LABORATORY FOR
COMPUTER SCIENCE

MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TM-296

CELLULAR AUTOMATA SUPERCOMPUTERS FOR FLUID DYNAMICS MODELING

NORMAN MARGOLUS
TOMMASO TOFFOLI
GERARD VICHNIAC

DECEMBER, 1985

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Technical Memo LCS-TM-296

Cellular Automata Supercomputers for Fluid Dynamics Modeling*

Norman Margolus, Tommaso Toffoli, and Gérard Vichniac
Laboratory For Computer Science
Massachusetts Institute of Technology, Cambridge, MA 02139

December 12, 1985

Abstract We report recent developments in the modeling of fluid dynamics, and give experimental results (including dynamical exponents) obtained using cellular automata machines. Because of their locality and uniformity, cellular automata lend themselves to an extremely efficient physical realization; with a suitable architecture, an amount of hardware resources comparable to that of a home computer can achieve (in the simulation of cellular automata) the performance of a conventional supercomputer.

1 Introduction

A cellular automaton is a discrete dynamical system consisting of finite-state variables, or *cells*, arranged on a uniform grid. The overall dynamics is specified by a finite *rule*, by which at every time-step each cell computes its new state from the current state of its neighborhood.

Surprisingly enough, cellular automata can faithfully model *continuum* systems such as fluids; unlike differential equations, they can be realized *exactly* by digital hardware.

Modeling with cellular automata is poorly supported by conventional scientific computers, whose architecture is optimized for the arithmetic treatment of continuum models. With a more appropriate architecture one can easily gain a performance factor of at least 10,000 in the simulation of cellular automata; this gain is of such magnitude that new classes of conceptual models have become computationally accessible.

The idea of using discrete lattices for modeling physical phenomena is not new. However, recent theoretical and technological developments have turned models based specifically on cellular automata into practical computational tools. On one hand, methods have been found for constructing cellular automata that are microscopically reversible (and thus support a realistic thermodynamics), obey exact conservation laws, and model continuum phenomena[1–10]. On the other hand, general-purpose machines well-suited to such fine-grained modeling are becoming generally available (see Appendix).

2 Hydrodynamic modeling

Differential equations such as the Navier-Stokes equation capture important macroscopic aspects of fluid dynamics; however, what one implements on a digital computer is not the equation itself, but a finitary model obtained from it by truncation and round-off.

It is possible to arrive at an analogous macrodynamics starting directly from a discrete microscopic model—a cellular-automaton idealization of the motion and collisions of individual particles. Models of this kind can give rise to the Navier-Stokes equation in the macroscopic limit[9], as had been indicated as early as 1973 by Pomeau and coworkers[1,2] in a theoretical analysis of a lattice gas model (hereafter, HPP gas).

Frames (a), (b), and (c) of Fig. 1 are taken directly from the display screen of CAM-6[12], a cellular automata machine. They show the evolution of an HPP

*This research was supported by grants from the Defense Advanced Research Projects Agency (N00014-83-K-0125), the National Science Foundation (8214312-IST), and the Department of Energy (DE-AC02-83ER13082).

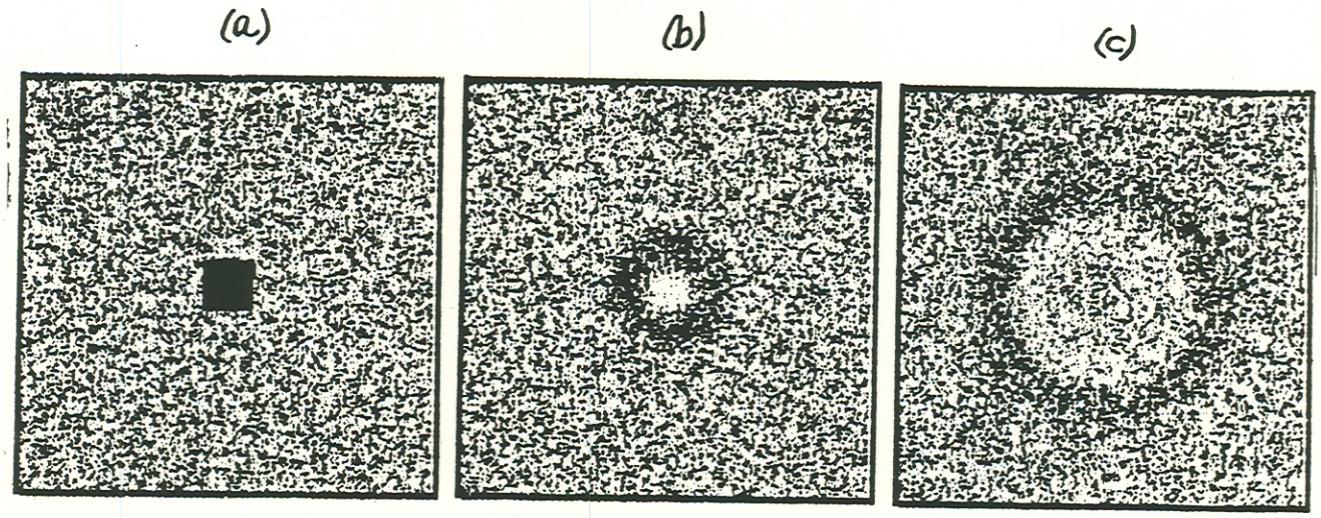

Figure 1: Wave propagation in the HPP gas. To enhance contrast, only sites that contain either 3 or 4 particles are shown.

gas consisting of 2^{16} sites (256×256) each of which can hold up to four particles (one traveling in each of four allowed directions). The evolution rule is simply that particles travel straight at unit velocity (one cell per time-step) unless exactly two particles collide head-on, in which case they scatter at right angles.

The initial state (a) was constructed with a bit-occupancy of 50% except in the middle, where we have a block of 4096 particles (100% occupancy in an area of 32×32 sites). Thus, the simulation involves about 130,000 particles. Frames (b) and (c) show the state after 30 and 90 steps—0.5 and 1.5 seconds at CAM's rate of 60 frame-updates/sec. Despite the fact that particles travel in only four directions the wave is circular, and moves at a speed that agrees with the theoretically predicted[2] value of $1/\sqrt{2}$.

The HPP rule is exactly reversible. To go back in time from frame (c) one need only transform the state of every cell so as to interchange the UP and DOWN particle information, and similarly for RIGHT and LEFT. If one then proceeds with the same rule the simulation will retrace its steps back to (a).

Boundary conditions such as sources, sinks, and obstacles of any shape are introduced by using additional bits of state at each site to mark selected areas, and extending the rule so as to take the values of these bits into account[10].

3 Time correlation functions

We have measured the time autocorrelation function for velocity at a given site

$$\nu(t) = \lim_{T\to\infty} \frac{1}{TNQ} \sum_{t'=1}^{T} [a_{ij}^{q}(\phi^{t+t'}C) - \bar{a}][a_{ij}^{q}(\phi^{t'}C) - \bar{a}]$$

where C is an initial configuration, ϕ is the transition rule (such as HPP), $a_{ij}^{q}C$ has a value of 1 or 0 depending on whether or not there is a particle moving in the q direction at the (ij) site in the configuration C, $\bar{a}$ is the average particle occupancy (per site) for each direction, and there is an implied summation over all N sites and all Q allowed directions.

The actual correlation experiments were performed using a cellular-automata realization of the HPP rule[13] that spreads each site over four one-bit cells[4,10]. Figure 2(a) shows the measured values of $\nu(t)$ for the HPP model, using $T=2^{13}$ and a density of $\bar{a}=1/4$. Finite-size effects show up past $t=256$ (the space is 256×256 and has periodic boundary conditions), but already before that point the predicted[14] asymptotic exponent of $-2/3$ is attained within $\pm 3\%$. The same exponent was obtained for $\bar{a}=1/6$, $\bar{a}=1/8$.

The exponent $-2/3$, which is characteristic of one-dimensional gases[14], arises from extra conservations (momentum on each row and column). The "TM" gas[10,11]—a similar model which also uses only four directions—avoids these extra conservations by having collisions occur with a nonzero impact parameter (resulting in right-angle scattering from two adjacent rows to two adjacent columns, or vice versa). This gas is shown in Fig. 2(b), using the same density $\bar{a}$ as for HPP but (to improve the statistics) using

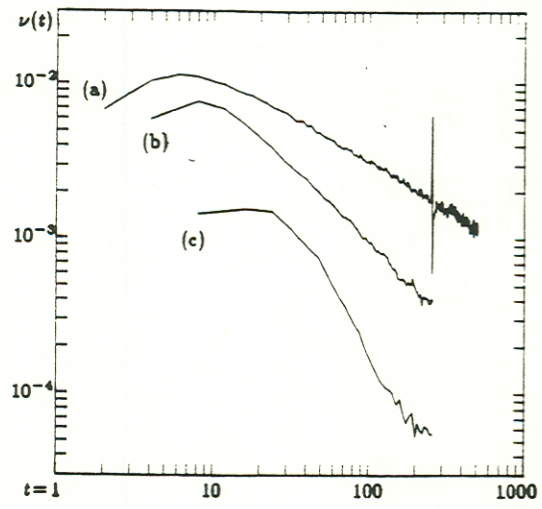

Figure 2: Time correlation function, $\nu(t)$, for the HPP gas (a), the TM gas (b), and the FHP gas (c); $\bar{a}$=1/4 in all cases.

T=2^{16}; note that the asymptotic exponent for this gas is close to -1, characteristic of a true two-dimensional gas ($\nu \sim t^{-d/2}$, for $d \geq 2$)[15]. The same exponent was obtained for $\bar{a}$=1/6, $\bar{a}$=1/8.

As is noted in [9], a shadow of the $\pi/2$ rotational symmetry persists in the macroscopic behavior of the HPP gas (and possibly also in TM). A more refined hydrodynamical model (the "FHP gas"[9,16]), which uses a hexagonal grid, avoids this problem[17]. Fig. 2(c) shows the behavior of $\nu(t)$ for this gas, using T=2^{17} and again using the density $\bar{a}$=1/4. The asymptotic exponent appears to be close to -2, which is surprisingly large and calls for a theoretical explanation—the same exponent was obtained for $\bar{a}$=1/6, $\bar{a}$=1/8. As above, this experiment used an implementation in which collision sites are spread out—in this case over four two-bit cells[10].

4 Hardware

A machine such as CAM-6 uses an amount of digital logic comparable to that of a home computer and achieves, in the simulation of cellular automata, a performance comparable to that of a CRAY-1; several modules can be ganged together—with a proportionate increase in performance—for larger two-dimensional arrays or for three-dimensional simulations.

To achieve maximum speed, the rule is internally stored as a look-up table. The rule as written by the user consists of a few lines describing in a high-level language how the new value of a cell depends on the current value of its neighbors. The problem of translating such a description into an appropriate look-up table is taken care of by the machine's designers; since the efficiency of this compilation process in no way affects the speed of the simulation, one's full attention can be kept on conceptual issues.

The above performance is achieved with a dedicated architecture but conventional circuitry and components. In this architecture two-dimensional planes are processed *serially* (with a substantial amount of pipelining); a third dimension is achieved by stacking planes and operating them in parallel[18]. Since sites in each plane are updated one by one, and corresponding sites in adjacent planes are handled at the same time in a "lockstep" fashion, communication between planes entails a few wires rather than the millions of physical interconnections required by a fully-parallel implementation. This approach makes extensive simulation of three-dimensional models of hydrodynamics immediately practical.

One may simultaneously run (on different planes of the same machine) two copies of the same system that are identical except for a given spatial or temporal offset between them. Since sites are processed serially, correlations between corresponding sites are easily detected and accumulated "on the fly," thus eliminating the need for storage- and computation-intensive post-processing. This is, in fact, how the time correlation experiments presented in this paper were conducted.

5 Conclusions

Cellular automata machines are well suited to a large class of computational models of physics having both theoretical and practical interest. The performance they offer provides strong encouragement for the development of models of this kind; conversely, the usefulness of these models will stimulate technology to provide further performance in this direction.

Because of the speed-of-light constraint, locality of interconnection is an important advantage of the cellular-automaton paradigm. A fully-parallel implementation of specific two-dimensional cellular automata having 10^{12} sites and an update cycle of 100

picoseconds for the *whole array* will be feasible in one decade and within easy reach in two; one using 10^{16} sites (the Avogadro number "in two dimensions") is not inconceivable. Thus, we can look forward to computational tools that *directly span the gap between the microscopic and the macroscopic world.*

Appendix

CAM-6[12] is derived from CAM-5, a general-purpose cellular automata machine[19,20,21] designed at the MIT Laboratory for Computer Science. CAM-7, a much larger version, is under development[18]; the basic module of this machine (1/4 billion sites, 10 billion site-updates per second) will be capable of three-dimensional simulations (a 512×512×512 cube) or much larger two-dimensional simulations (16K×16K).

The Connection Machine[22] and other heavily parallel architectures can be adapted to perform usefully in some cellular automata contexts[23].

6 Acknowledgements

We would like to thank G. Maturana, U. Frisch, L. Kadanoff, G. Parisi, Y. Pomeau, and S. Zaleski for useful discussions.

References

[1] J. Hardy, Y. Pomeau, and O. de Pazzis, "Time evolution of a two-dimensional model system. I. Invariant states and time correlation functions," *J. Math. Phys.* 14 (1973), 470.

[2] J. Hardy, O. de Pazzis, and Y. Pomeau, "Molecular dynamics of a classical lattice gas: Transport properties and time correlation function," *Phys. Rev.* **A13** (1976), 1949.

[3] T. Toffoli, "Cellular automata as an alternative to (rather than an approximation of) differential equations in modeling physics," *Physica* **10D** (1984), 117.

[4] N. Margolus, "Physics-like models of computation," *Physica* **10D** (1984), 81.

[5] G. Y. Vichniac, "Simulating physics with cellular automata," *Physica* **10D** (1984), 96.

[6] Y. Pomeau, "Invariant in cellular automata," J. Phys. **A17** (1984), L415.

[7] M. Creutz, "Deterministic Ising Dynamics," *Annals of Physics* (in press).

[8] G. Y. Vichniac, "Cellular automata models of disorder and organization," in *Disordered Systems and Biological Organization* (E. Bienenstock, F. Fogelman, and G. Weisbuch, ed.), Springer-Verlag (in press).

[9] U. Frisch, B. Hasslacher, and Y. Pomeau, "A lattice gas automaton for the Navier Stokes equation," Preprint LA-UR-85-3503, Los Alamos National Laboratory (1985), submitted to *Phys. Rev. Lett.*

[10] T. Toffoli, N. Margolus, *Cellular Automata Machines: A new tool for modeling*, to be published by MIT Press, Cambridge, MA.

[11] Uses the same neighborhood as the BBMCA[4]. The rule is: rotate contents of neighborhood 90° clockwise on even time-steps and counterclockwise on odd time-steps unless there are exactly two ones and they are on a diagonal.

[12] SYSTEMS CONCEPTS, *CAM-6 Technical Manual*, San Francisco (1986).

[13] Since the new state of a site in the HPP model depends only on that of the nearest neighbors and not on its current state, the lattice reduces to two uncoupled sublattices. Our realization naturally yields a single sublattice, which is computationally more efficient, and has other technical advantages for use on a cellular-automata machine.

[14] Y. Pomeau and P. Résibois, "Time dependent correlation functions and mode-mode coupling theories," *Phys. Reports* **19** (1975), 63.

[15] B. Alder and T. Wainwright, "Decay of the Velocity Autocorrelation Function," *J. Phys. Soc. Jap.* **26** (1968), 267; *Phys. Rev.* A 1 (1970), 18. This discussion treats both our case (correlation at fixed site) and the standard fixed particle case.

[16] D. d'Humières, P. Lallemand, and T. Shimomura, "Lattice Gas Cellular Automata, a

New Experimental Tool for Hydrodynamics," Preprint LA-UR-85-4051, Los Alamos National Laboratory (1985), submitted to *Phys. Rev. Lett.*

[17] In our implementation, head-on collisions result in clockwise scattering on even time-steps, counterclockwise on odd steps; three particles colliding symmetrically bounce back, all other cases go straight through. The wraparound aligns two of the three directions with the major axes of a torus.

[18] T. Toffoli, N. Margolus, "The CAM-7 Multiprocessor: A Cellular Automata Machine," Tech. Memo LCS-TM-289, MIT Laboratory for Computer Science (1985).

[19] T. Toffoli, "CAM: a high-performance cellular-automaton machine," *Physica* **10D** (1984), 195.

[20] B. Hayes, "The cellular automaton offers a model of the world and a world unto itself," *Sc. Am.* **250**:3 (1984), 12.

[21] J. B. Tucker, "Cellular automata machine: the ultimate parallel computer," *High Technology* **4**:6 (1984), 85.

[22] D. Hillis, *The Connection Machine*, MIT Press, Cambridge, MA (1985).

[23] J. B. Salem and S. Wolfram, "Thermodynamics and Hydrodynamics with Cellular Automata" (1985), submitted to *Phys. Rev. Lett.*