

Three-dimensional visualization of many-body system dynamics

by M. Bernaschi
E. Marinari
S. Patarnello
S. Succi

This paper describes a graphic rendering system for use in visualizing the behavior of three-dimensional physical systems. The tool is general and allows the user to characterize a great variety of phenomena. The only requirement is that the physical system be represented by variables defined on quantifiable positions (sites) within a three-dimensional grid. The variables may be discrete (e.g., binary), real, or even complex numbers. The first part of the paper gives a technical description of the graphic program, which is based on a graPHIGS¹ interface; two versions of the code (in the C and FORTRAN languages) are available. The hardware platform consists of an IBM 5080 graphic workstation with a 5081 high-resolution monitor which can be driven either by a machine employing IBM System/370¹ architecture with VM/XA¹ (in our case a 3090¹ processor running under VM/XA) or by a RISC System/6000¹ workstation [we have used both an IBM RT²

System and (recently) an IBM RISC System/6000 processor] running under AIX.² The second part of the paper describes three different examples of the application of this tool: discrete spin models, quantum chromodynamics (QCD), and three-dimensional turbulence. For spin systems and QCD, the physical problem consists in understanding the nature of the phase transition from disordered to ordered states of the system. In both cases a direct (i.e., through visualization) investigation of the system configurations reveals valuable information about properties such as the order of the transition, the behavior of the correlation length, and phase coexistence. We note, however, that the meaning of the site variables is very different in the two cases. In particular, for QCD the site variables are complex numbers, which we code by using a color table to represent the phase of the number and pixel size to represent a value proportional to the modulus. This kind of coding is also used for three-dimensional turbulence. Here the analysis can show where dissipation phenomena take place in the fluid and characterize the geometrical nature of the set of dissipative structures.

¹ graPHIGS, System/370, VM/XA, 3090, and RISC System/6000 are trademarks of International Business Machines Corporation.

² RT and AIX are registered trademarks of International Business Machines Corporation.

©Copyright 1991 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

Introduction

Computer simulations are adding a new dimension to the world of scientific investigation, and computational physics has emerged as an invaluable third alternative to the two traditional avenues of analytic theory and direct physical experimentation. The impressive growth of computing power made possible by technological advances permits today's scientists to attack complex problems whose solution would have been inconceivable only a few years ago. This increased power has begun to satisfy an increasingly compelling need not only to "crunch numbers" and solve equations, but also to display and manage the voluminous data produced by numerical simulation. A clear and global visual representation improves enormously the quality of numerical approaches to the comprehension of physical problems. In this respect, a key role is played by computer graphics, which has in itself gained the rank of a software engineering discipline.

In this paper we present a new graphic tool, which we have used to analyze the behavior of a variety of complex systems arising in different domains of theoretical physics. The underlying physical problems are quite diverse, but our technique can be applied to all of them, the only requirement being that they can be formulated in terms of local variables defined in a discrete three-dimensional space. The test problems discussed in this paper include the study of phase transitions in magnetic systems, the origin of quark confinement in hadronic matter, and the nature of three-dimensional fluid turbulence.

The paper is organized as follows: In Section 1 we briefly describe the main technical features of the code and the computing environment in which the code runs. In Section 2, applications to many-body systems at equilibrium are reported, with particular attention to statistical physics and quantum chromodynamics (the theory which describes interactions among quarks). Section 3 is devoted to the use of this tool in the context of three-dimensional fluid flows (that is, of chaotic systems that are not in equilibrium). Finally, Section 4 outlines ongoing and future developments for this project.

1. Graphic tool

The tool we have developed is primarily oriented toward the visualization of data obtained from simulations of physical systems defined on a lattice. Redefining on a lattice a theoretical model that was originally formulated as a continuum is often a prerequisite for numerical simulation. It is also necessary, of course, to extrapolate back to the continuum case the results obtained on the lattice, and our tool can be very effective for that purpose as well.

Two directives have been always followed during the design and implementation of the tool: simplicity of use (we need something that simplifies our work, not a further complication) and enough flexibility and generality that the tool can be applied with minor effort to the visualization of data obtained from simulations of many different physical systems. We chose as display station for the tool an IBM 5080 graphics workstation with a 5081 high-resolution (1024×1024) monitor. The choice was justified by the powerful capabilities of the 5080 monitor; zooming, 3D rotations, clipping, and perspective view can all be effectively executed in an interactive way. On the other hand, the 5080 is not a standalone workstation; it must be guided by a host machine, either a System/370¹ architecture with VM/XA¹ or a powerful workstation such as the RT² System or RISC System/6000¹ processor. Our goal was to be able to use the tool on both hardware platforms, which requires a high degree of code portability. In order to achieve this goal, we wrote both a FORTRAN and a C version of the tool, always using *grPHIGS*¹ as the graphics interface. However, we actually did most of the work on the RT because we consider the AIX² environment to be more suitable for application development. Moreover, AIX offers much more support for interprocess communications, and we were interested in this aspect for reasons to be explained later.

The structure of the tool can be rapidly summarized as follows (the description is more faithful to the RT version): As a general rule we tried to split the code into several "atomic" parts (subroutines in FORTRAN or functions in C) so that changes and updates could be easily located. The first part of the code reads the data to be displayed. We always used unformatted data to speed up the reading and reduce disk space occupancy. When data are moved from System/370 machines (where they are often produced) to RTs, a set of routines makes conversions between the respective internal formats. The next step is the initialization of the *grPHIGS* environment, which involves the specification of things such as the display station and the color table that will be used. The most interesting parts of the code are those used to define the *grPHIGS* structures and manipulate the structures themselves. We manipulate the data to extract all possible information; part of the information is displayed in graphic form and another part is shown in text form on the screen. Three different output primitive elements are employed: blocks of pixels, polymarkers, and polylines. For instance, when we show results from a simulation of a system of spins that can take complex values, we represent every spin as a pixel block, the size of the block being proportional to the modulus and the color indicating the phase of the complex number according to a scale shown in a legend on the screen.

Finally, we report the "global magnetization" of the system as a text string on the higher part of the screen (Figures 1-4, shown later). The construction of *grPHIGS* structures is the CPU-bound part of this visualization code. For the more complex cases (e.g., when we deal with large lattices or with complex-valued spins), the RT could not offer enough computing power. (In the last section we briefly present possible solutions to this problem.) In general, the rendered *grPHIGS* structures represent 3D lattices; parameters (such as visualization angles) defining their representation have values at this stage that are read from an input file. As we define this "default" representation form for the structures, we also initialize and set all the 5080 devices (dials, mouse, tablet, etc.) that will be used later. When this operation is completed, the graphic structures are sent to the 5080 to be displayed; the features of the 5080 can then be exploited for interactive manipulation of the displayed lattice.

As an aid to better understanding of the usefulness of the tool, we provide some examples of typical applications. In Section 2 we find that one of the most interesting observable attributes of a spin system is the correlation length. With our graphic representation, the order of magnitude of the correlation length can be extrapolated from the size of those spin blocks (spins which in turn are represented by pixel blocks) that are distinguished by the same color. The view we have at the beginning of the display phase is global (Figures 1-6, shown later); at this level, for instance, one can check by zooming whether a certain area of the lattice is *coherent*, or whether there are *impurities* indicated by spins of a different color. Then, by clipping and realigning the lattice, we can isolate one or more planes so that it is possible to distinguish among 1D, 2D, and 3D coherence zones. Finally, 3D rotations allow us to look at the entire lattice regardless of its initial position, without having to restart the program from the beginning.

Another interesting and useful feature of the tool is the ability to elaborate and manipulate more than one configuration at a time. This means that we can read a certain number of different configurations of the same physical system (e.g., configurations related to different temperatures), manipulate these configurations, and then by changing in rapid sequence the configuration that is displayed, observe the *evolution* of the system. It is possible to gain an immediate comprehension of the simulation dynamics with this kind of visualization. The sequence structure in the "fast visualization" is circular so that, at the end, the sequence restarts itself from the first displayed configuration; moreover, changes such as zooming or 3D rotations, made interactively during the visualization, can be preserved during the rest of the sequence. The transition from one configuration to

another can be made simply by pressing one of the buttons on the 5080 mouse, and we stress that the updating of the image is immediate, since the definition of the image structures is made only once, at the beginning. The number of configurations that we can see in sequence is limited only by the size of the 5085 memory buffer. Because the memory used by a given configuration depends on both the lattice size and the complexity of the physical system, we typically have from 4 to 16 different configurations.

A few comments on limitations related to the choice of chromatic scale are in order. Quantitative judgments based on particle color can be biased by the different chromatic significance of the three (red, green, and blue) base colors, as well as by the fact that in the absence of depth-cueing, some pixels at the "back" of the image may be hidden by those at the "front." Improvements planned or in development are described in the last section.

Real physical problems studied with the aid of the tool are introduced and discussed in Sections 2 and 3.

2. Phase transitions and critical phenomena

• Discrete spin models

Equilibrium statistical mechanics is one of the tools used by physicists to describe many-body systems. In dealing with a system with many elementary components, one must distinguish between the *microscopic* state of the system and its *macroscopic* description. A microscopic state is defined in terms of the specific values of the coordinates which describe the motion of each elementary component. For a gas of N atoms this would correspond to providing the $3N$ values of the position and the $3N$ values of the velocity for each of the atoms. In practice, one cannot track the microscopic description of the system in real time. However, there is another approach which is successful: Assuming that the evolution of the system over time (which is a trajectory in the $6N$ -fold space of the system coordinates) can be reproduced by a statistical distribution, one can determine the *probability* that a given microscopic state will occur. Pursuit of this strategy reveals that the probability *distribution* for a microscopic state is a simple function of a single, macroscopic quantity, the *energy* E of the system. Let S be a microscopic state with energy E ; the probability that state S will occur during the evolution of the system is

$$p(S) = \frac{e^{-E/k_B T}}{Z}, \quad (1)$$

where Z is the *partition function* of the system,

$$Z = \sum_S e^{-E/k_B T}, \quad (2)$$

which normalizes the probability $p(S)$ so that $\sum_S p(S) = 1$. In this formula, the temperature T of the system appears, multiplied by a dimensional factor, the Boltzmann constant k_B , whose numerical value depends on the system of units chosen. Since one can choose a system where $k_B = 1$, we omit the Boltzmann constant hereafter in our notation. The probability $p(S)$ is often referred to, however, as the "Boltzmann weight" of the microscopic state.

The solution provided in this framework consists in labeling a microscopic state in terms of a single, macroscopic variable (its energy) and expressing the probability for that state in terms of that variable. From this basic relation one is able in principle to compute all macroscopic quantities as average values. Given the dependence of the Boltzmann weight on T , it is quite evident that when temperature is large compared with the typical energy scales for the system, the probability is the same for any given microscopic state. When T decreases, however, the role of energy E becomes important, and as T tends to 0, the states most likely to be entered by the system are those of *lowest energy*.

Statistical physics applied to many-body systems has been found very useful in the study of *phase transitions*. A phase transition is characterized, roughly speaking, by the transition from one state having a certain degree of order to another state of greater or lesser order under the influence of a control parameter, which in most situations is the temperature T . An example is the phase transition which characterizes a material: A liquid that becomes solid when cooled down clearly becomes more ordered (organized in a crystalline structure). Conversely, when the temperature of such a liquid is increased, it becomes a vapor or gas, which is a state of increased disorder. This description is of course rather general: In a real phase transition there is usually a many-body system, which may exhibit both ordered and disordered states (a thorough review of the subject is provided by [1]).

Most systems in phase transition show certain peculiar properties called *critical phenomena*. Looking at the optical properties of a pot of boiling water, one can see regions where light is scattered in an unusual fashion, a feature known as the *critical opalescence* of the system. The bubbles of vapor which grow in the liquid behave as "scattering centers" which affect the refractive behavior of the medium. These bubbles are evidence that at the transition point, phenomena of *phase coexistence* take place: Some regions of the more disordered phase (the gaseous one) appear inside the more ordered one (the liquid). Another feature to note in the example of boiling water is the presence of gaseous bubbles of *all sizes* in the system, a property which is often referred to as *scale invariance*. All length scales are dynamically relevant: Any event at a point P is able to influence any other

point P' , no matter how far apart the two points may be. This concept can be made more rigorous: A *correlation length* can be mathematically defined corresponding to the maximum distance over which cooperative effects have long-range influence. The divergence of the correlation length at the point of transition is another interesting property of critical systems.

Phase coexistence and scale invariance do *not* occur in all phase transitions. Those transitions which show such properties are defined as *second-order* transitions, as opposed to those showing *first-order* behavior. An example of the latter is the transition from the liquid to the solid phase in materials (one cannot see "bubbles" of ice of all sizes inside a body of freezing water). The distinction between first-order and second-order behavior is one of the key points in understanding systems at criticality, as we discuss below.

The class of critical phenomena we consider in this paper is related to the *magnetic* properties of matter. Magnetism is a very well-known phenomenon. A bar of iron can behave as a magnet and affect with its action other metallic materials with similar magnetic properties, because it produces a *magnetic field* in its neighborhood. Magnetic properties are strongly affected by the temperature of the system, and for most magnetic materials there exists a *critical temperature* T_c . For $T > T_c$, thermal disorder prevails, and the magnetic properties of the material are somehow suppressed. At $T = T_c$, ordering phenomena start to take place, and at temperatures lower than T_c even an infinitesimal magnetic field is able to order the system, which then becomes a magnet. At a microscopic level, cooperative effects take place; each atom of the system has its own magnetic moment, or *spin*, and behaves as a tiny magnet. Below T_c the interactions among the spins tend to align them in the same direction, and the macroscopic effect which stems from this cooperation is the magnetic field produced by the magnet.

Many models have been introduced to describe magnetic phase transitions, the most popular being the *Ising model* [1, 2], which was introduced by E. Ising in 1925. Here the magnetic moment of the atom is modeled as a binary variable; i.e., an *Ising spin* may take only two values, which in general are chosen to be +1 or -1. This means that one is considering only the projection of the magnetic moment (which, strictly speaking, is a vector in three-dimensional space) along one direction. Thus, the only remaining degree of freedom is the orientation of the spin, which may point "up" (corresponding to the state +1) or "down" (state -1). This approximation may seem very crude, but in fact the model at this stage shows many interesting features.

To study an interacting spin system, it is natural to assign such spins to a regular crystalline structure, or

lattice. If each point of this regular structure is labeled with an index i , each site will be occupied with an Ising spin s_i . In considering the interaction among these magnetic moments, the simplest notion is a *local* interaction, which couples sites which are neighbors on the lattice and is proportional to the product of the two spin variables. This leads to the following expression for the *total energy* of the system:

$$E = - \sum_{\langle ij \rangle} K s_i s_j. \quad (3)$$

The coupling constant K gives the strength of the interaction, which is homogeneous over all the lattice. The constraint $\langle ij \rangle$ limits the sum to nearest-neighbor sites of the lattice. The energy distribution favors configurations where spins are aligned ($s_i = s_j$). In fact, the lowest-energy states of the system are completely ordered: Either all spins are +1 (we define this state as S_+) or they are -1 (state S_-).

The critical behavior of the Ising model is very simple: At high temperature, the action of thermal disorder is such that each spin points randomly up or down, because it is not strongly influenced by neighboring spins. On the other hand, at very low temperatures the probability of either S_+ or S_- is exponentially dominant, and the system becomes ordered. This competition between *entropic* and *energetic* effects becomes maximal at an intermediate temperature, which is the critical temperature T_c of the system.

In this context the symmetry of the system is the reversal of the spin $s_i \rightarrow -s_i$. This is the same symmetry group as the permutation group for two elements, the Z_2 group. A disordered state has no preferred orientation, as spins randomly point in the +1 or the -1 direction. The dynamics of the system below T_c , however, are such that it chooses one of the two states S_+ or S_- . In this case the symmetry is lost; reversing all spins simply changes one ordered state into the other.

There is a very simple quantity which precisely characterizes the disordered state as well as S_+ and S_- . This is the *magnetization density* of the system:

$$m = \frac{1}{N} \sum_i s_i = \frac{1}{N} (N_+ - N_-). \quad (4)$$

Here N is the total number of spins, N_+ is the number of spins in the state +1, and N_- is the number of spins in the state -1. The magnetization gives the degree of alignment of the spins of the system. In a disordered state, on average, $N_+ = N_-$ and $m = 0$. For S_+ , m is 1, whereas for S_- , m is -1. Magnetization is also referred to as the *order parameter* of the system. Its behavior at the critical point is another criterion for distinguishing second-order from first-order transitions. Since in the former case the two phases coexist, this transition must

be *continuous* at T_c . So must be the behavior of the order parameter as a function of temperature, across the transition from one phase to the other. Conversely, in the case of a first-order transition the change in the order parameter is abrupt; it is essentially a step function from the value it takes in the ordered phase to that in the disordered phase.

A set of quantities that are important in describing phase transitions are the so-called *critical exponents*. These characterize the behavior of many relevant physical attributes at the critical temperature T_c . An example is the aforementioned correlation length ξ , which becomes infinite at T_c . The critical exponent ν describes this divergence:

$$\xi \sim |T - T_c|^{-\nu}. \quad (5)$$

Other exponents are defined at the critical point. Many *response functions*, which measure the sensitivity of the system to external solicitations, become infinite. One example is the *specific heat* of the system, whose divergence near T_c is described by the exponent α :

$$C \sim |T - T_c|^{-\alpha}. \quad (6)$$

The set of critical exponents that one measures on a system defines the analytical behavior at the critical point and thus identifies the phase transition in a precise way. Thus, one is able to classify all second-order phase transitions by looking at the set of critical exponents at T_c . Another important element in the theory of critical phenomena is the property of *universality*: The set of exponents of a phase transition are uniquely defined by the dimensionality of the model and by its internal symmetry. Other microscopic details of the interaction do not change behavior at T_c , and two models defined in the same number of space dimensions, with the same kind of internal symmetry, fall in the same *universality class*. This also explains why even a simplified model such as the Ising model mimics a magnetic phase transition quite well.

The Ising model has been studied for different values of the spatial dimensionality d . Most real magnetic materials exist in $d = 3$ (and we can also perform some revealing experiments in $d = 2$). However, understanding the behavior of models defined in other dimensions is relevant for at least two reasons: It lends insight to the study of the three-dimensional model, and it is also useful in the analysis of different universality classes.

In $d = 1$ the model represents a chain of interacting spins. The critical properties in this case are trivial, because ordering takes place only at $T = 0$; no phase transition occurs at any finite temperature. The two-dimensional model is much more interesting; its critical behavior was determined by means of an *exact* analytical computation by Onsager [3] in 1944. At this stage, taking

correlations at all scales into account makes the computation highly nontrivial. No such exact solution exists for Ising models in higher (finite) dimensions. The critical temperature and some critical exponents in $d = 2$ are shown in Table 1.

The Ising model in $d = 3$ also undergoes a phase transition at finite temperature, which has been studied so far using approximate techniques. Some estimates for critical quantities are also shown in Table 1.

Among the many approximate techniques used to study these complex systems, numerical methods are becoming increasingly important. This is primarily because modern supercomputers can perform numerical computations which were impossible ten years ago. This is particularly true for Monte Carlo simulations [4, 5], which are used to reproduce the Boltzmann probability distribution of the system. This is obtained by generating a chain of microscopic states, with Markovian probability rules for the elementary transition [the $(i + 1)$ th state of the chain depends on the i th state]. When the simulation is long enough (and statistical equilibrium is achieved), the probability of generating a state in the chain converges to the Boltzmann weight, which we have defined in Equation (2).

Thus, Monte Carlo simulation is in principle a wonderful tool, because it permits the simulation of a process in which temperature is varied, and ultimately the phase transition can be reproduced. One can measure all interesting averages and some of the critical exponents. Many difficulties are encountered, however, in simulating complex systems, the two most important ones being probably the following:

- The very method of generating states via a Markov chain induces correlations among such states; these are particularly relevant near the critical temperature (and, eventually, correlation times diverge at T_c) and strongly bias the average values of measured quantities. This problem is known as *critical slowing down*.
- Since the computational power and memory of a computer, though large, are finite, we can afford only simulations of *finite* systems (i.e., a finite number of lattice points); the more complex the system, the smaller the size of the lattices we are able to simulate. Such *size effects* are again a strong source of errors in the measurements.

The problem of discriminating among first- and second-order transitions is strongly affected by these limitations. Looking at the behavior of the order parameter (e.g., magnetization in the Ising model), how can one identify a discontinuity for a finite system? True singular behavior does not occur for a system with a finite number of degrees of freedom. To circumvent this difficulty, many

Table 1 Correlation length (ν) and specific heat (α) critical exponents for Ising model in two and three dimensions. The value $\alpha = 0$ means that in $d = 2$ the specific heat diverges as a logarithm.

d	ν	α
2	1	0
3	0.628	0.116

strategies have been introduced. The most commonly used is *finite size scaling*, in which the dependence of physical quantities on the size of the system is used to obtain information on the transition [6]. As an example, the inverse critical temperature $\beta_c = 1/T_c$ has the following dependence on the size of the system:


$$\beta_c(L) - \beta_c = L^{-1/\nu}. \quad (7)$$

Here $\beta_c(L)$ is the inverse critical temperature for a finite system of linear size L , whereas β_c is the same quantity for the infinite system and ν is the critical exponent for the correlation length ξ . Similar scaling laws, which can be used to determine critical exponents, hold for many physical quantities. This is an ingenious method, as it exploits a limitation of the strategy (the finite size of the simulated system) to recover useful information.


However, it is not absolutely reliable, because simulation within the scaling region (near T_c) is subject to critical slowing down, which means longer simulation times. It also implies that correlation length ξ is very large.

Moreover, if one is simulating systems with $\xi < L$, the fact that scaling holds for nontrivial critical exponents cannot be ignored, even for a first-order phase transition. One is forced once again to larger and larger systems.

Thus, the problem of understanding the order of a phase transition via a numerical simulation is *per se* very challenging. It makes sense to try to analyze it for relatively straightforward situations, such as the Ising model and related approximations. Consider the model in three dimensions; all computations (both numerical and analytical) consistently predict for this model a second-order phase transition. It would be useful to compare its behavior with that of a (possibly very similar) model which is known to undergo a first-order transition. Such a model exists, the so-called Potts model [7]. This is a generalization of the Ising model, in that instead of considering only two possible states for the spin variable, one can think of a spin which may take q possible values. For each value of q a model is defined ($q = 2$ is the now-familiar Ising case). The three-state Potts model is the simplest generalization that can be defined, and the interaction favors configurations where the spins are all in the same state. The model has an internal Z_3 symmetry (permutations of a group of three elements).



(a)



(b)

Figure 1

Display of two configurations from a Monte Carlo simulation for an Ising model (see text). Configuration (a) is a typical ordered, low-temperature state (spins are mostly in the red state), whereas configuration (b) refers to a disordered, high-temperature state, which the system visits later during the simulation. The pictures show that, at the critical temperature, the two phases coexist.

The energy is defined as

$$E = - \sum_{\langle ij \rangle} K \delta_{\sigma_i \sigma_j} \quad (8)$$

Here the spin variable σ_i may be in three different states (one can choose them to be $-1, 0, 1$) and the Kronecker δ function takes value 1 if $\sigma_i = \sigma_j$ and 0 otherwise. Note that the Ising model with energy defined as the product of two spin variables [Equation (3)] can easily be remapped into this one (apart from an overall constant term, which redefines ground-state energy).


There is a general consensus that the three-state Potts model in three dimensions ($q = 3, d = 3$) undergoes a first-order phase transition. Recently a very accurate numerical study [8] definitely confirmed this scenario. Thus it is clear that testing the effectiveness of criteria for distinguishing second-order from first-order phenomena does not require very complicated models. By comparing the Ising model to the three-state Potts model (i.e., comparing $q = 2$ to $q = 3$) one can analyze the two situations without dealing with very different systems.

• Visualization of results


The graphic tool described in Section 1 has precisely the purpose of comparing the dynamics of the two systems. The idea is to show directly the *configurations* of the systems as they evolve during the Monte Carlo simulations, which is a literal way of viewing these

numerical "laboratories." Together with the important analysis of global and average quantities (e.g., the behavior of magnetization or specific heat), we claim that this graphic analysis can greatly improve our understanding of a phase transition.

In this perspective the key feature to investigate is phase coexistence. As previously explained, for a system in which a second-order phase transition takes place, one should see at $T = T_c$ that the ordered and disordered phases coexist; "bubbles" of one phase should grow inside the other. In **Figures 1** and **2** (for the Ising and Potts models, respectively), we show some output from the visualization part of our codes. The pictures are quite self-explanatory; both sets refer to systems of 16^3 degrees of freedom. In other words, we map our system on a cubic grid, the linear size of the cube being 16 grid points. The representation of the value of each degree of freedom is based on colors: two possible values (red and blue) for the up and down states of the Ising model, and three (red, blue, and green) for the Potts model. We show the configurations of the system during a Monte Carlo simulation at $T \sim T_c$. For the Ising model the "snapshots" were taken at $\beta = 1/T = 0.45$; for the Potts model β was 0.551 in Figure 2(a) and 0.553 in Figure 2(b). During the simulation the system is able to jump from one ordered phase (i.e., one color) to the other. This is true for both systems, but a basic difference is that in



(a)



(b)

Figure 2

Display of two configurations from a Monte Carlo simulation for a Potts model (where the spin can exist in three states, coded as red, green, and blue). In this case the two configurations are in the critical region at different temperatures (inverse temperatures are displayed in the pictures). Here the system does not jump from ordered to disordered states, but persists in the ordered phase.

the Potts case this transition takes place as a competition between two ordered states which both belong to the low-temperature region, as shown in Figures 2(a) and 2(b). In the case of the Ising model this is not true; when going from one ordered state to the other, the system spends a considerable amount of time in a disordered, high-temperature phase configuration like the one in Figure 1(b). This is direct evidence that in this case the two phases coexist, as they should for second-order behavior.

The other quantity that one can examine quite instructively in this context is the correlation length of the system. In all figures it is clear that there are regions of coherence (where all points are of the same color). By zooming and cutting sections, the geometry of these coherent regions can be better understood.

- *Quantum chromodynamic model*

Another physical problem that we considered as a test ground for this tool is quantum chromodynamics (QCD), a theory which describes strong interactions of matter. Understanding QCD is one of the most challenging problems of modern theoretical physics. Again, numerical simulations are one of the widely used tools of investigation, and the complexity of the problem is such that very powerful computers, optimized for this class of problem, have been built to perform the simulations [9, 10]. One of the quantities that play a major role in

this framework is the set of particle fields which mediate interactions among quarks, the so-called *gauge fields*.

Quarks are supposed to be, in our universe, permanently confined within the observable particles that constitute matter (protons and neutrons, for example). But at high matter densities (i.e., at high temperatures such as those thought to exist during the first moments of life of the universe, which physicists are now trying to reproduce in high-density plasma), a phase transition could deconfine some quarks. The study of such a phenomenon has crucial importance. One can build from elementary gauge fields an effective system that strongly resembles the three-state Potts model (since quarks can appear as three different types, called colors), and universality concepts can then be applied.

- *Visualization of the complex variable*

We have examined the output configurations of some very long numerical simulations performed by the APE group on the APE computer [9]; some examples are shown in Figures 3–5. The behavior of the system was again analyzed on a cubic lattice by looking at its critical properties. The variable defined on the site is in this case a complex variable (so is the gauge field), so our graphic representation had to deal with two degrees of freedom, a modulus and a phase. As briefly described in Section 1, the representation we chose for a complex variable on a




Figure 3

Three examples taken from a simulation of quantum chromodynamics (QCD—see text). Here the site variable is a complex field whose magnitude varies with position and is proportional to the size of the pixel. The phase of this complex number is color-coded, and it appears from these pictures that site angular variables are concentrated around three fundamental, ‘‘Potts-like’’ phase angles.

site is as follows: Each variable is associated with a block of pixels, the size of the block being proportional to the modulus of the complex number. The phase is described by means of a color scale; phases near 0 degrees are represented as red, those near 60 degrees as green, and those near 120 degrees as blue. In principle, variables whose phase is intermediate between two of these three values could well exist, but from the universality argument one expects that ‘‘Potts-like’’ configurations will

be dynamically selected close to the phase transition. This is indeed confirmed by our visualizations, which show only site variables with phase near the three ‘‘basic’’ angles.

• *Spin glasses*

We conclude this section by considering a third physical system, which is again closely related to the Ising model in which interactions are deterministic and

homogeneous. However, one can conceive some sort of class of "amorphous" models, where interactions strongly depend on space position and translational invariance is explicitly broken. These are known as *spin glasses*. The energy for an Ising spin glass is

$$E = - \sum_{\langle ij \rangle} K_{ij} s_i s_j. \quad (9)$$

The interaction strengths K_{ij} vary from site to site and are randomly extracted from a probability distribution. For example, K_{ij} can be a Gaussian random variable with an average value of 0. Thus *disorder* is explicitly present in the coupling strengths. The real materials which inspired this model are alloys of magnetic with nonmagnetic materials (such as CuMn). The behavior of this model is very complex to disentangle. An approximate solution (the *mean field*) has already shown some very complex and rich properties which have been precisely analyzed by means of an ingenious analytical method (for a review of that subject, see [11]). However, the nature of the phase transition in three dimensions is still an open issue. One source of complexity manifests itself at the very microscopic level. A spin s_i interacts with its neighbors with the interactions K_{ij} . These forces may impose competing constraints; e.g., if two K are both negative, and the two neighbors are in opposite states, the spin s_i cannot "choose" one state without violating one constraint. This property is referred to as *frustration*, for obvious reasons. This local property can be examined by looking at the value of the link energy $K_{ij}s_i s_j$. If this energy is positive, the link is said to be frustrated. A recent set of numerical simulations has benefited greatly from the analysis of "frustration maps" generated during the running of the simulations, such as the one shown in **Figure 6**. From this picture it is clear that frustration is not locally concentrated, but rather widely distributed throughout the system. This rules out some proposed models for $3d$ spin glasses [12] and makes more plausible that many crucial features of the mean-field solution are also preserved for finite-dimensional models [13].

Three-dimensional turbulence

• 3D fluid dynamics

The problems discussed so far are concerned with systems at thermodynamic equilibrium. Our graphic tool, however, can also be profitably employed in the analysis of nonequilibrium systems such as, for example, the motion of a fluid in a three-dimensional space. This is particularly true in view of the recent finding that hydrodynamic phenomena can be simulated by means of discrete lattice methods (lattice gas cellular automata [14]) which bear many technical similarities to the lattice methods adopted in statistical mechanics. Therefore, all




Figure 4

Another QCD visualization from a different perspective view, to permit identification of coherent regions and their dimensionalities.




Figure 5

Zooming of one QCD configuration. This makes it possible to see the different values of the modulus of the complex site variable as the size of the "block pixel" varies from site to site.




Figure 6

Frustration map for a spin glass system (see text).

the simulations of fluid flows presented in this section have been carried out by means of a lattice gas code.

As a matter of fact, even though the basic equations of fluid dynamics were established more than a century ago by Claude Navier and Gabriel Stokes, the rules which govern the overall behavior of fluids are, to a large extent, still poorly understood. The reason is that, in many instances, fluids are *turbulent*, in the sense that they alternate regions of strongly chaotic motion with other regions where the fluid motion is quite smooth and quiescent. This "dual" nature makes hydrodynamic turbulence one of the most enigmatic and challenging phenomena of macroscopic physics.

The Navier-Stokes equations take the form

$$\frac{D\mathbf{u}}{Dt} = \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\frac{\nabla p}{\rho} + \nu \nabla^2 \mathbf{u}, \quad (10)$$

where $\mathbf{u}$ is the velocity field, p is the fluid pressure, and ρ is the fluid density. The physics of hydrodynamic turbulence is ultimately dictated by the competition between nonlinear convective effects [the term $(\mathbf{u} \cdot \nabla)\mathbf{u}$], which couple disparately different spatial scales, and viscous damping (the term $\nu \nabla^2 \mathbf{u}$) associated with molecular dissipation. The relative strength of these two mechanisms, usually expressed by a single dimensionless parameter, the Reynolds number, is commonly accepted as a good measure of the degree of turbulence of a given fluid. The Reynolds number is defined as the ratio

$$N_{\text{Re}} = \frac{UL}{\nu}, \quad (11)$$

where U is a typical speed of the fluid, L a typical length, and ν the molecular viscosity. When the Reynolds number is of order 1, the fluid is said to be "laminar"; i.e., it moves smoothly and uniformly in a fairly predictable fashion. However, for increasing values of N_{Re} , the aforementioned chaotic motions take over and start to dominate the dynamics of the fluid. This is normally understood by formulating fluid turbulence in classical nonlinear field theory, which, as such, gives rise to problems in which the relevant dynamics are spread over an increasingly wider range of mutually interacting modes as the Reynolds number increases. More precisely, at a given Reynolds number N_{Re} , the spectrum of turbulence covers the whole range from the macroscopic scale L up to the dissipative scale $l \sim LN_{\text{Re}}^{-5/3}$, which is the typical scale at which energy starts to be dissipated. A turbulent fluid can be regarded as a collection of elementary excitations whose number grows rapidly with increasing Reynolds number. Some of these elementary excitations reorganize coherently into ordered and smooth aggregates of vorticity that exhibit a longevity well in excess of the decay time one would expect from viscous dissipation. The persistence of these coherent structures, a sort of metastable state of the theory, is intimately related to the existence of invariants of motion, i.e., quantities which are preserved in the course of the evolution. On the contrary, however, where the elementary excitations interact randomly, these invariants break down. Coherence is replaced by chaos, and dissipation becomes the dominant mechanism.

Most flows in nature are fully turbulent, with values of the Reynolds number ranging from $N_{\text{Re}} \sim 10^7$ for the air flow past a car, to 10^{10} in geophysical phenomena, to several orders higher in astrophysical plasmas. From a practical point of view, a major consequence of turbulence (and possibly its most distinctive property) is the great difficulty encountered in trying to formulate reliable predictions of the temporal evolution of a turbulent system in terms of its initial conditions.

This is the basic reason for the tremendous recent impact of supercomputers in advancing fluid-dynamics research. On one side, this impact entails all the aspects related to the development of faster algorithms for vector and parallel architectures. Equally important, however, is the ability to develop a suitable graphic representation of the enormous amount of data produced by three-dimensional fluid simulation. In addition to the most immediate variables such as density, pressure, and velocity, there are other quantities of direct physical relevance. Among these, of particular interest are the vorticity $\hat{\omega}$, the helicity h , and the dissipation D . These

are defined as follows:

$$\begin{aligned}\bar{\omega} &= \nabla \times \bar{u}, \\ h &\equiv \bar{\omega} \cdot \bar{u}, \\ D &\equiv |\nabla \bar{u}|.\end{aligned}\quad (12)$$

Vorticity is especially useful in two-dimensional turbulent flows, as is easily seen once the Navier-Stokes equations are recast in terms of a vorticity-evolution equation:

$$\frac{D\bar{\omega}}{Dt} = \nu \Delta \bar{\omega} - (\bar{\omega} \cdot \nabla) \bar{u}. \quad (13)$$

The first term on the right-hand side represents the viscous dissipation, while the second one is the "vortex-stretching" term which can be either positive (vorticity source) or negative (vorticity sink). In two dimensions, the vortex-stretching term vanishes ($\bar{\omega}$ and $\bar{u}$ are necessarily perpendicular); as a result, in the inviscid case ($\nu \rightarrow 0$), vorticity is locally conserved. This means that $\bar{\omega}$ does not vary in a reference frame moving with the local speed of the fluid. (This property is referred to as "topological" invariance.) The result is that the best graphical representation of developed 2D turbulence is one which shows isovorticity contours; it permits immediate identification of those spatial regions where the coherent structures (vortices) tend to concentrate. In three dimensions, however, the vorticity-stretching term becomes active and vorticity is no longer a topological invariant. This means that the vortex tubes (surfaces of isovorticity) undergo complex topological transformations which may even lead to the development of singularities in the vorticity field. In particular, there is considerable speculation that these singularities might occur as a fractal set [15].

In such a complicated scenario, the primary source of information is flow visualization, which directly implies the strategic value of graphic algorithms that can follow these complex transformations as far as possible through the course of evolution of the fluid (see also the paper by M. Briscolini and P. Santangelo in this issue [16]).

Another quantity of primary interest, possibly playing a very profound role in the physics of three-dimensional turbulence, is fluid helicity [17]. The reasoning, which evokes analogies with the concepts of statistical field theory discussed in the first part of this paper, is as follows.

Starting from the Navier-Stokes equations, one can prove that at zero dissipation there is an uncountable set of fluid flows with the same energy but a different helicity. Hence, the helicity can serve as a label for a degenerate set of flows in much the same way as the magnetization does for the ground state of the Ising Hamiltonian (with the important difference that the set

of degenerate flows is uncountable, as it must be in view of the fact that the inviscid Navier-Stokes equations are invariant under a *continuous* group of transformations). It can be proved that, among the inviscid solutions of the Navier-Stokes equations, there is a special class of solutions for which the vorticity field tends to align with the velocity field, thereby achieving maximal helicity. These solutions are called "Beltrami" flows and are considered the "ordered" states of the theory. In the absence of dissipation, these configurations are topologically "frozen" in the fluid and cannot transform from one into another because they are separated by an infinitely high topological barrier (note the correspondence between dissipation in a fluid and temperature in a magnetic system). However, even an infinitesimal amount of dissipation is sufficient to bring the topological barrier down to a finite height, thus generating a nonzero tunneling probability between distinct topologies. This means that a whole ensemble of flows becomes accessible to the fluid and, by invoking a principle of minimum dissipation, that Beltrami flows should be the most frequently visited configurations in fully developed turbulence ($\nu \rightarrow 0$, $N_{re} \rightarrow \infty$).

In conclusion, one postulates that "Beltramization" of the flow should normally occur where dissipation is minimal or, differently stated, that helicity h and dissipation D should be strongly anticorrelated.

Again, it is apparent that three-dimensional visualization (possibly concurrent) of these two fields can be of crucial value in validating or disclaiming such a speculative scenario.

• Graphic representation

In the previous section we have shown that there are several fields which can be visualized to aid in the comprehension of three-dimensional turbulence. We now present a series of examples which show how our tool can be used in this context.

In general, the field is systematically displayed in the complex "order parameter" format,

$$\Psi = A + iB, \quad (14)$$

where A and B are "dummy" fields whose actual content can be selected case by case according to the specific need. The quantity A normally corresponds to an amplitude and is associated with the size of the pixels, while B is meant to describe the cosine of a real angle and is normally associated with a color scale.

As a first example, we display a series of three pictures (Figures 7, 8, and 9) corresponding to the temporal evolution of two offset orthogonal vortex tubes. In this case the identification is

$$A = |\bar{\omega}|, \quad B = \frac{\omega_M}{|\bar{\omega}|} \equiv \cos \theta_M, \quad (15)$$




Figure 7

Perspective view of the vorticity field at $t = 0$. Vorticity is maximum on the axes of the two red offset orthogonal tubes and decays thereafter with a Gaussian profile.




Figure 8

The configuration of Figure 7 after 600 steps of the Navier-Stokes solver. The axes of the vortex tubes bend under the effect of convective effects, while the cores expand and diffuse away from the axes as a result of viscous dissipation.

where ω_M designates the maximum of the three components of the vector $\vec{\omega}$. The amplitude A is discretized over ten values of the pixel size which go linearly from the minimum to the maximum value of vorticity. The quantity B is designated to distinguish between counterclockwise rotation (positive vorticity, red color) and clockwise rotation (negative vorticity, blue color).

Figure 7 contains 32^3 sites, a number sufficient to show the cores of the two offset orthogonal vortex tubes. According to the physical picture outlined in the previous section, once this configuration is evolved, the nonlinear terms promote a conservative, long-range interaction which would make the tubes attract, stretch in a complicated way, and possibly merge after a sufficient length of time. This is in competition with viscous dissipation, which tends to "broaden" the tubes and spread them over the space, thus ultimately annihilating them. Figures 8 and 9 offer a neat visualization of the aforementioned competition. In fact, we see in Figure 8 that the extremities of the two vortex cores tend to bend under the effects of advection and vortex stretching. At the same time the vortex cores undergo a strong diffusion, as evidenced by the broadening of the red regions. As time proceeds further, viscous dissipation becomes more and more dominant and, as shown in Figure 9, the system tends to lose the memory of its initial well-organized shape. Such a neat dominance by dissipative effects is due to the low resolution of the lattice gas solver, in this case 32^3 grid points corresponding to a Reynolds number of the order of 10 [18].

This kind of visualization can offer a useful complement to the conventional representation of isovorticity surfaces, especially when the interaction tends to generate complicated nesting which is difficult to trace with connectivity-path algorithms. An additional potential advantage is having the whole range of vorticity values available through the size of the pixels. For smooth configurations, this is equivalent to viewing several isosurfaces at the same time (even if less smooth); for chaotic configurations it can help in localizing the regions where vorticity gradients tend to concentrate, a stage at which zooming capabilities are likely to become essential.

Next, we show an example of "hybrid" representation in which the terms A and B refer to two distinct physical quantities. This is particularly useful in highlighting the presence of spatial regions where A and B exhibit correlation or anticorrelation effects. In view of the discussion presented in the previous section, we let

$$A = |\Delta \vec{u}|, \quad B = \frac{h}{|\vec{u} \cdot \vec{\omega}|}, \quad (16)$$

so that A represents the intensity of dissipation while B is a measure of the fluid "Beltramization," i.e., the tendency of the vorticity field to align with the velocity field. The color table has been chosen so as to yield red for positive h ($\vec{\omega}$ parallel to $\vec{u}$), blue for negative (antiparallel), and green for $h = 0$ (perpendicular). As a result, Beltrami structures should appear as red or blue spots in the picture.

Figure 10 shows the same fluid configuration as Figure 7 in the helicity-dissipation representation. From this figure we first notice a large majority of green points, with small fluctuations of red color and practically no blue regions. A tentative conclusion is that the fluid does not exhibit much Beltramization, which is not surprising because dissipation is by far the dominant mechanism. Because of the moderate size of the red pixels which do appear, however, one might speculate that where Beltramization does occur, the dissipation is rather contained. It should be stressed that these considerations are purely qualitative; one color might affect pixel size differently than another, and pixels representing points near the "front" of the lattice might, depending on the drawing order, obscure those representing points near the "back." The former difficulty is alleviated to some extent by allowing the user to swap colors in a simple way, while the latter is considerably mitigated by the ability to rotate the picture in real time. In any case, it is clear that a careful inspection of a whole series of *different* pictures of the *same* data are required before quantitative assessments can be made.

The same representation is also extremely well suited for checking the speculation that dissipation should take place on a highly irregular set of fractal measure. To test this point, Navier-Stokes simulations with at least 256^3 grid points would be required, a resolution that cannot be accommodated within the present memory capability of the 5085 graphics processor. Several alternatives are available, however. For instance, one could select a subset of the computational domain, either by decimation (one graphic node for each four or eight computational nodes) or by spatial averaging, but in order to control the loss of information on the short scales, both of these procedures should be repeated over an ensemble of such subsets.

Conclusions and future plans

We have described both the tool, as it currently exists, and some examples of its utilization. However, even if we consider the tool extremely useful as it is, many improvements and extensions are still possible. One of our primary goals is to achieve a visualization in real time of many-body system dynamics. The approach we are following is essentially based on the exploitation of




Figure 9

The configuration of Figure 7 after 1200 steps. Due to the low Reynolds number, dissipative broadening definitely prevails over convection.




Figure 10

The configuration of Figure 7 in the helicity-dissipation picture. The dissipation is proportional to the size of the pixels, while helicity is represented in a color scale: red for positive (velocity parallel to vorticity), blue for negative (antiparallel), and green for values around zero (perpendicular). Note the dominance of green pixels, which indicates that helicity production is strongly inhibited by dissipation. Note also that most of the red pixels are small, indicating that where helicity is maximal, dissipation tends to be small (Beltrami flow).

cooperative processing in a heterogeneous environment. We have a mainframe (3090) and a number of workstations of different types (RISC System/6000 processors, RT PCs³, and PS/2s³) all connected in a LAN. For the time being we are using token ring as physical support and TCP/IP as the communication protocol. The "cooperative processing" concept can be described as follows: One of the machines just mentioned (typically the 3090 processor) runs a simulation; as soon as a new set of results is produced it is sent through the LAN to another machine running the process that produces and handles the graphic output. A "socket" abstraction is used at the application level. This allows both processes to regard the LAN as being very similar to a standard I/O unit. Moreover, there is no reciprocal dependency, as there would be in an RPC model of interprocess communication; the simulation and the graphics actually run in parallel. In order to hide the details of communication as much as possible from the user, a library is under development which will allow both C and FORTRAN codes to be easily adapted to the distributed environment. Since all of our machines support graPHIGS, we could use, in principle, a PS/2 or an RT PC as workstation for the graphics, but the RISC System/6000 processor offers performance so much higher than the others that, apart from the development phase, it is the workstation we use for the production of graphic output. The presence of the 24-bit Color Graphics Processor on the RISC System/6000 (possibly with the optional 24-bit z-buffer) will allow the addition of new features to the tool, such as the removal of hidden lines and depth cueing (with this facility, the points on a line further from the eye are less bright, so that there is a fading effect with increasing distance). Another key point is the possibility that the user can modify the behavior of the simulation by changing the value of the parameters or suspending and then restarting the simulation itself. In the typical scenario the user works at a RISC System/6000 monitor and has four windows (X-Windows⁴) active on the screen; two windows are the consoles of the workstation and of the machine running the simulation, one contains the output from a graPHIGS manipulation of the data, and the last allows the user to send feedback to the simulation. The results we have obtained in this "new" environment so far, although preliminary, are encouraging as to the quality of the animation we are able to achieve.

In summary, the tool we have developed allows us to extract from numerical simulations much more information than was previously possible, and to present this information in a clearer and more intuitive manner.

The availability of more and more efficient hardware will improve both the quality of the graphic output and the degree of interactivity of the tool.

Acknowledgments

We thank A. Fernandez, F. Lozupone, M. P. Lombardo, F. Piccolo, N. Sourlas, and A. Tarancon for very profitable exchanges. We gratefully acknowledge M. Briscolini for providing us with the initial conditions of the fluid dynamics simulations and for useful discussions concerning visualization in this framework. We also enjoyed interesting discussions with N. Cabibbo, who provided us with some very useful suggestions.

References

1. *Phase Transitions and Critical Phenomena*, C. Domb and M. S. Green, Eds., Academic Press Ltd., London, 1974.
2. E. Ising, *Zeit. Phys.* **21**, 613 (1925).
3. L. Onsager, *Phys. Rev.* **65**, 117 (1944).
4. *Applications of the Monte Carlo Method in Statistical Physics*, K. Binder, Ed., Springer-Verlag, Berlin, 1984.
5. G. Parisi, *Statistical Field Theory*, Addison-Wesley Publishing Co., Redwood City, CA, 1988.
6. M. E. Fisher, in *Critical Phenomena*, M. S. Green, Ed., Academic Press, Inc., New York, 1971.
7. R. B. Potts, *Proc. Camb. Phil. Soc.* **48**, 106 (1952).
8. R. V. Gavai, F. Karsch, and B. Peterson, "A Study of the Correlation Length Near a First Order Phase Transition: The 3D Three-State Potts Model," *Nucl. Phys. B* **322**, 738 (1989).
9. P. Bacioli et al., "The APE Collaboration: Order of the Deconfining Phase Transition in Pure-Gauge QCD," *Phys. Rev. Lett.* **61**, 1545-1548 (1988).
10. F. R. Brown, N. H. Christ, Y. Deng, M. Gao, and T. J. Woch, "Nature of the Deconfining Phase Transition in SU(3) Lattice Gauge Theory," *Phys. Rev. Lett.* **61**, 2058-2061 (1988).
11. M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond*, World Scientific Press, Singapore, 1987.
12. D. A. Fisher and D. A. Huse, "Ordered Phase of Short-Range Ising Spin-Glasses," *Phys. Rev. Lett.* **56**, 1061 (1986).
13. S. Caracciolo, G. Parisi, S. Patarnello, and N. Sourlas, "3D Ising Spin-Glasses in a Magnetic Field and Mean-Field Theory," *Europhys. Lett.* **11**, 783 (1990).
14. U. Frisch, B. Hasslacher, and Y. Pomeau, "Lattice Gas Automata for the Navier-Stokes Equation," *Phys. Rev. Lett.* **56**, 1505-1508 (1986).
15. B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman & Co., San Francisco, 1977.
16. M. Briscolini and P. Santangelo, "Animation of Computer Simulations of Two-Dimensional Turbulence and Three-Dimensional Flows," *IBM J. Res. Develop.* **35**, 119-139 (1991, this issue).
17. E. Levich and E. Tzvetkov, "Certain Problems in the Theory of Developed Hydrodynamics," *Phys. Rep.* **128**, 1 (1987).
18. F. Higuera, S. Succi, and R. Benzi, "Lattice Gas Dynamics with Enhanced Collisions," *Europhys. Lett.* **8**, 517 (1989).

Received November 19, 1989; accepted for publication November 5, 1990

³ RT PC and PS/2 are registered trademarks of International Business Machines Corporation.

⁴ X-Windows is a trademark of MIT.

Massimo Bernaschi *IBM European Center for Scientific and Engineering Computing (ECSEC), Via Giorgione 159, 00147 Rome, Italy.* Dr. Bernaschi received a degree in physics from the University of Rome "Tor Vergata." He subsequently worked in a project of the National Institute for Nuclear Physics (called APE collaboration) that built a parallel computer for lattice gauge theory. In 1987 he joined IBM ECSEC. Currently his main interests are in the areas of statistical mechanics and the use of cooperative and distributed processing concepts for scientific and engineering computing.

Enzo Marinari *University of Rome "Tor Vergata," Via E. Carnevale, Rome, Italy.* Professor Marinari received his doctorate in physics in 1980 from the University of Rome "La Sapienza." From 1982 to 1984, he was a member of the Service de Physique Théorique, CEN-Saclay, France. He currently holds the position of Professor at the University of Rome "Tor Vergata," where he teaches theoretical physics. Professor Marinari has held assignments at the École Normale Supérieure, CERN, Brookhaven National Laboratories, Edinburgh University, the Jerusalem University, and the Hamburg Imperial College, as well as several other research posts in theoretical physics. In 1988 he received the Fondazione Dr. Giuseppe Borgia prize from the Accademia dei Lincei (Classe Scienze Fisiche); this prize recognizes outstanding work by a physicist who is not yet 35 years of age.

Stefano Patarnello *IBM European Center for Scientific and Engineering Computing (ECSEC), Via Giorgione 159, 00147 Rome, Italy.* Dr. Patarnello received his degree in physics in 1983 from University of Pisa, Italy. He joined IBM in 1984 as a member of European Center for Scientific and Engineering Computing in Rome. His main research interests have been in statistical physics, where he has been engaged mainly in the study of Monte Carlo methods for simulation of critical phenomena, stochastic optimization techniques, such as the use of simulated annealing in engineering applications, and neural networks, where he has primarily investigated new computational models and their application in many areas of information science. During his career, Dr. Patarnello has also worked in numerical methods for resolution of fluid dynamics equations in two dimensions and in the simulation and performance evaluation of architectures for supercomputing.

Sauro Succi *IBM European Center for Scientific and Engineering Computing (ECSEC), Via Giorgione 159, 00147 Rome, Italy.* Dr. Succi received his doctorate in nuclear engineering from the University of Bologna in 1979 and a Ph.D. in plasma physics from the Swiss Federal Institute of Technology in 1987. In 1980 he was at the Department of Fast Breeder Reactors of the National Committee of Nuclear Energy (CNEN) in Bologna, where he developed algorithms for the optimal control of nuclear accidents. Subsequently, he joined the Neutral Injection Team of the Max Planck Institut für Plasmaphysik in Garching (FRG), working on Monte Carlo simulation codes for the heating of thermonuclear plasmas. In 1982 he moved to the Plasma Physics Center of the Swiss Federal Institute of Technology in Lausanne, where he was in charge of the utilization and development of finite element codes for the study of plasma magnetohydrodynamics and plasma kinetics. In 1986 Dr. Succi joined IBM ECSEC, where he is currently involved in projects of theoretical and industrial computational fluid dynamics. He has authored numerous papers in the fields of kinetic theory, plasma physics, and computational fluid dynamics.