around the object. The details of the streamlines close to the object determine its local growth. Due to the growth of the aggregate a straightforward decomposition (for example partitioning of the lattice in equal sized slices or boxes) would lead to strong load imbalance. To solve this problem we have tested two strategies to obtain a more equal distribution of the load over the processors:

1. Box decomposition in combination with scattered decomposition.
2. Orthogonal Recursive Bisection (ORB) in combination with scattered decomposition.



Fig. 4. Decomposition of an irregular shaped object in a 2D lattice. In this case 100 blocks are scattered over 4 processors.

The idea is to decompose the grid in a large number of partitions, much larger than the number of available processors, using either a box decomposition or ORB. Next, the partitions are randomly assigned to a processor. An example of a scattered decomposition over 4 processors of an irregular shaped object in a 2D lattice is shown in Fig. 4. In this example the lattice is divided into 100 blocks, where each block is randomly assigned to one of the four processors. Most of the computation is done in the blocks containing exclusively fluid nodes. The scattering of the blocks over the processors leads to a spatial averaging of the load, where decreasing block sizes cause a better load balancing, but also an increasing communication overhead. [17] Especially in simulations in which the shape of the object cannot be predicted,

redundant scattered decomposition is an attractive option to improve the load balance in parallel simulations [18, 19].

We have compared both decomposition strategies by computing the load balancing efficiency:

$$\varepsilon_{load} = l_{\min} / l_{\max} , \quad (2)$$

where l_{\min} is the load of the fastest process and l_{\max} the load of the slowest process. The two decomposition strategies were tested by using two extreme morphologies of the aggregates, a very compact shape and a dendritic shaped aggregate, and by measuring the load balancing efficiency during one iteration of the parallel LBM on 4 processors. The results are shown in Fig. 5. These results indicate that the combination of redundant ORB or box decomposition in combination with a scattered decomposition of partitions on processors may indeed improve load balancing. However, a disadvantage is that although the load balancing efficiencies increase with the number of redundant blocks, the communication overhead also increases. Furthermore, this test was for a single iteration only. We are currently working on testing these methods for real growing objects.




Fig. 5. The load balancing efficiency as a function of the total number of redundant blocks scattered on 4 processors, for box and ORB decomposition and for a compact and dendritic aggregate.

4. Cases: Flow in Complex Geometry

We applied our distributed particle simulation environment as described above for a large number of simulations of fluid flow in complex geometries. Here we show two representative examples, that of flow in random fiber networks and of flow in a static mixer reactor.

The first example is flow in a random fiber network, as drawn in Fig. 6. The fiber network is a realistic model of paper, and the question was to obtain the permeability of the network as a function of the volume fraction of fibers. Simulations were performed on 32 nodes of a Cray T3E using our parallel LBM environment described in the previous section. We obtained permeability curves that are in very good agreement with experimental results (see [20]).

Another impressive example is flow in a Static Mixer Reactor [21]. In such a mixer high viscosity fluids are mixed by letting them flow around a complex arrangement of internal tubes. A typical mixer is shown in Fig. 7. Here, LBM simulations and conventional Finite Element simulations were compared, which agreed very well. The simulation results also agree very well with experimental results. This shows that LBM, which is much easier to parallelize and much easier to extend with more complex modeling compared to Finite Element, (multi-species flow, thermal effects, reactions), is very suitable in real life problems involving complex flow.




Fig. 6. A random fiber network.




Fig. 7. A Static Mixer Reactor. Flow lines resulting from LBM simulations are also shown.

5. Conclusions

Within the general concept of Cellular Automata we have developed a distributed mesoscopic particle simulation environment for fluid flow. Parallel Lattice Gas Automata and Lattice Boltzmann methods have been realized, and we showed that by carefully taking load balancing into account it is possible to achieve very high parallel efficiencies. By means of two realistic examples of flow in complex geometries the power and potencies of such a distributed particle simulation environment were demonstrated.

References

1. Toffoli, T., Margolus, N.: Cellular Automata Machines - A New Environment for Modelling. MIT Press, Cambridge, MA (1987)
2. von Neumann, J.: Theory of Self-Reproducing Automata. In: A.W. Burks (eds.): University of Illinois Press, Champaign, IL (1966)
3. Talia, D., Sloot, P. (eds.): Special Issue: Cellular Automata: Promises and Prospects in Computational Science. Future Generation Computer Systems, (1999)
4. Sloot, P.M.A., Kaandorp, J.A., Hoekstra, A.G., Overeinder, B.J.: Distributed Cellular Automata : Large Scale Simulation of Natural Phenomena. In: A. Zomaya (eds.): Solutions to Parallel and Distributed Computing Problems: Lessons from Biological Sciences. John Wiley & Sons, (2001)

5. Mange, D., Tomassini, M. (eds.): Bio-inspired Computing Machines. Lausanne (1998)
6. Moore, C., Nordahl, M.G.: Lattice Gas Prediction is P-Complete. Santa Fe Institute, 97-04-034 (1997)
7. Wolfram, S.: Cellular Automata and Complexity. Addison-Wesley, (1994)
8. Frish, U., Hasslacher, B., Pomeau, Y.: Lattice-gas automata for the Navier-Stokes equation. Phys. Rev. Lett. **56** (1986) 1505
9. Chopard, B., Droz, M.: Cellular Automata Modelling of Physical Systems. Cambridge University Press, (1998)
10. Rothman, D.H., Zaleski, S.: Lattice-Gas Cellular Automata, Simple Models of Complex Hydrodynamics. Cambridge University Press, Cambridge (1997)
11. Chen, S., Doolen, G.D.: Lattice Boltzmann Method for Fluid Flows. Ann. Rev. Fluid Mech. **30** (1998) 329
12. Dubbeldam, D., Hoekstra, A.G., Sloot, P.M.A.: Computational Aspects of Multi-Species Lattice-Gas Automata. In: P.M.A. Sloot, Bubak, M., Hoekstra, A.G. & Hertzberger, L.O. (eds.): Proceedings of the International Conference HPCN Europe '99. Springer Verlag, Berlin (1999) 339-349
13. Kandhai, D., Koponen, A., Hoekstra, A.G., Kataja, M., Timonen, J., Sloot, P.M.A.: Lattice Boltzmann Hydrodynamics on Parallel Systems. Comp. Phys. Comm. **111** (1998) 14-26
14. Simon, H.D.: Partitioning of unstructured problems for parallel processing. Computing Systems in Engineering **2** (1991) 135-148
15. Kaandorp, J.A., Lowe, C., Frenkel, D., Sloot, P.M.A.: The effect of nutrient diffusion and flow on coral morphology. Phys. Rev. Lett. **77** (1996) 2328-2331
16. Schoneveld, A., de Ronde, J., accepted for publication in Fut. Gen. Comp. Syst., 1999.
17. de Ronde, J., Schoneveld, A., Sloot, P.M.A.: Load balancing by redundant decomposition and mapping. In: H. Liddell, Colbrook, A., Hertzberger, B. & Sloot, P. (eds.): High Performance Computing and Networking (HPCN'96). (1996) 555-561
18. Machta, J., Greenlaw, R.: The parallel complexity of growth models. Journal of Statistical Physics **77** (1994) 755-781
19. Machta, J.: The computational complexity of pattern formation. Journal of Statistical Physics **70** (1993) 949-967
20. Koponen, A., Kandhai, D., Hellin, E., Alava, M., Hoekstra, A., Kataja, M., Niskanen, K., Sloot, P., Timonen, J.: Permeability of three-dimensional random fiber webs. Phys. Rev. Lett. **80** (1998) 716-719
21. Kandhai, D., Vidal, D., Hoekstra, A., Hoefsloot, H., Iedema, P., Sloot, P.: Lattice-Boltzmann and Finite-Element Simulations of Fluid Flow in a SMRX Static Mixer. Int. J. Num. Meth. Fluids **31** (1999) 1019-1033