

Thermodynamically reversible generalization of diffusion limited aggregation

Raissa M. D'Souza

Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Norman H. Margolus

Center for Computational Science, Boston University, Boston, Massachusetts 02215

and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

(Received 7 December 1998)

We introduce a lattice gas model of cluster growth via the diffusive aggregation of particles in a closed system obeying a local, deterministic, microscopically reversible dynamics. This model roughly corresponds to placing the irreversible diffusion limited aggregation model (DLA) in contact with a heat bath. Particles release latent heat when aggregating, while singly connected cluster members can absorb heat and evaporate. The heat bath is initially empty, hence we observe the flow of entropy from the aggregating gas of particles into the heat bath, which is being populated by diffusing heat tokens. Before the population of the heat bath stabilizes, the cluster morphology (quantified by the fractal dimension) is similar to a standard DLA cluster. The cluster then gradually anneals, becoming more tenuous, until reaching configurational equilibrium when the cluster morphology resembles a quenched branched random polymer. As the microscopic dynamics is invertible, we can reverse the evolution, observe the inverse flow of heat and entropy, and recover the initial condition. This simple system provides an explicit example of how macroscopic dissipation and self-organization can result from an underlying microscopically reversible dynamics. We present a detailed description of the dynamics for the model, discuss the macroscopic limit, and give predictions for the equilibrium particle densities obtained in the mean field limit. Empirical results for the growth are then presented, including the observed equilibrium particle densities, the temperature of the system, the fractal dimension of the growth clusters, scaling behavior, finite size effects, and the approach to equilibrium. We pay particular attention to the temporal behavior of the growth process and show that the relaxation to the maximum entropy state is initially a rapid nonequilibrium process, then subsequently it is a quasistatic process with a well defined temperature.

[S1063-651X(99)11007-9]

PACS number(s): 05.70.Ln, 61.43.Hv, 68.70.+w

I. MICROSCOPIC REVERSIBILITY AND PATTERN FORMATION

Pattern formation is an intrinsically dissipative process [1], however, the laws of physics are microscopically reversible: there is no dissipation at the microscopic scale. In this paper we describe a simple system which organizes into patterns through microscopically reversible dynamics, hence it also models how dissipation arises (i.e., how information flows between the macroscopic and the microscopic degrees of freedom). This system provides a clear example of how to reconcile the macroscopic irreversibility that gives rise to patterns with the microscopic reversibility adhered to by physical processes. Motivated by the desire to understand this general issue, we study specific details of the model, focusing on transitions in the resulting growth morphology and the approach to thermodynamic equilibrium.

We have previously observed several examples of reversible cellular automata dynamics which produce large scale order through microscopically reversible dynamics [2–4]. In contrast other research in the field of pattern formation has focused on irreversible microscopic mechanisms, with examples ranging from crystal growth [5,6], to Turing patterns in chemical reactions [7], to patterns formed by growing bacterial colonies [8], to kinetic growth problems [9,10].

Here we model cluster growth by reversible aggregation (RA), generalizing the irreversible diffusion limited aggregation model (DLA) [10] to include contact with a heat bath.

Particles, which are initially diffusing on a two-dimensional lattice, stick upon first contacting a cluster member and release heat which then diffuses about a superimposed lattice representing the heat bath. The two subsystems exchange only heat and together form a closed system. The release of heat transfers entropy from the aggregating system (which is becoming ordered) into the heat bath (which was initially empty). When the heat bath is nearly empty the model is essentially equivalent to the canonical DLA formulation (analogous to a supercooled gas crystallizing in a far from equilibrium situation). Hence the RA growth cluster initially resembles a typical DLA cluster. As the heat bath becomes populated, singly connected cluster members are able to absorb heat and evaporate. As the effect of evaporation becomes significant the RA and DLA models diverge. The RA dynamics is exactly invertible: at any point we can invert the dynamics and run backwards, observing the flow of heat from the heat bath back into the gas-crystal system until we recover the exact initial condition.

The population levels of the heat bath and of the aggregate initially grow linearly in time, quickly reaching stable values which remain very nearly constant for the remainder of the evolution. The energy of each subsystem is a function only of the population levels, independent of the physical configuration of the particles. Hence, once the population levels stabilize, the rate of energy exchange (which is entirely in the form of heat) between the heat bath and the

gas-aggregate system becomes so slow that we can characterize the subsequent dynamics as a quasistatic process, with a well defined temperature at all times.

The aggregate mostly forms while the heat bath is at a lower temperature than in the quasistatic steady state. Hence, after the population levels stabilize, the cluster slowly anneals. The cluster morphology, which initially resembles a DLA cluster, gradually becomes more spread out and tortuous, until it ultimately resembles a branched polymer with quenched randomness. The two timescales that characterize the growth process are separated by two orders of magnitude. Initially, the population levels quickly reach a quasistatic steady state. Subsequently, the aggregate slowly anneals until reaching the ensemble of configurations corresponding to the highest entropy macrostate (the branched polymer).

Aside from insight into microscopically reversible mechanisms that give rise to macroscopic patterns, the development of invertible dynamics and algorithms has technological significance in pushing down the barrier to atomic scale computing. Each bit of information erased at temperature T releases at least $T\Delta S = k_B T \ln 2$ units of heat into the environment [11]. Heat is created in proportion to the volume of the computer, yet heat leaves the computer only in proportion to the surface area. Hence, as logic gate density in computers increases, the use of an invertible dynamics (which does not erase information and hence does not need to produce heat) will be required to keep the mechanical parts from burning up [12,13]. From a more pedagogical viewpoint, discrete computer models of reversible microscopic dynamics provide a laboratory for studying nonequilibrium statistical mechanics and the approach to equilibrium. These models let us explore physically plausible dynamics for nonequilibrium systems (i.e., discrete dynamics which are microscopically reversible and thus automatically obey Liouville's theorem). A particularly instructive example of this approach is the formulation of a dynamical Ising model [14–16]. However, more widely used in physics are discrete, reversible models of fluid flow such as the hpp and fhp lattice gases [17,18]. For a recent discussion of modeling physical phenomena with reversible computer models see Ref. [3]. For a recent discussion of macroscopic irreversibility and microscopic reversibility see Ref. [19]. For a recent discussion of techniques for the explicit construction of reversible models in statistical mechanics see Ref. [20]; but note that closely related techniques were discussed in the early 1980's by Fredkin (as discussed in Ref. [21]).

The initial sections of this manuscript describe our model, the middle contain a mathematical formulation of the model, and the final, the empirical results. Specifically, Sec. II describes the detailed dynamics, including the subtleties of constructing an invertible dynamical model and implementation issues. In Sec. III we discuss the macroscopic limit of an analytic formulation of the model and establish the reaction-diffusion equations describing the system. In Sec. IV we treat the reaction-diffusion equations in the mean field limit and compare predictions for equilibrium densities of particles to empirical measurements. Empirical measurements of temperature are presented in Sec. V A, with emphasis on the quasistatic nature of the annealing portion of the growth process. In Sec. V B we study the evolution of the fractal dimension of the clusters and thus quantify the change in

growth morphology as the clusters relax to the maximum entropy state. We conclude with a discussion of limitations and possible modifications of our model.

II. MODELING AGGREGATION

A. Diffusion limited aggregation

Diffusion limited aggregation [10] is a conceptually simple model which serves as a paradigm for some aspects of kinetic growth phenomena. Several comprehensive reviews of DLA have been written. In particular see Ref. [22] for a clear presentation of the basics, Ref. [23] for details on physical mechanisms, and Ref. [24] for a review of real-space renormalization group approaches to DLA.

The typical scenario for DLA begins with a vacant two-dimensional lattice initialized with a single stationary seed particle, which is the nucleation site for a growth cluster. Moving particles are introduced from the edges of the lattice, following random walks along the lattice sites. When a moving particle lands on a site adjacent to a stationary seed particle (an active site) it sticks (i.e., the moving particle aggregates and becomes a stationary seed particle). The frozen aggregate particles constitute the solid (crystal) phase, and the moving particles constitute the gas phase. Aggregation hence consists of a particle undergoing an irreversible transition from gas to solid. Gas particles are usually introduced in a serial manner: only one gas particle is diffusing at a time. However, to take advantage of parallel computational resources, parallel models of DLA have been studied in which multiple particles are diffusing at once [25,26]. In the dilute particle limit, before a substantial fraction of the particles have aggregated, these models recover the serial DLA model exactly.

With the first aggregation event the DLA cluster grows from one single to two adjacent sites. The presence of the second cluster member eliminates certain paths along which random walkers could approach the first, with the effect that the probability of sticking at either end (tip) of the cluster is enhanced, whereas the probability of sticking along the edge of the cluster is reduced. As particles continue to aggregate creating new cluster tips and edges, the probability to stick at the tips continually outweighs the probability to stick along the edges. This leads to branching. A second influence on the growth morphology comes from shadowing: the probability for a particle to diffuse into the center of the growth cluster before encountering an active site becomes negligible as the cluster grows in size. Hence the outer tips grow most rapidly. As a result the growth aggregate rapidly assumes a bushy and branching, random fractal structure, resembling frost on a window pane, the branching of neurons, and many other branched structures found in nature.

B. Reversible aggregation

Our goal is to introduce a reversible, deterministic model of growth by aggregation, where reversible means that from any state of our system we can recover the previous state exactly. We must address the subtleties of making each component step invertible, including steps which realize stochastic processes. As we discuss below, the same mechanisms that are employed in our model in order to ensure exact conservation of energy, particle number, and other con-

straints, also make it easy to incorporate invertibility. The stochastic component of the model is diffusion, which is modeled as a sequence of invertible “random walks” based on a deterministic algorithm using an invertible pseudorandom number generator.

Information about exactly when and where a particle undergoes a phase transition is stored in the heat particles. The idea of storing information in a heat bath was introduced by Creutz [14] to explore the connection between the microcanonical and canonical ensembles for the dynamical Ising model. Heat bath techniques have been used on occasion since then to construct reversible computer models of physical phenomena [27].

1. Overview of the model

To construct the reversible aggregation model we begin with a parallel DLA system (similar to that described above) and add degrees of freedom at each site, corresponding to a distributed heat bath. The latent heat released during each aggregation event can then be explicitly represented. In the RA dynamics, whenever a random walking gas particle lands on a site with exactly one nearest neighbor crystal particle, it will stick only if there is room locally in the heat bath to accept the latent heat it will release as the particle transitions from the gas phase to the crystal phase. The heat is released in quantized units called heat particles, with one heat particle released for each aggregation event. These heat particles diffuse amongst themselves (i.e., they undergo random walks along the lattice sites, independently of the gas particles). Explicitly modeling the latent heat released upon aggregating provides a mechanism for modeling the inverse process: a diffusing heat particle which contacts a susceptible crystal particle (a crystal particle which has only one nearest neighbor crystal particle) is absorbed while the crystal particle evaporates to become a gas particle which then diffuses away.

The restriction on the dynamics that aggregation and evaporation events can occur only when exactly one nearest neighbor is a crystal particle means that only one crystal bond is ever formed or broken when a single lattice site is updated. As each aggregate particle contributes one crystal bond to the aggregate, and there is no further potential energy contribution, the energy of the aggregate is a function only of the number of aggregate particles, independent of their configuration. Moreover, this constraint has two direct implications for the growth morphology. The first is that evaporation can only occur for particles which are singly connected to the growth cluster, and so the aggregate cannot break off into disconnected clusters. The second is that it introduces an excluded volume (i.e., no closed loops can be formed), thus we might expect the equilibrium cluster configuration to be similar to that of a polymer. Note that introduction of an evaporation mechanism in the RA model mitigates the shadowing effect that was important in determining the DLA growth morphology: crystal particles within the cluster can evaporate, thus introducing gas particles into the interior of the aggregate.

2. The detailed dynamics

The RA model is constructed with seven bits of state at each lattice site. One bit, $N_c(\vec{x}, t)$, denotes the presence or

absence of a crystal particle at that site [i.e., $N_c(\vec{x}, t) = 1$ indicates presence, $N_c(\vec{x}, t) = 0$ absence]. Two bits, $N_g^\gamma(\vec{x}, t)$ where $\gamma = \{1, 2\}$, denote the presence or absence of each of two gas particles. Two bits, $N_h^\gamma(\vec{x}, t)$, denote the presence or absence of each of two heat particles. The final two bits, $\xi_g(\vec{x}, t), \xi_h(\vec{x}, t)$, are independent binary pseudorandom variables. The dynamics of the model consists of two kinds of steps: diffusion steps alternating with interaction steps.

The same kind of diffusion process is applied to the gas and heat subsystems simultaneously and independently, while the crystal particles remain unchanged. A given diffusion step consists of two parts: mixing and transport. During the mixing portion of the step, a binary random variable is used to determine whether or not the two particle bits of that species at the site (\vec{x}, t) are interchanged:

$$\begin{aligned} N_i^1 &= (1 - \xi_i)N_i^1 + \xi_i N_i^2, \\ N_i^2 &= (1 - \xi_i)N_i^2 + \xi_i N_i^1, \end{aligned} \quad (1)$$

where $i = g$ or $i = h$. During the transport portion of the step, every site replaces its first particle bit ($\gamma = 1$) with the first particle of its neighbor a distance k away on one side, and its second particle ($\gamma = 2$) with the one from the same distance neighbor on the opposite side. At even time steps, we use horizontal neighbors (i.e., the diffusion moves particles horizontally):

$$\begin{aligned} N_i^1(\vec{x}, t+1) &= N_i^1(\vec{x} + k\hat{x}, t), \\ N_i^2(\vec{x}, t+1) &= N_i^2(\vec{x} - k\hat{x}, t). \end{aligned} \quad (2)$$

At odd time steps we use vertical neighbors [i.e., substitute \hat{y} for \hat{x} in Eq. (2)]. The only differences between the gas and heat diffusion are (1) each uses a separate binary random variable to control its mixing and (2) the distance of the neighbor particle to be copied k can be chosen separately for each subsystem—this allows us to independently control the diffusion constants for the heat bath and for the gas (see Ref. [28]).

Diffusion steps alternate with steps in which the two subsystems—gas-crystal and heat bath—interact allowing aggregation and evaporation. The rule at a single lattice site during an interaction step is that exactly one particle can aggregate or evaporate provided that (a) there is exactly one crystal particle at one of the four nearest neighbor sites, (b) there is room at the site for a crystal particle (for aggregation) or for another gas particle (for evaporation), and (c) the heat bits at the site can absorb (for aggregation) or supply (for evaporation) one unit of heat. Since the gas and heat particles will undergo a mixing step before transport, it makes no difference which of the two available gas particle positions a crystal particle is moved into when it evaporates, or which of the two possible heat particle positions a unit of heat gets put into. Defining this precisely will, however, become important when we discuss invertibility.

The interaction rule described thus far would be sufficient if we updated just one lattice site at a time. If, however, all sites on the lattice are updated simultaneously, then the global dynamics no longer obeys the “single bond”

TABLE I. The various phases of one step of the RA dynamics. Each phase is applied over the entire lattice simultaneously.

(1) Interact gas/heat/crystal at even sites
(2) Mix gas and mix heat separately
(3) Transport gas and heat horizontally
(4) Interact gas/heat/crystal at odd sites
(5) Mix gas and mix heat separately
(6) Transport gas and heat vertically

constraint—that at any site where particles aggregate or evaporate exactly one crystal bond is formed or broken. For example, suppose that the tip of a crystal branch evaporates just as a gas particle condenses next to it. Each of these events would separately preserve the constraint, but the two simultaneous events result in the addition of a disconnected crystal particle which has no other crystal particle immediately adjacent to it. We can easily avoid this difficulty by holding the values at the adjacent sites fixed during a step in which we let the subsystems interact at a given lattice site, since the interaction step has a nearest neighbor range. In other words, we perform a *checkerboard updating*: all of the lattice sites in which the x and y lattice coordinates add up to an even number are updated using our single site interaction rule, while the odd sites are held fixed, and then vice versa. Since nearest neighbors are held fixed during an interaction, the constraint is obeyed locally and thus it is also obeyed globally. The overall dynamical rule is summarized in Table I—the various phases of the rule are applied consecutively.

Every phase of the rule described in Table I can be inverted. The transport portion of the step can be run backwards by simply moving all particles back into the sites they came from [i.e., inverting the directions of the transport in Eq. (2)]. The mixing portion of the step is easy to invert, given the same “random” binary variables that were used to determine which pairs of bits were originally swapped. We simply swap exactly those pairs again. The pseudorandom portion of the system (which supplies the random data) can simply follow some invertible dynamics that is independent of the rest of the system—the rest of the system looks at the state of this subsystem but does not affect it—so this pseudorandom subsystem can be run *backwards* independently of everything else.

Making the interaction steps invertible is also straightforward. When a single gas particle turns into a crystal particle, we put the heat token that is released into the corresponding heat bit (i.e., the heat particle with the same value of γ), and thus we remember which of the two gas particle bits was initially occupied. If the corresponding heat bit is already occupied, the particle is not allowed to aggregate (even if the other heat bit is unoccupied). Similarly, a crystal particle is allowed to evaporate only if it can move into the gas bit with the same value of γ as the heat token being absorbed. If there are two gas particles at a site we impose the constraint that the particle with $\gamma=1$ attempts to aggregate first; likewise if there are two heat particles at a site, the crystal particle attempts to absorb the $\gamma=1$ particle first. This does not introduce a bias to the growth since we are mixing the $\gamma=1$ and $\gamma=2$ variables in an unbiased manner at each time step. With these refinements, our interaction rule *applied to a*

single site is its own inverse: if we apply it twice (without a diffusion step in between) we get back the state we started with. Since the interaction rule is applied in a checkerboard fashion, sites are updated independently: if we apply the rule a second time to the same checkerboard, it will undo the first application at every site.

Thus an inverse step consists of applying the inverses of the rule phases in the opposite of the order listed—once one phase is undone, the previous phase can be undone. Each inverse step undoes one step of the forward time evolution. As we watch the inverse evolution, we see each heat particle retrace its path, to be in exactly the right location at the right time to uncrystallize the crystal particle which originally released it. Particles unaggregate and undiffuse and unevaporate in a manner that exactly retraces their behavior in the forward evolution.

3. Implementation

The RA model was implemented on a special purpose cellular automata machine, the CAM-8 [29], which was designed to efficiently perform large-scale uniform, spatially arrayed computations. On this machine, all simulations must be embedded into a lattice gas framework [17,18], in which uniform data movement (data-advection) alternates with processing each site independently (site update). For a two-dimensional model such as ours, sheets of bits move coherently during the advection phase: corresponding bits at each site all move in the same direction by the same amount. The boundaries are periodic—bits that shift past the edge of the lattice reappear at the opposite edge. After moving the bits, we perform the site update phase. During this phase, the bits that have landed at each lattice site are updated in a single operation by table lookup: the bits at each lattice site are used as an index into a table that contains a complete listing of which state should replace each possible original state. Both the data movement and the lookup table can be freely changed between one lattice updating step and the next.

Our model requires seven bits of state to appear at each site in our $L \times L$ lattice. Using random data generated by a serial computer, the bits which correspond to the gas particles are initially randomly filled with a 4% density of particles and the bits which correspond to the binary random variables with a 50% density of particles. One crystal particle is placed at the center site of the lattice. The heat bath is initially empty.

The dynamics on the pseudorandom subsystem is very simple: each of the two random bit planes (each consisting of all the η_g 's or all of the η_h 's) are simply shifted by some large amount at each time step. We could choose the amount and direction of each shift at random for each step of updating, using a reversible random number generator running on the workstation controlling the simulation. Instead, the simulations discussed here simply shift the bit planes by a large and fixed amount at each step, making sure that the x and y components of the two shifts are all mutually coprime, as well as being coprime with the overall dimensions of the lattice. Thus to run the random subsystem backwards, we just reverse the direction of the shifts.

The checkerboard updating is accommodated by adding an eighth bit to each lattice site, and filling these bits with a checkerboard pattern of ones and zeros. In our rule the vari-

ous subsystems are allowed to interact only at sites marked with a one. To change which checkerboard is marked for updating, we simply shift the checkerboard-marker bitplane by one position in the $+x$ direction.

The rule described in Table I turns into two lattice-gas steps on CAM-8. The first three phases listed in Table I are done during one step, and the next three in the second step. The data movement is part of each step: note that each of the two steps uses the same lookup table applied to each lattice site, but slightly different data movement. To run backwards, we use the inverse lookup table, and the inverse data movement. Note that in the discussion of experimental results, everything in Table I is counted as a single step.

CAM-8's event counting hardware was used to monitor simulation parameters while the simulations ran. Including event counting, the eight-processor CAM-8 performed about 10^8 site update operations per second for this model.

III. THE MACROSCOPIC LIMIT

The dynamics of the RA model, described in detail above, can be succinctly presented in an analytic framework. We develop this framework first in terms of the discrete space, time, and occupation number variables. We then ensemble average over the occupation numbers and take the continuum limit of the space and time variables to establish the reaction-diffusion equations for the system.

As discussed in Sec. II B 2 there are seven bits of state: $N_c(\vec{x}, t)$, $N_g^\gamma(\vec{x}, t)$, $N_h^\gamma(\vec{x}, t)$, $\xi_g(x, t)$, and $\xi_h(x, t)$, where $\gamma = \{1, 2\}$. They correspond respectively to one bit of crystal, two bits of gas particles, two bits of heat particles, and two bits of random data. $N_i^\gamma(\vec{x}, t) = 1$ indicates the presence of species i at location \vec{x} and time t , in channel γ , and $N_i^\gamma(\vec{x}, t) = 0$ indicates the absence. The absence or presence of a crystal particle is denoted by $N_c(\vec{x}, t) = \{0, 1\}$, respectively. The total number of particles of species i at time t present on the lattice is denoted by $\mathcal{N}_i(t) = \sum_{\vec{x}, \gamma} N_i^\gamma(\vec{x}, t)$, where the sum is over all of the lattice sites and the two particle channels.

There are no external sources or sinks for any of the three species represented (the gas, crystal, and heat species). Energy is only exchanged between the gas crystal and the heat bath subsystems. Thus the complete system is isolated. Conservation of the total number of gas and crystal particles implies that

$$\mathcal{N}_g(t) + \mathcal{N}_c(t) = \mathcal{N}_g(0) + \mathcal{N}_c(0). \quad (3)$$

Conservation of the total energy of the system implies that

$$\begin{aligned} \mathcal{N}_g(t)\varepsilon_g + \mathcal{N}_c(t)\varepsilon_c + \mathcal{N}_h(t)\varepsilon_h \\ = \mathcal{N}_g(0)\varepsilon_g + \mathcal{N}_c(0)\varepsilon_c + \mathcal{N}_h(0)\varepsilon_h, \end{aligned} \quad (4)$$

where ε_i represents the energy (kinetic and potential) per particle of species i (notice that there is no configurational contribution to the energy of the crystal). As discussed in Sec. II each aggregation event releases one heat particle (likewise each evaporation event absorbs one heat particle), thus $\varepsilon_h = \varepsilon_g - \varepsilon_c$ and moreover

$$\mathcal{N}_h(t) = \mathcal{N}_c(t) - \mathcal{N}_c(0). \quad (5)$$

[Note that $\mathcal{N}_h(0) = 0$ and $\mathcal{N}_c(0) = 1$ in our experiments.]

To facilitate the description of the dynamics, we introduce a functional equation which is $+1$ at any site where a particle is about to crystallize, -1 at a site where a particle is about to evaporate, and 0 otherwise. The functional $\mathcal{F}^\gamma[N_g^\gamma(\vec{x}, t), N_h^\gamma(\vec{x}, t), N_c(\vec{x}, t), \{N_c(\vec{x} + \hat{e}_k, t)\}]$ is evaluated on a neighborhood of lattice sites surrounding some given position \vec{x} at a given time t [the notation $\{N_c(\vec{x} + \hat{e}_k, t)\}$ refers to the set of values of N_c for the nearest neighbors of the point \vec{x}]:

$$\begin{aligned} \mathcal{F}^\gamma(\vec{x}, t) &= N_g^\gamma(\vec{x}, t)[1 - N_h^\gamma(\vec{x}, t)][1 - N_c(\vec{x}, t)] \\ &\times \sum_{j=1}^d N_c(\vec{x} + \hat{e}_j, t) \prod_{k \neq j} [1 - N_c(\vec{x} + \hat{e}_k, t)] \\ &- [1 - N_g^\gamma(\vec{x}, t)] N_h^\gamma(\vec{x}, t) N_c(\vec{x}, t) \sum_{j=1}^d N_c(\vec{x} + \hat{e}_j, t) \\ &\times \prod_{k \neq j} [1 - N_c(\vec{x} + \hat{e}_k, t)]. \end{aligned} \quad (6)$$

Here \hat{e}_j and \hat{e}_k are the vector lattice directions of the nearest neighbor cells, and d the number of distinct lattice directions. For a two-dimensional square lattice (i.e., the lattice used for the present implementation) $d=4$ and the vector lattice directions are $\{\hat{x}, -\hat{x}, \hat{y}, -\hat{y}\}$.

The first term in Eq. (6) equals 1 if a gas particle in channel γ is present at site \vec{x} and time t , a heat particle in channel γ is absent at site \vec{x} and time t , there is no crystal particle already at that site, and only one crystal particle is present at a nearest neighbor site. It is zero otherwise. The second term equals 1 if there is no gas particle in channel γ present at site \vec{x} and time t , there is a heat particle in channel γ present at site \vec{x} and time t , there is a crystal particle present at that site, and only one crystal particle is present at a nearest neighbor site. It is zero otherwise. The first and second terms are mutually exclusive (a heat particle in channel γ cannot be simultaneously present and absent, nor can a gas particle).

The dynamics consists of making the changes indicated by \mathcal{F}^1 and then \mathcal{F}^2 , then applying a random permutation to mix $\gamma=1$ and $\gamma=2$, and then performing the streaming step to move the particles. The permutation mixes the $N_i^1(\vec{x}, t)$ and $N_i^2(\vec{x}, t)$ components in an unbiased way, so it is simpler to discuss the dynamics of a combined variable $N_i(\vec{x}, t) = N_i^1(\vec{x}, t) + N_i^2(\vec{x}, t)$. Likewise, if we let $\mathcal{F}(\vec{x}, t) = [\mathcal{F}^1(\vec{x}, t) + \mathcal{F}^2(\vec{x}, t)][1 - \mathcal{F}^1(\vec{x}, t)\mathcal{F}^2(\vec{x}, t)/2]$, the interaction portion of the dynamics at a single lattice site can be written:

$$\begin{aligned} N_c(\vec{x}, t+1) &= N_c(\vec{x}, t) + \mathcal{F}(\vec{x}, t), \\ N_g(\vec{x}, t+1) &= N_g(\vec{x}, t) - \mathcal{F}(\vec{x}, t), \\ N_h(\vec{x}, t+1) &= N_h(\vec{x}, t) + \mathcal{F}(\vec{x}, t). \end{aligned} \quad (7)$$

The $[1 - \mathcal{F}^1(\vec{x}, t)\mathcal{F}^2(\vec{x}, t)/2]$ factor in the definition of $\mathcal{F}(\vec{x}, t)$ reflects the fact that only one transition occurs at a given site, even if two gas or two heat particles are present.

To construct continuous variables from the discrete ones, we consider the average short-term behavior of the system over an ensemble of independent realizations which all have the same set of local particle densities. For each discrete variable, we let $n_i(\vec{x}, t) \equiv \langle N_i(\vec{x}, t) \rangle$ and denote the average of the functional as $\langle \mathcal{F} \rangle$. (This technique of averaging over many independent realizations, i.e., establishing the one particle density function, is commonly used to derive the lattice Boltzmann equation starting from discrete particle models of hydrodynamics [30,31].)

With this notation, the average propagation of the gas and heat particles can be expressed as

$$n_i(\vec{x}, t+1) = \frac{1}{4} [n_i(\vec{x} - \hat{x}, t) + n_i(\vec{x} - \hat{y}, t) + n_i(\vec{x} + \hat{x}, t) + n_i(\vec{x} + \hat{y}, t)], \quad (8)$$

To establish the continuum limit, we Taylor expand. The terms involving the first derivatives cancel, leading to the result

$$\begin{aligned} n_i(\vec{x}, t+1) &= n_i(\vec{x}, t) + \sum_j \frac{|\Delta x|^2}{4} \frac{\partial^2}{\partial x_j^2} [n_i(\vec{x}, t)] + O(\Delta x^3) \\ &= n_i(\vec{x}, t) + \frac{|\Delta x|^2}{4} \nabla^2 n_i(\vec{x}, t), \end{aligned} \quad (9)$$

where $i = g$ or $i = h$. Note that to order Δt the above equation is the standard diffusion equation

$$\frac{\partial}{\partial t} n_i(\vec{x}, t) = \frac{|\Delta x|^2}{4\Delta t} \nabla^2 n_i(\vec{x}, t). \quad (10)$$

As discussed in Sec. II B 2, we can control the length of each diffusion step separately for the heat and for the gas particles. The heat particles execute random walks composed of individual steps of length k , whereas the gas particles execute walks of step length unity. Thus if the $|\Delta x|^2$ that appears in Eq. (9) and in Eq. (10) refers to the gas subsystem, then $|\Delta x|^2$ for the heat subsystem (and hence its diffusion constant) is a factor of k^2 larger.

To proceed further, we will make the approximation that our average variables are independent. Then we are allowed to replace the average of a product of variables by the product of the average for each variable: $\langle ab \rangle = \langle a \rangle \langle b \rangle$, for a, b independent variables. This is the assumption of molecular chaos, which is also used to derive the lattice Boltzmann equation. With this approximation the average of the functional \mathcal{F}^γ is simply

$$\begin{aligned} \langle \mathcal{F}^\gamma \rangle &= \{ n_g^\gamma(\vec{x}, t) [1 - n_h^\gamma(\vec{x}, t)] [1 - n_c(\vec{x}, t)] \\ &\quad - [1 - n_g^\gamma(\vec{x}, t)] n_h^\gamma(\vec{x}, t) n_c(\vec{x}, t) \} \sum_j n_c(\vec{x} + \hat{e}_j, t) \\ &\quad \times \prod_{k \neq j} [1 - n_c(\vec{x} + \hat{e}_k, t)]. \end{aligned} \quad (11)$$

Similarly, we can write down an expression for $\langle \mathcal{F} \rangle$.

To obtain the continuum limit of these averaged equations, we again use Taylor series approximations. In the diffusive regime, $\Delta t \sim (\Delta l)^2$, so we truncate the expansions at these appropriate orders. Let $\tilde{\mathcal{F}}$ be the continuum limit of $\langle \mathcal{F} \rangle$ (which we will not write out explicitly). Then from Eq. (7), we obtain

$$\frac{\partial}{\partial t} n_c(\vec{x}, t) = \frac{1}{\Delta t} \tilde{\mathcal{F}}. \quad (12)$$

The other reaction-diffusion equations for our system can be obtained by proceeding as we did in Eq. (9). For example, under the full dynamics (which consists of both the interaction and diffusion phases),

$$\begin{aligned} n_g(\vec{x}, t+1) + n_c(\vec{x}, t+1) &= n_c(\vec{x}, t) + \frac{1}{4} [n_g(\vec{x} - \hat{x}, t) \\ &\quad + n_g(\vec{x} - \hat{y}, t) + n_g(\vec{x} + \hat{x}, t) \\ &\quad + n_g(\vec{x} + \hat{y}, t)], \end{aligned} \quad (13)$$

since any particles present at a site at time $t+1$ were either already there at time t , or moved there. Expanding this exactly as in Eq. (9) and using Eq. (12), we get

$$\frac{\partial}{\partial t} n_g(\vec{x}, t) = \frac{|\Delta x|^2}{4\Delta t} \nabla^2 n_g(\vec{x}, t) - \frac{\partial}{\partial t} n_c(\vec{x}, t) \quad (14)$$

and, similarly,

$$\frac{\partial}{\partial t} n_h(\vec{x}, t) = \frac{k^2 |\Delta x|^2}{4\Delta t} \nabla^2 n_h(\vec{x}, t) + \frac{\partial}{\partial t} n_c(\vec{x}, t). \quad (15)$$

Note that if $k=1$, and we add the last two equations the variable $n_g(\vec{x}, t) + n_h(\vec{x}, t)$ obeys the diffusion equation, unaffected by the interaction between the subsystems (i.e., if we remove the distinctions between gas and heat, the combined variable simply diffuses without interacting).

To test the consistency between the microscopic diffusive dynamics implemented in our model and macroscopic descriptions given by Eqs. (14) and (15), we empirically measured the diffusion coefficient for gas and for heat particles as they diffuse about the space. Each particle should execute a random walk. The variance of the distance from the origin in the \hat{x} or \hat{y} direction σ_i^2 is proportional to the diffusion coefficient in that direction D_i , where $i = g$ or $i = h$. The exact relation is $D_i = \sigma_i^2 / 4p$, where p is the number of steps taken. For an unbiased random walk, the variance of the net displacement from the origin is $\sigma_i^2 = k^2 p$, thus $D_i = k^2 / 4$. For the gas particles ($k=1$) we find $D_g = (0.996 \pm 0.009) / 4$. For the heat particles, with $k=3$, we find $D_h = (9.00 \pm 0.08) / 4$. Thus the ratio of the heat to the gas diffusion length $D_h / D_g = 9.0 \pm 0.1$, agreeing with the theoretically predicted value of k^2 .

IV. THE MEAN FIELD LIMIT

The mean field limit corresponds to the ‘‘well-stirred reaction,’’ meaning that the reacting species are uniformly

spread throughout the space, and thus each particle feels the presence of the mean concentration of each species. For our system in equilibrium the gas particles and the heat particles *are* uniformly distributed throughout the space; it is only the crystal particles which do not obey this assumption. A uniform distribution means that there are no concentration gradients [$\nabla n_i(\vec{x}, t) = 0$ for all \vec{x} and i], and thus $\nabla^2 n_i(\vec{x}, t) = 0$. Also we can drop the explicit \vec{x} notation from the argument of the variables: $n_i(\vec{x}, t) = n_i(t)$.

Once the population level of the heat bath has reached the quasistatic steady state, the concentrations of all three species will remain essentially constant and the systems will have a well defined temperature from then on, as discussed in Sec. V A. We denote the time to reach the steady state (i.e., the time for the subsystems to reach the same temperature) as τ_T . We can now drop the explicit time notation from the arguments of the variables in steady state: $n_i(t > \tau_T) = n_i$. In this regime $dn_i/dt = 0$ and likewise $dn_i^j/dt = 0$, thus $\langle \mathcal{F}^\gamma \rangle = 0$:

$$\begin{aligned} 0 &= \langle \mathcal{F}^\gamma \rangle \\ &= 4[n_g(1-n_h)(1-n_c) - (1-n_g)n_h n_c] \\ &\quad \times [n_c(1-n_c)^3]. \end{aligned} \quad (16)$$

There are three solutions to Eq. (16). Each solution corresponds to fixed point of the dynamics. Only one is in the regime of interest. The fixed point at $n_c = 0$ corresponds to the presence of only gas particles. The fixed point at $n_c = 1$ is not allowed by the aggregation conditions (the aggregate can not have any closed loops). The remaining fixed point predicts that the equilibrium condition is $n_g[1-n_h][1-n_c] = [1-n_g]n_h n_c$. Noting that in the mean field limit ensemble averages equal spatial averages [i.e., $n_g(t) = \mathcal{N}_g(t)/L^2$], the constraints described in Eqs. (3) and (5) can be written respectively as $n_g(t) = \mathcal{N}_g(0)/L^2 - n_c(t) + 1/L^2 \approx \mathcal{N}_g(0)/L^2 - n_c(t)$ and $n_c(t) = n_h(t) + 1/L^2 \approx n_h$. After incorporating these relations the equilibrium condition can be expressed as

$$\frac{n_c}{1-n_c} = \mathcal{N}_g(0) \frac{1}{L^2} + O(n_c^3). \quad (17)$$

Figure 1 is a plot of the equilibrium value $n_c/(1-n_c)$ versus the initial density of gas particles, $\mathcal{N}_g(0)/L^2$. The solid line is the mean field prediction, Eq. (17). The points were obtained empirically from our simulations of three different system sizes $L = 128, 256,$ and 512 . The agreement between the three system sizes should be noted. The mean field approach makes predictions about the overall density of the system (hence the equilibrium temperature, as described below), but it does not make predictions about the growth morphology.

V. EMPIRICAL ANALYSIS

A. Temperature

The gas-aggregate subsystem and the heat bath subsystem together form a thermodynamically isolated system. These two subsystems are allowed to exchange energy only between themselves, and this energy is purely in the form of

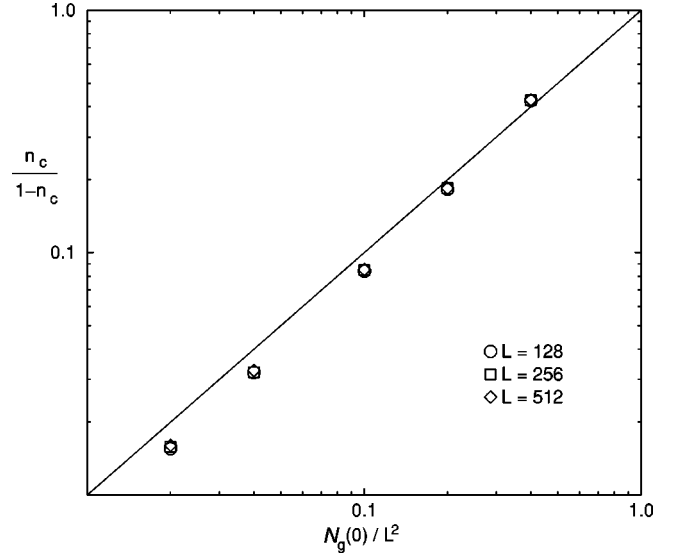


FIG. 1. The empirically determined equilibrium value of $n_c/(1-n_c)$ as a function of the initial gas density $\mathcal{N}_g(0)/L^2$ for systems of size $L = 128, 256,$ and 512 . The solid line is the mean field prediction. Note that the error bars are the same size as the points.

heat (ΔQ). As discussed in Secs. II B 1 and III the energy of the aggregate is a function only of the number of aggregate particles and is independent of their configuration. The total internal energy of the gas and heat particle species also is a function only of the number of particles of each species. Hence if invariant average population densities are achieved there is no further net exchange of heat between the subsystems, and they have by definition attained the same temperature.

The standard expression for the temperature of a two-level system [32], such as the heat bath in the RA model, follows directly from combining the definition of temperature ($1/T = \Delta S / \Delta E|_V$) with the microcanonical definition of entropy [$\Delta S = k_B \ln(\Omega_f / \Omega_i)$, where Ω denotes the number of microstates consistent with the macroscopic variables]:

$$\frac{1}{T} = - \frac{k_B}{\varepsilon_h} \ln \left(\frac{n_h}{1-n_h} \right). \quad (18)$$

Directly computing the temperature of the other subsystem is more difficult, but we can infer its temperature from that of the heat bath (note the gas particles are free to diffuse over the crystal, hence there is no change in the accessible volume for the heat particles as the crystal changes size and conformation: the crystal does no work on the gas, $PdV = 0$).

The approach to temperature equilibrium and a closeup of the subsequent fluctuations in temperature are shown in Fig. 2. Figure 2(a) plots the mean occupancy of the heat bath versus the time step into the simulation, with the corresponding temperature (in units of $k_B T / \varepsilon_h$) displayed on the right vertical axis. The initial growth is linear, with a slope of about 1.8. It then levels off near the quasistatic steady-state density of $n_h = 0.031$ (indicated in the figure as the dashed horizontal line). The results are averages over several independent realizations for three different system sizes $L = 128, 256,$ and 512 . Note that the three systems reach the same

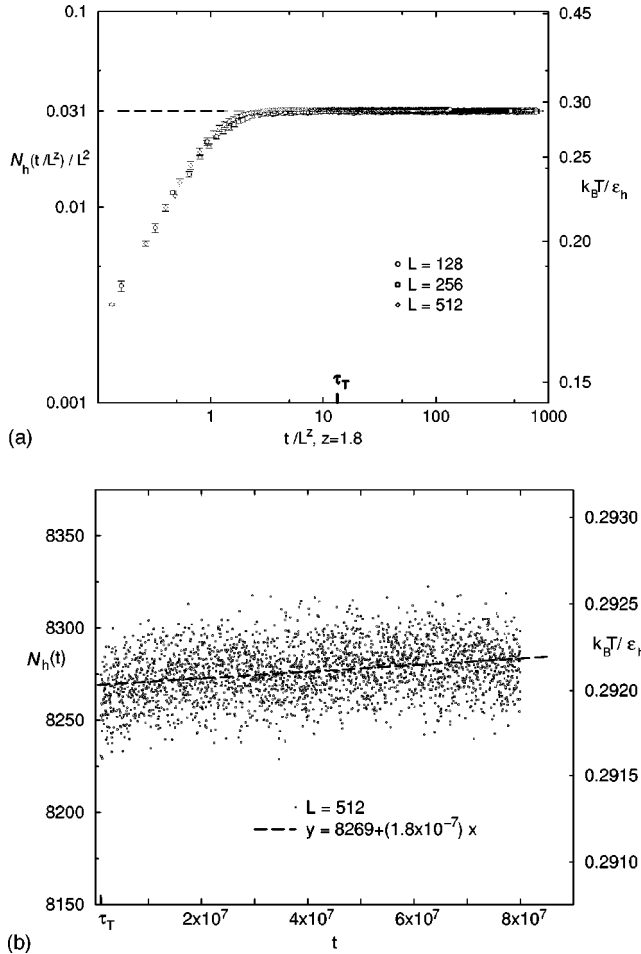


FIG. 2. (a) The mean density of heat bath particles as a function of time into the simulation, plotted for $L=128, 256,$ and 512 . The corresponding temperature is given on the right vertical axis. The initial growth of the heat bath density is linear, with a slope of about 1.8. Note that the steady-state density of the heat bath, and hence the temperature, is equal for all three system sizes, yet the time to equilibrate scales with the system size as $\tau_T \sim 10L^{1.8}$. (b) The actual average values of the total population of the heat bath as a function of time for every third time measured beyond τ_T . The dashed line is the result of a linear least squares regression on all of the data. Note the slight drift upward with time, of about two particles per 10^7 steps.

steady-state densities and hence the same temperature, but that the time to equilibrate depends on the system size. The data for the three systems was collapsed onto one curve by rescaling the time axis by the factor L^z , with $z=1.8$. The time to reach the equilibrium temperature is $\tau_T \sim 10L^z$. Note that this scaling behavior has an exponent which is slightly smaller than the diffusion exponent: the diffusion time is proportional to L^2 . As discussed in Sec. V B, the fractal dimension at time τ_T , $d_f(gt_T) \approx 1.8$: the time seems to scale with the fractal dimensionality instead of the Euclidean dimensionality of the space.

To study the details of the subsequent fluctuations we focus on the largest system size, 512×512 . As mentioned, during the initial period the growth of the heat bath population (and the size of the aggregate) is linear. It levels off at about 8270 particles on average, in a time which is less than 10^6 steps. The population continues to grow extremely

slowly after this point, rising by an average of 1.8 ± 0.1 particles every 1×10^7 steps, as determined by a linear regression based on about 8000 data points taken at equally spaced intervals in the regime where $t > \tau_T$. The probability that the slope is actually zero is 3×10^{-8} , as determined by a t-statistic comparing the ratio of the obtained slope to the sum of squares differences. Figure 2(b) shows a scatter plot of every third of the 8000 data points, overlaid by a straight line indicating the results of the linear regression on all 8000 points (only every third point is shown for visualization purposes: showing all the points results in a dense black cloud). The actual number of particles in the heat bath is indicated on the left axis, the corresponding temperature is given on the right. Although the temperature of the heat bath is not constant, it is very nearly so. Once the population levels stabilize the subsequent dynamics (i.e., the relaxation to the maximum entropy state for the crystal) is clearly a quasi-static process. The crystal does continue to exchange heat with the heat bath when it anneals, but the net heat exchange is essentially zero (the net heat exchange rate is $\sim 2 \times 10^{-7}$ particle per update of the space).

We measured also the fluctuations in population levels of the heat bath for the $L=256$ and 128 systems, in the corresponding regimes. Consistent with the $L=512$ system, we find the population rises by an average of 2.1 ± 0.1 particles every 1×10^7 steps.

B. Fractal dimension

The aggregate formed primarily while the heat bath contained less energy than its equilibrium level. Hence, if we continue running the dynamics, the cluster anneals; it evolves from a DLA-like cluster to a quenched branched random polymer structure. To quantify the cluster structure we calculate the fractal dimension of the aggregate and specifically how the fractal dimension changes as a function of the time into the simulation.

We measure the fractal dimension using a box-counting procedure which requires that we first establish the center of mass of the growth aggregate (which is typically not the initial seed particle, as the center of mass diffuses about the space as the aggregate anneals). An imaginary window box of edge length l is defined and centered on the center of mass. The number of lattice sites within that window that contain a crystal particle $\mathcal{N}_c(l)$ is tallied. The window size is increased and the count retallied. This procedure is iterated until the number of crystal particles contained no longer increases with window size. Before saturation, the number of particles contained should increase with some power of the window size

$$\mathcal{N}_c(l) \propto l^{d_f}. \quad (19)$$

The exponent d_f is the fractal dimension.

The RA cluster should initially resemble a parallel, irreversible DLA cluster of equivalent mass. Figure 3(a) shows a typical RA cluster for the $L=512$ system at the time $t = \tau_T$, which is the time for the mass of the RA cluster to stabilize at essentially its final mass ($\mathcal{N}_c \approx 8270$). Figure 3(b) shows a typical DLA cluster of equivalent mass. Both systems were initialized with a 4% density of diffusing gas particles. The gas particles still present at this stage of the evo-

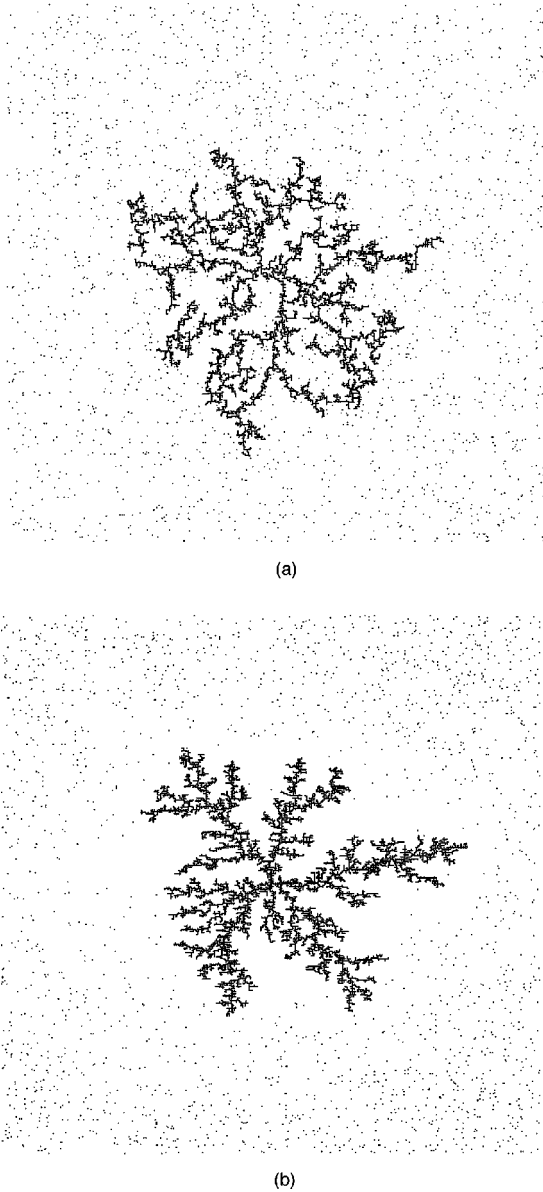


FIG. 3. Two growth clusters of the same mass $\mathcal{N}_c \sim 8270$. (a) A cluster grown via the RA model, pictured at time $t = \tau_T$, where τ_T is the time for the heat bath and gas-aggregate system to reach the same temperature. (b) A parallel DLA cluster. Note the gas particles, which are shown as the small dots. For the RA system the gas particles are distributed throughout the space, yet for the DLA system very few gas particles penetrate the region defined by the edges of the cluster.

lution are shown as the small dots in the figure. Note that for the RA system the gas particles are distributed throughout the space, yet for the DLA system very few gas particles penetrate the region defined by the edges of the cluster. The RA cluster has experienced enough annealing by the time $t = \tau_T$ to have a fractal dimension different than that of the DLA cluster, yet the radii of both clusters are comparable and approximately equal to a quarter of the length of the system ($r \approx L/4$). The RA cluster morphology is still far from its final morphology.

Figure 4 shows the box-counting results obtained for both models in the regime described above and pictured in Fig. 3. The top curve is for DLA, the bottom for RA. Both models

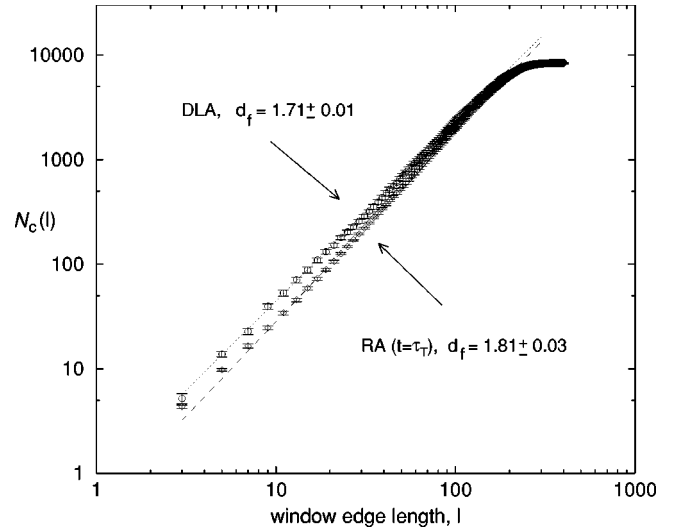


FIG. 4. The number of aggregate particles contained in a box of length l , as a function of l . The slope of the line is the fractal dimension. The top curve is for parallel DLA clusters of mass $\mathcal{N}_c \approx 8270$. The bottom curve is for the RA clusters sampled at $t = \tau_T$. Examples of these clusters are pictured in Fig. 3.

were implemented on a $L = 512$ size system. The vertical axis is the mass contained in the window $\mathcal{N}_c(l)$, the horizontal is the window size l . The curve for the DLA system is the average of ten independent DLA clusters of mass $\mathcal{N}_c \approx 8270$. The curve for the RA system is the average of ten independent RA clusters sampled at time $t = \tau_T$. The slope of the curve corresponds to the fractal dimension and was determined via a linear least squares fit. Consistent with past numerical studies of DLA [22,33], we find that the fractal dimension for DLA clusters is $d_f^{\text{DLA}} = 1.71 \pm 0.01$ (for difficulties associated with determining the fractal dimension of DLA see the detailed discussion in Ref. [24]). For the RA clusters the fractal dimension is $d_f^{\text{RA}}(t = \tau_T) = 1.81 \pm 0.03$. A line with this slope is shown overlaying each respective curve [34].

The RA cluster is less dense than the DLA cluster in the area immediately surrounding the initial aggregation site, however, the radii of both clusters are comparable. Several of the aggregate particles in the RA cluster have annealed away from the center to occupy the region between the center and the edge of the cluster. Hence, at the time depicted in Fig. 3, the RA cluster has a higher fractal dimension than the DLA cluster. The only constraints on the cluster are the number of particles and the connectivity. As there are more ways to have connected clusters of a specified particle number in an area of larger radius, the RA cluster evolves from the dense, bushy DLA-like structure shown in Fig. 3, to a tenuous structure which occupies more of the available lattice (with the initial increase in fractal dimension being a transient behavior). We ultimately expect to observe a diffuse structure with just a few meandering vines which can access more of the available configuration space.

As the time into the simulation advances, the density of the growth aggregate decreases, the radius of the aggregate increases, and hence the fractal dimension decreases. Figure 5 shows a typical RA growth cluster at the time $t = 80\tau_T$ timesteps. Note that the structure does resemble meandering

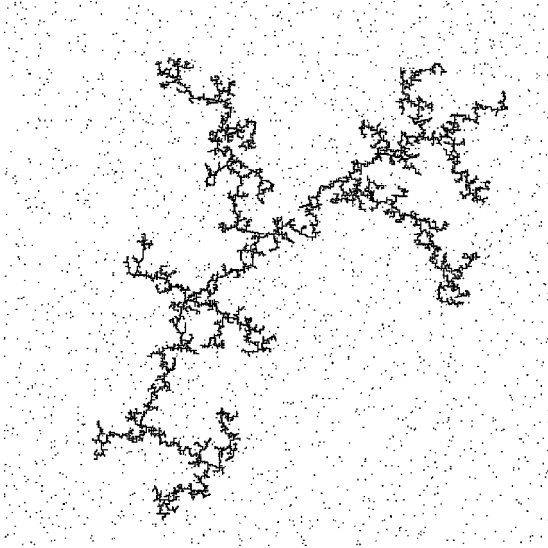


FIG. 5. A growth cluster grown via the RA model, pictured at time $t=80\tau_T$. The fractal dimension for this cluster ($d_f=1.63 \pm 0.02$) has seemingly reached the asymptotic value.

vines. Also the radius of the cluster is comparable to half of the lattice size ($r \approx L/2$).

Figure 6 is a plot of the average fractal dimension for RA clusters as a function of time into the simulation, for all three system sizes. The measurements reported below are averages over five independent realizations for the $L=512$ system, ten independent realizations for the $L=256$ system, and ten for the $L=128$ system (i.e., averages over either five or ten independently generated, large clusters). The data points and standard errors shown in the plots are the average and standard error over the set of independent realizations.

The fractal dimension is initially very close to the fractal dimension for DLA. We then observe a slight increase in the fractal dimension as the cluster center begins to anneal (an example is the RA cluster shown in Fig. 3), then a gradual decrease in the fractal dimension until it converges upon an equilibrium value. The solid line is drawn to denote the equilibrium value upon which results for the three system sizes

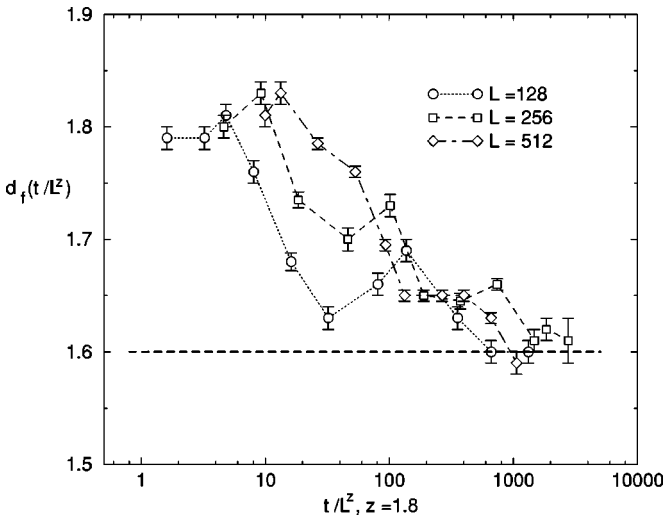


FIG. 6. The average fractal dimension of the RA growth clusters as a function of time into the simulation for all three system sizes.

are converging $d_f=1.6$. Using Flory-type scaling arguments it has previously been shown that a quenched branched polymer obeys the scaling relationship $N \sim R^{d_Q}$, with $d_Q = [2(D+2)]/5$ [35]. Here R represents the characteristic end-to-end distance of a polymer, and D the dimension of the space. R can be taken in direct analogy to l in Eq. (19), which defines the end-to-end distance of the window of interest. For $D=2$ the exponent $d_Q=1.6$. We should note that an exact result was obtained for quenched polymers in $D=2$, $d_Q=1.64$ [36], which is slightly larger. Flory-type scaling has also been studied for annealed branched polymers and the scaling exponent determined to be $d_A=(3D+4)/7$ [37]. For $D=2$, $d_A=1.43$. One might expect to observe a crossover from quenched to annealed behavior for the equilibrium RA growth clusters as we go from the large to the small system sizes, but we did not see this for the system sizes investigated.

Note that the time axis in Fig. 6 is rescaled by $L^{1.8}$ in order to match that of Fig. 2. Neither the fractal dimension nor the equilibrium temperature exhibit finite size effects as far as we can determine within the precision of our measurements.

VI. DISCUSSION

We have presented a microscopically reversible model which exhibits macroscopic pattern formation. In this model, entropy initially grows rapidly with time, and then subsequently grows exceedingly slowly—the slow relaxation can be characterized as a quasistatic isothermal process. The morphology of the aggregate formed by this dynamics changes markedly with time, evolving from a pattern having a conformation and fractal dimension similar to that of an irreversible DLA system, to a pattern characteristic of a branched quenched random polymer.

The RA model is an extension of the standard DLA model: we model the latent heat released when a gas particle aggregates in addition to modeling the gas and crystal particles. Since the dynamics is local and microscopically reversible, we are realistically modeling the flow of heat and the creation of entropy in this system and thus we model the thermodynamic behavior of growing clusters.

The model presented is simple and amenable to theoretical analysis. Given the constraints that we set out, an even simpler model could be constructed with just a single heat bath particle at each site, and a single gas particle at each site. In this case, diffusion would be performed by block partitioning [3]. This simpler model would have two fewer bits of persistent state at each lattice site than the RA model, and would be slightly easier to analyze theoretically. It would, however, be less computationally efficient: for a given lattice size, the volume available to the gas and heat particles would be reduced, but the computational effort required by each step of the simulation would be unchanged.

There are some simple variants of the RA model which might merit study. For example, we have only investigated situations in which the temperature is set by the size of the final aggregate. It would be interesting to study morphology

in situations where there is independent control of temperature and aggregate size. It would also be interesting from a thermodynamic perspective to modify the model by introducing a gas-crystal exclusion: gas particles would collide with the aggregate, but not diffuse over it. Thus there would be an excluded volume for the gas particles, and the crystal would do work on the gas as it grew.

ACKNOWLEDGMENTS

The work at MIT was supported by DARPA under Contract No. DABT63-95-C-0130. The authors would like to thank M. Smith, A. Erzan, and Y. Bar-Yam for discussions useful in formulating this model, as well as M. Kardar, B. Boghosian, and H. Mayer for helpful critiques.

-
- [1] G. Nicolis and I. Prigogine, *Self-Organization in Nonequilibrium Systems* (Wiley-Interscience, New York, 1977).
- [2] T. Toffoli and N. H. Margolus, *Cellular Automata Machines: A New Environment for Modeling* (MIT Press, Cambridge, MA, 1987).
- [3] N. H. Margolus, in *Feynman and Computation*, edited by A. Hey (Addison-Wesley, New York, 1998).
- [4] R. M. D'Souza, Ph.D. thesis, Massachusetts Institute of Technology, 1999 (unpublished).
- [5] M. J. Vold, *J. Colloid Sci.* **14**, 168 (1959).
- [6] J. M. Kim and J. M. Kosterlitz, *Phys. Rev. Lett.* **62**, 2289 (1989).
- [7] J. E. Pearson, *Science* **261**, 189 (1993).
- [8] E. Ben-Jacob *et al.*, *Phys. Rev. E* **53**, 1835 (1996).
- [9] M. Eden, in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, edited by J. Neyman (University of California Press, Berkeley, 1961), Vol. 4.
- [10] T. A. Witten and L. M. Sander, *Phys. Rev. Lett.* **47**, 1400 (1981).
- [11] R. Landauer, *IBM J. Res. Dev.* **3**, 183 (1961).
- [12] C. H. Bennett, *Int. J. Theor. Phys.* **21**, 905 (1982).
- [13] M. P. Frank and T. F. Knight, Jr., *Nanotechnology* **9**, 162 (1998).
- [14] M. Creutz, *Ann. Phys. (N.Y.)* **167**, 62 (1986).
- [15] Y. Pomeau, *J. Phys. A* **17**, L415 (1984).
- [16] G. Vichniac, *Physica D* **10**, 96 (1984).
- [17] J. Hardy, O. de Pazzis, and Y. Pomeau, *Phys. Rev. A* **13**, 1949 (1976).
- [18] U. Frisch, B. Hasslacher, and Y. Pomeau, *Phys. Rev. Lett.* **56**, 1505 (1986).
- [19] J. L. Lebowitz, in *Proceedings of the XIII Stiges Conference*, Lecture Notes in Physics, Vol. 445, edited by J. J. Brey, J. Marro, J. M. Rubi, and M. San Miguel (Springer, Berlin, 1995).
- [20] D. Levesque and L. Verlet, *J. Stat. Phys.* **72**, 519 (1993).
- [21] N. H. Margolus, *Physica D* **10**, 81 (1984).
- [22] H. E. Stanley, in *Fractals and Disordered Systems*, edited by A. Bunde and S. Havlin (Springer-Verlag, New York, 1996).
- [23] J. S. Langer, *Rev. Mod. Phys.* **52**, 1 (1980).
- [24] A. Erzan, L. Pietronero, and A. Vespignani, *Rev. Mod. Phys.* **67**, 545 (1995).
- [25] R. F. Voss, *J. Stat. Phys.* **36**, 861 (1984).
- [26] T. Nagatani, *J. Phys. Soc. Jpn.* **61**, 1437 (1992); K. Moriarty, J. Machta, and R. Greenlaw, *Phys. Rev. E* **55**, 6211 (1997).
- [27] M. A. Smith, Ph.D. thesis, Massachusetts Institute of Technology, 1994 (unpublished).
- [28] B. Chopard and M. Droz, *J. Stat. Phys.* **64**, 859 (1991).
- [29] N. H. Margolus, in *Pattern Formation and Lattice-Gas Automata*, edited by A. Lawniczak and R. Kapral (American Mathematical Society, Providence, 1996).
- [30] B. M. Boghosian and W. Taylor, *Phys. Rev. E* **52**, 510 (1995).
- [31] B. Hasslacher, *Discrete Fluids* (Los Alamos, 1987).
- [32] K. Huang, *Statistical Mechanics*, 2nd ed. (Wiley, New York, 1987).
- [33] T. C. Halsey and M. Leibig, *Phys. Rev. A* **46**, 7793 (1992).
- [34] Note that in the limit where the cluster size approaches infinity, a cluster grown via a parallel implementation of DLA will exhibit a crossover from fractal to two dimensional. In our simulations of parallel DLA the initial density of gas particles is dilute enough, and the cluster size small enough, that we do not see any crossover effects and we obtain smooth scaling curves as shown in Fig. 4.
- [35] M. Daoud, P. Pincus, W. H. Stockmayer, and T. Witten, *Macromolecules* **16**, 1833 (1983).
- [36] G. Parisi and N. Sourlas, *Phys. Rev. Lett.* **46**, 871 (1981).
- [37] A. M. Gutin, A. Yu. Grosberg, and E. I. Shakhnovich, *Macromolecules* **26**, 1293 (1993).