



Diffusion-limited aggregation: a kinetic critical phenomenon?

LEONARD M. SANDER

Diffusion-limited aggregation (DLA) is a model which represents noisy growth limited by diffusion. This process is quite common in nature and the simple algorithm gives a good representation of the large-scale structure of many natural objects. The clusters grown in the computer and the real objects in question are tenuous and approximately self-similar. A good deal is known about the algorithm, but a complete theory is not yet available. I review the current state of knowledge about the model, its applications and theoretical analysis of the results.

1. Introduction

In 1981 Tom Witten and I wrote an article [1] with the title above — without the question mark. It introduced a computer algorithm which we dubbed diffusion-limited aggregation (DLA). DLA represents noisy growth limited by diffusion, e.g. crystallization in a random environment. I remember at the time that I thought that the idea was pretty interesting mainly because I wanted to play with computer graphics — we had just acquired a new Hewlett-Packard pen plotter (now hopelessly obsolete). No one expected that eighteen years later the subject would still be alive.

However, a large number of scientists quickly became fascinated with DLA and started to do research on the subject. Our original paper has been cited almost 2000 times, and the literature is still growing unabated. From time to time, I have felt like Conan Doyle, who repeatedly wanted to kill off Sherlock Holmes — without success. In fact, after many years of working in other fields I am amazed to see that three of my recent papers have been ... about DLA.

What makes this subject so interesting and, in fact, rather peculiar are three facts.

- (i) The extremely simple process seems to seize the essential ingredients of a great many natural phenomena with very little physical input.

- (ii) It produces clusters of intriguing complexity which look very much like real objects which are random, tenuous and approximately self-similar. The mathematical fact that the simple algorithm makes self-similar (fractal) clusters is remarkable. The fact that things very like this occur rather commonly in nature is still more remarkable.
- (iii) The simple process in the algorithm has resisted analysis despite the fact that the model is very widely known. This is a *devilishly* difficult model to solve, even approximately.

The algorithm is the following: suppose we start with a nucleation centre, a single 'particle' of radius a which we locate at the origin of coordinates. Then release another similar particle at a random point some distance away. The new particle is allowed to diffuse, i.e. take steps of length a in random directions, until it happens to be within a of the first. Then it is stopped and added to the aggregate at the point of contact. Then a second walker is released and allowed to diffuse until it is within a of either of the first two, and so on. A DLA cluster of size N is the result of the addition of $N - 1$ particles to the original centre. The results of some computer runs of this type are shown in figures 1–3 for two-dimensional DLA. These are approximately fractal, as we will see below.

The algorithm has a very appealing simplicity, far simpler than some models in quantum field theory that can be worked through in complete detail. This has led many mathematically-inclined scientists to assume that it must be trivial. One of my friends told me that when he first heard Tom Witten talk about DLA he decided to go home

Author's address: H. M. Randall, Laboratory of Physics, The University of Michigan, Ann Arbor, MI, USA, 48109.




Figure 1. A small DLA cluster; $N=10\,000$




Figure 2. A medium-sized DLA cluster; $N=100\,000$

and exactly solve the problem that weekend. He could not do it, and no one has succeeded since then.

The scheme was motivated by physics, though not the right physics, as it turned out. Tom and I were interested in some experiments of Forrest and Witten [2] on a kind of coagulated aerosol. These are not described by DLA, but by a related process that I will talk about below. But, from the early work of Brady and Ball on electrodeposition [3] we began to see real objects with a spooky resemblance to

computer generated DLA clusters. See figures 4–6 for some examples. This, more than anything else has kept interest alive in the model despite its horrid difficulty.

Witten and I suggested that, as in the case of other geometric fractals (e.g. percolation clusters and self-avoiding walks), there should be a renormalization theory for this subject rather like that of critical phenomena. However, now, despite massive efforts we still have no idea whether this idea is correct. The sophisticated expansion and renormalization methods that have been so successful in field theory and phase transition theory have not worked here. I will give some reasons for these difficulties below. In fact, very little is really clear about the theory of DLA. Nevertheless it seems, if anything, more interesting now than it originally did. For a recent comprehensive review with many references see the book of Meakin [4].

In this article I will try to introduce this subject to a general community. This will necessarily be a personal view; a comprehensive treatment of the vast literature of this subject would be impractical. I will review the essentials of the model, talk about what natural shapes it can describe, and explain the very sophisticated computer simulations methods now available. Finally, I will outline what we can say about a theory.

2. Why it works—screening and the Mullins–Sekerka instability

The first question that arises in looking at a picture like figure 1 is to ask why the cluster has such a rough surface. If we think of the cluster as a coastline, there are very deep ‘fiords’. Why don’t they fill up? There are two complementary ways to understand this: if we start with a rough surface for the cluster, we can see that the fiords are *screened* and the cluster remains rough. Additionally, if we start with a smooth surface and a compact object without fiords, a growth instability roughens the surface.

Basically, the reason that fiords do not fill up in the cluster is that random walkers coming from the outside and hit one of the branches before they can go very deep inside—the fiords are screened. The situation would be quite different if the walkers moved in a straight line: as we will see below; then the cluster would be compact. This is a simple idea, but it means that the large-scale structure of a DLA cluster is dominated by non-local effects. Any approximation that neglects this feature will not work, and that makes the formulation of an analytical theory very tricky. Screening is easily stated, although it turns out to be one of the most elusive concepts to incorporate correctly in a theory.

We now turn the question and ask how the cluster gets to be very rough in the first place? We could speculate that a smooth outline, e.g. a disc, would continue to grow smoothly, and never get into the trap of screening. However, this is not so: it is easy to test that any initial

condition is soon forgotten in the growth [5]. If we start with a smooth shape it roughens immediately because of growth instability intrinsic to diffusion-limited growth.

This instability was discovered in the context of metallurgy by Mullins and Sekerka [6]. To see what they did, I restate the problem of diffusion-limited growth in continuum terms; this is known as the Stefan problem (see [7]), and is the standard way to idealize crystallization in the diffusion-limited case. Suppose that we have a density $u(\mathbf{r}, t)$ of particles that diffuse until they reach the growing cluster where they deposit. Then we have:

$$\partial u / \partial t = v \nabla^2 u, \quad (1)$$

$$\partial u / \partial n \propto v_n. \quad (2)$$

That is, u should obey the diffusion equation; v is the diffusion constant. The normal growth velocity, v_n , of the interface is proportional to the flux onto the surface, $\partial u / \partial n$. It is useful to estimate the size of the term $\partial u / \partial t$ by noting that if there is a typical velocity of growth, v , then $\partial u / \partial t \propto v \partial u / \partial x$. Now $|\nabla^2 u| \approx (v/D) |\partial u / \partial n|$. In the DLA case we launch one particle at a time, so that the velocity goes to zero. Hence equation (1) reduces to the Laplace equation,

$$\nabla^2 u = 0. \quad (3)$$

We are to solve an *electrostatics* problem and advance the surface proportional to the electric field at each point. This is called the quasi-static or *Laplacian* growth regime. In the Laplacian regime the diffusion constant drops out of the problem.

In addition to these conditions we need a boundary condition for u . For DLA this is pure absorption at the surface. In continuum terms:

$$u_s = 0. \quad (4)$$

(In the case of crystallization equation (4) is replaced by the Gibbs–Thompson boundary condition:

$$u_s = d_0 \kappa, \quad (5)$$

where κ is the curvature of the surface and d_0 is a measure of the surface tension. It expresses the well-known fact that crystals have a higher melting point when they have a curved surface. We will return to the significance of this term below.)

The boundary condition on u far from the cluster depends on the dimensionality. In dimensions greater than two we can take $u \rightarrow u_\infty$. However, in $d = 2$ we must generalize this. I will adopt the convention that the ‘charge’ on the aggregate is fixed at unity: $u(r) \rightarrow \log(r)$.

Now why do equations (1)–(4) not simply describe a smooth surface which advances in time? In fact, if we start with a flat surface it does advance in time (with $v \propto t^{1/2}$) but this solution is not stable [6, 7]. To see why this is so it is sufficient for our purposes to reason qualitatively. Suppose

we start with a flat surface with a small bump. Consider the electrostatic interpretation above. We are asked to find the potential, u , near a grounded conductor with a bump, and then advance the different parts of the surface at a speed proportional to $\partial u / \partial n|_s$, the surface electric field. The field is largest near the bump, as is known in elementary electrostatics: this is the principle of the lightning rod. Thus the bump grows larger.

In DLA bumps form on the surface due to the shot noise on the arrival of particles. They grow as a result of the instability and then *tip-splitting* occurs, in this case, again because of noise. The proliferation and interaction of tips through screening appears to give rise to a fractal, but the details of that process still remain obscure.

3. Examples of DLA-like growth

Our motivation in introducing the model was to describe growth where the limiting step is diffusion to the surface of the growing object. Such processes are quite common in nature. For example, when a crystal grows by deposition of matter of a new layer on a crystal surface an atom lands on a terrace, wanders about (diffuse) until it finds an island and attaches. Suppose we idealize the attachment process, and assume that the atom sticks with 100% probability where it first arrives, and does not rearrange later. Clearly the DLA model should have something to do with the shape so formed. Crystallization from a solution, electro-deposition, and many other processes can be idealized the same way under some conditions. We will see in this section that the results of such processes often really do look like DLA clusters.

3.1. Crystallization and island growth on surfaces

Materials scientists often grow surfaces using molecular beam epitaxy. In this technique a beam of material is generated in high vacuum and is allowed to fall on a clean surface. The atoms from the beam stick to the surface — they are then known as adatoms. At all but the very lowest growth temperatures they diffuse on the surface until they find a stable site and start to form a new layer. Usually through the nucleation and growth of islands. In the proper temperature regime (see below), islands with striking tenuous shapes are formed which look very much like DLA clusters.

This phenomenon is easily understood from the discussion above and the continuum description of equations (1)–(5). Island growth clearly has a diffusive instability [8]: the adatoms move by diffusion and stick to the edges of islands that are already nucleated. For an example, see figure 4 for DLA-like islands on Rh [9]. Often, as the temperature is raised, compact islands are formed because there is adequate diffusion around the island periphery to fill in

the fiords. This is how the difference between equations (4) and (5) manifests itself: if there is negligible rearrangement, surface tension, which favours compact islands, plays a small role in the growth.




Figure 3. A large DLA cluster, $N=100\,000$. The radius of this object is roughly 15 times that of the cluster in figure 1. The colours represent the time of arrival: thus white is the first $1/10$ of N , grey the second $1/10$, etc. There are ten colours in all.




Figure 4. Islands on the surface of Rh observed with a scanning tunnelling microscope. Each island is about 50 \AA across. The grey levels represent heights: there are a few multilevel islands. *Courtesy of R. Clarke*

However, the story only begins there. The interaction between the island shape and the formation of new layers has been studied intensively for about a decade [9]. Basically, if the islands are DLA-like, they usually do not trap adatoms on top. Thus we have island nucleation and growth followed by merging of islands until a new layer is formed. This is *layer-by-layer* growth. Compact islands, however, often have an Ehrlich–Schwoebel barrier [11], that is, there is a barrier for adatoms to fall off. (Why this barrier is suppressed for DLA-like islands is not precisely known — it may be that the extra edge length is enough to allow escape of the second layer adatoms.) Then new islands can be formed on top of the existing ones and the growth is not layer-by-layer. At still higher temperatures the adatoms can jump the Ehrlich–Schwoebel barrier, and layer-by-layer growth returns. This re-entrant layer-by-layer growth created considerable interest in the materials science community.

3.2. Viscous fingering

Equations (2)–(3) are a system which has been intensively studied in another context, that of fluid flow. In a classic paper Saffman and Taylor [12] considered the displacement of a viscous fluid (like oil) by an inviscid one (like air) under conditions of creeping flow, for example in a porous medium, or between thin parallel plates with a small gap (a Hele–Shaw cell). They showed that an air bubble does not simply displace oil: because of an instability identical in form to the Mullins–Sekerka instability, it forms a steady-state finger in the middle of the channel.

The relationship to equations (1)–(5) is easy to see: we need only point out that flow in porous media or a Hele–Shaw cell is described by an empirical rule called D’Arcy’s law which gives for the fluid velocity, $\mathbf{V} = -K \nabla P$, where P is the pressure in the viscous fluid. The proportionality constant depends on the viscosity. Since most fluids are almost incompressible,

$$\nabla \cdot \mathbf{V} = -K \nabla^2 P = 0. \quad (6)$$

We take the zero of pressure to be that in the inviscid fluid (since it has small viscosity, its pressure there is approximately constant). We have

$$P_s = \sigma \kappa, \quad (7)$$

resulting from the pressure drop due to the curvature of the interface, where σ is the surface tension. At the interface D’Arcy’s law reads

$$V_n = -K \partial P / \partial n. \quad (8)$$

These three equations are the same as equations (2), (3) and (5). Presumably DLA has a close relationship to Hele–Shaw flow *in the presence of noise and for small surface tension*.

Paterson pointed out this relationship [13] and another fascinating fact: experiments on the Hele–Shaw problem look like DLA in some cases with the air bubble playing the role of the cluster. Specifically, when noise plays a role, in a porous medium or in fast flow in a Hele–Shaw cell, then the Saffmann–Taylor finger breaks up into subfingers that strikingly resemble DLA clusters. Even more striking is the fact that for *radial* flow there is no need to introduce external noise, and radial Hele–Shaw flow makes patterns which simply look like DLA clusters. See figure 5 for an example.

3.3. Electrodeposition

When Witten and I started playing with our model in the early eighties, I began to wonder if it was just a mathematical curiosity, or whether we could see some real example of DLA growth in nature. My doubts were laid to rest once and for all by a very elegant paper in Nature by Ball and Brady [3] who made electrodeposits of copper on the end of a thin wire in diffusion-limited conditions. The process is the familiar one, the reduction of Cu from CuSO_4 . By examining electron micrographs and from indirect evidence Ball and Brady showed that the clusters seemed to be self-similar over five orders of magnitude in size.

This was an impressive achievement, but the deposits were fragile and could only be rather small. The aesthetic appeal of the subject took a big jump when a Japanese group [14] and two American groups [15,16] made two-dimensional deposits which could be quite large because they were supported on a glass cell. I will describe the Michigan experiment [16] as an example.

We made a thin cell by confining a film of electrolyte between two plexiglass plates spaced by about 0.1 mm. There was a cathode in the centre and a ring anode of about 10 cm diameter around the periphery. The metal deposited on the cathode, and we found that for slow growth (voltages less than 1 V and small concentrations of electrolyte) we indeed got structures that looked like DLA, see figure 5. (Note, however, that our voltages are always very high compared to a conventional electrochemistry experiment). For other conditions we produced other patterns: for example, for the highest voltages we used (on order 5 V) we could make ordered crystalline arms.

It soon became evident that the situation was quite complicated, even for the low-voltage case. In fact, we showed that our cells were ohmic, so that the drift current, not the diffusion current was the main contribution to the growth, unlike the case of the Ball–Brady experiment. However, as I pointed out above, the pattern formation in




Figure 5. A radial viscous fingering pattern. Air is injected through the tube in the centre and displaces fluid (glycerin) which is confined between two plates held 1 mm apart. The pattern is about 20 cm across.




Figure 6. A zinc electrodeposit produced in a thin cell. The electrolyte is confined between two plexiglass plates held 0.1 mm apart. The cathode is inserted through a hole in the plate, and there is a ring anode (not shown). The pattern is about 3 cm across.

this case is expected to be similar, and the DLA pattern is really no surprise.

There are many fascinating details of this process which I have glossed over. In fact, [15,16] have spawned a minor industry which attempts to take into account the real details of the chemistry in this case. For more information the reader should consult [17].

3.4. Bacteria colonies and other examples from biology

The growth of bacteria colonies in Petri dishes is a common experimental technique. Normally the overall shape of the colony is rather uninteresting: starting from an infection site a rough, more or less round colony develops. However, this is not necessarily true if the growth takes place under conditions of stress. In 1989 Matsushita and collaborators [18] showed that if the food supply of a colony of *bacillus subtilis* was reduced, the overall outline of the growth looked very much like a DLA cluster! See figure 6 for an example (with another type of bacterium). The reasoning was that multiplication of the bacteria at the surface could be bottlenecked by the diffusion of the food supply (in this case, peptone) and that the growth of the tips of arms of the colony is enhanced. The bacteria inside are not nourished and, in fact, become dormant.

This observation gave rise to a large amount of activity, and a peculiar incursion of physicists into biology. When examined in detail, the situation is very complicated indeed. Many effects have been considered such as chemotaxis, mutations and the secretion of fluid by the bacteria. Bacteria can do many more things than just exhibit




Figure 7. A colony of *Paenibacillus dendritiformis* bacteria, T morphotype, grown on hard agar and under severe starvation. The pattern is about 10 cm across. *Courtesy of E. Ben-Jacob*

unstable growth. For example they can grow spirals. A recent informal review is given by Ben-Jacob [19] who reviews both experiments and modelling. However, it remains true that the general observation of Matsushita *et al.* remains valid: the shape of colonies in some limiting situations is like DLA, and this is embodied in models such as the ‘communicating walkers’ model of Ben-Jacob and collaborators [19] which has a diffusive instability [20].

Other branching patterns in biological systems look very much like DLA, and a number of workers have tried to see whether this is mere coincidence, or corresponds to a real effect. For example, Family *et al.* [21] measured the fractal scaling of the blood vessels in the retina, and found that the self-similarity was very much like a two-dimensional DLA cluster. Quite recently Fleury and Schwartz [22] have given a model based on the formation of blood supply in an embryo which gives DLA patterns.

I am aware of several other examples of biological patterns which are governed by diffusion-limited processes and give rise to DLA-like patterns (e.g. in the growth of certain sponges [23]), and there are undoubtedly a large number that I have not heard of. The application of DLA and related models for the case of living systems is still in its infancy.

4. DLA and diffusion-limited growth

There is a good deal of general theory associated with diffusion-limited growth without disorder [24]. In this section I will mention some results from this theory and see how they help us understand DLA.




Figure 8. The radius of gyration of a DLA cluster as a function of N .

4.1. *Orderly viscous fingers*

The mathematical theory of viscous displacement in Hele–Shaw flow is very well developed [25]. The general picture that emerges is as follows. Suppose we consider the development of a finger in a channel geometry where we inject air from one end. Then, independent of the exact form of the initial condition a finger will develop which moves down the channel and which occupies a fixed fraction of its width. This fraction goes to one-half as the surface tension goes to zero. The selection of the width is thought to depend on the existence of surface tension in an essential way. This solution to the problem is known as the Saffmann–Taylor finger.

The disorderly clusters of fingers that look like DLA are formed in cases where noise plays a role. Theory [25] shows that the threshold for stability of a single finger decreases as the velocity increases and as surface tension decreases. Fast fingers have small thresholds and, as a practical matter, quickly destabilize. Arneodo and collaborators [26] discovered a fascinating and unexplained relationship between disorderly fingers of this type, DLA, and the orderly Saffmann–Taylor finger. They superimposed a collection of disorderly fingers and an ensemble of DLA clusters grown in a channel. Then they chose a reasonable criterion for tracing an average outline of the fuzzy superpositions. This turned out to be exactly the Saffmann–Taylor finger for small surface tension! Levine and collaborators have used this insight as a motivation to develop a mean-field theory [27] based on earlier work of Witten and Sander and others [5,28] which tries to describe the superposition of members of the ensemble of DLA clusters. There remain many ambiguities in this area.

For the radial case with surface tension very little is known from a mathematical point of view. It appears that there is no stable steady state — the noise threshold may be zero. We associate this lack of a steady state with the presence of a dynamically evolving pattern which eventually becomes fractal.

There are many known exact solutions to the problem without surface tension [29], but these are probably unstable. For a generic initial condition without surface tension the surface will develop cusps in finite time: that is, without surface tension the Stefan problem is ill-posed and some finite σ is necessary to regulate the problem.

4.2. *Dendritic growth and anisotropy*

Many mathematicians and mathematical physicists have studied the Stefan problem, equations (1), (2) and (5) [24], which models the diffusion-limited growth of a crystalline solid. The relevant results may be summarized as follows. Crystals grow in many modes, but one of them, the formation of dendritic tips (such as one of the branches of a snowflake) is of particular interest to us. In this case a single

‘finger’ can form, although usually with sidebranches, and it can grow so that its steady-state shape translates with constant velocity.

The steady-state tip is not a stable solution to the problem unless a new feature is introduced which corresponds to an important effect in real crystals, the anisotropy of the lattice. This manifests itself because different crystal faces have a different surface energy and we should write $u_s = d_0(\theta)\kappa$ in equation (5), where θ is the angle with respect to the crystalline axes. The surprise in all this is that an arbitrarily small anisotropy can stabilize single tip growth — this is important since the actual anisotropies in crystals are usually less than a few percent.

In the absence of anisotropy, numerical solutions [24] show a sequence of *tip-splittings* where tips form, split and reform. Thus crystal growth is stabilized by anisotropy, and in its absence, we expect tip-splitting. There are direct analogues of this statement in simulations of DLA, as we will now see.

5. Variations on the theme

When the DLA model was first proposed there was a natural scepticism in the community about the relevance and meaning of this simple process. One of the first questions we asked was how robust the process was. In the course of time many variations of the original model have been tried, and many of them lead to interesting insights about the process, and about disorderly diffusion-limited growth in general.

5.1. *Off-lattice and on-lattice DLA; anisotropy*

The DLA model in our original paper was not quite the one I described above. Instead, I had the random walkers and the cluster defined on a square lattice: particles could only live on lattice points. The computations were much easier in that case, and given the primitive state of computer resources at that time, the increase in speed was very much worthwhile. However, it was obvious that it was necessary to check whether the results were lattice dependent. I was very proud to be able to produce a small cluster on a triangular lattice which looked very much like the square lattice results, and looked like it had the same scaling. Indeed, for small clusters ($N < 100000$) the growth appeared to ignore the lattice.

The situation changed radically when the Brady–Ball algorithm (see below) was introduced and very large clusters became available so that lattice effects started to be visible [30]. Robin Ball visited Ann Arbor for a summer and we worked together to make some very large DLAs. We proved definitively that for large ($N \sim 10^6$) clusters the overall shape of the cluster became elongated along the x and y axes of the lattice. This is presumably related to the extreme sensitivity of dendritic growth to anisotropy.

Family and Hentschel [31] have given some arguments which show why the square lattice shows the extreme distortions that we found. In a very interesting development Ball *et al.* [32] showed how to understand the overall shape distortions in the special case where the sticking probability along the x and y directions were not the same. However, a complete understanding of the matter must await a theory of the growth—then we may be able to add the anisotropy back in as a perturbation. For the moment it is universal practice to work only with off-lattice DLAs.

5.2. Laplacian growth

Thus far we have regarded the Laplace equation (3) with the boundary condition of equation (4) as a representation of the probability of the arrival of a random walker at a point outside of the cluster. It is possible, however, to solve this equation directly using, say, the relaxation method. Pietronero and collaborators [33] did this, and modelled growth by saying that the probability to grow at a point on the surface of the cluster is given by:

$$P_s \propto \partial u / \partial n, \quad (9)$$

cf. equation (2). The model defined by equations (3), (4) and (9) might be expected to be the same as DLA neglecting some details, and the experiments of Paterson on viscous fingering (see above) also were in accord with this expectation. The simulations of [33] were consistent with the picture.

This development is quite significant for several reasons. First, we can regard this scheme, now known as *Laplacian growth*, as a model in its own right. Most modern theoretical developments on DLA go back and forth freely between particle language and Laplacian growth. Second, the authors pointed out (following a suggestion in the original Witten–Sander paper) that the equations stated can be viewed as a representation of dielectric breakdown (or lightning), where u represents the electrostatic potential in the insulator breaking down, and equation (9) says that the probability of further breakdown on the surface of the ionized region is taken as proportional to the electric field at that point. They also did measurements of real breakdown patterns (Lichtenberg figures) and showed that they had the fractal dimension of DLA. Finally, Pietronero and collaborators generalized the growth condition to read:

$$P_s \propto [\partial u / \partial n]^\eta, \quad (10)$$

where η is a parameter representing nonlinearities in the breakdown characteristics. They found that for each η there is a different fractal growth process with a different fractal dimension. For $\eta = 1$ the problem reduces to ordinary DLA.

5.3. Surface tension

The role of surface tension in diffusion-limited growth has been glossed over in all of our discussion so far. For the noise-free problem, as we saw above, surface tension plays a central role in the pattern formation. However, DLA has no surface tension and still gives rise to patterns similar to those seen, for example, in radial Hele–Shaw growth, where $\sigma \neq 0$.

There are two aspects of the problem to be considered. One, surface tension regulates the usual Hele–Shaw problem by eliminating small scale singularities. In DLA the finite particle size plays this role, although the regulation is quite different, and simply cuts off features which are too small. On the large scale, on the other hand, surface tension determines the overall shape of a Saffmann–Taylor finger and the fraction of the channel which it fills. For DLA the overall pattern is determined by noise, not σ . Thus the role of surface tension is secondary and usually only gives rise to clusters with thickened branches which look very much like noisy Hele–Shaw patterns.

This has been demonstrated numerically in various ways. The most primitive version of surface tension [5] is to idealize this effect in the most basic way possible, in terms of bond counting and to suppose that the sticking probability for a site was higher if there was more than one occupied neighbour. If we set $j =$ number of occupied neighbours to a growth site, we might put $p_j \propto p_0^{A-j}$ for a square lattice, with $p_0 < 1$.

However, this is not a complete description of macroscopic surface tension which corresponds to rearrangements of matter because the surface is in local equilibrium. Kadanoff [34] and others [35] have shown how to deal with this in a particle simulation. Basically, they allow particles to be re-emitted from the surface with a probability proportional to the local curvature. Then, if the particle is reabsorbed before it escapes to infinity, the growth is allowed.

5.4. Cluster–cluster aggregation

The original problem that Tom Witten and I thought we were solving had to do with the formation of a wispy bit of smoke by aggregation of small particles. This is an interesting process and gives rise to an interesting model, but it is not DLA. The reason is that aggregates in an atmosphere have some mobility: they can move and combine. In this case we have a different sort of physics, namely the aggregation of aggregates or cluster–cluster aggregation. A model for the process [36] produces fractals of a different sort than DLA. This model is directly applicable to many situations in colloid and aerosol chemistry.

In the model, one begins with a large collection of particles each of which is allowed to diffuse until it

encounters another. Then the resulting cluster continues to move until large clusters result. Cluster–cluster aggregates have a fractal dimension which is much smaller than that of DLA. Some results [4] are $D(2) \approx 1.43$, $D(3) \approx 1.75$ (compare DLA with 1.7, 2.5 respectively). The reason is simple: however since particles have much difficulty in wandering down fjords, clusters will have much more trouble and will stick near to the surface. This is true if the kinetics is such that at any time the cluster size distribution has a mean that increases in time so that a cluster does not encounter many individual particles, and the aggregation is dominated by aggregation of clusters of similar sizes. Any ‘normal’ kinetics turns out to have this property, as we will now discuss.

We noted that we must continue to let aggregates diffuse. In order to completely define the model, we must decide on how fast they are to diffuse, namely, what the diffusion coefficient of an aggregate is to be. If we take Stoke’s law of friction for the diffusion, the retarding force is proportional to the inverse radius of the cluster. In the simulations it is usual to take the diffusion coefficient, v , of the clusters to be a power law of the mass [36,37]:

$$v \propto M^\gamma . \quad (11)$$

For Stoke’s law $\gamma = D - 1$. In the simulations it is usual to take γ to be a free parameter. In fact, the fractal dimensions are independent of γ (and equal to the values above) if $\gamma < 1$.

In contrast to the DLA case, theory for this process has been quite successful, and its applications have turned out to be useful in many contexts [4]. The deep puzzles that plague the DLA model are absent here, and we can, without too much exaggeration, think of cluster–cluster aggregation as a solved problem.

6. Fractal scaling

The fact that DLA clusters look like macroscopic objects is really very odd. The biggest computer generated two-dimensional DLAs now have $\sim 10^7$ particles [38]. However, this is a tiny number compared to those for a real macroscopic object which has $\sim 10^{20}$ particles.

From the beginning we realized that what we were doing was unconventional. When I first started plotting pictures of DLAs (in 1980) we used our wonderful pen plotter and it made a lot of noise. My colleague, Bob Lewis, walked by one day and asked what the racket was about. I explained that we were looking at the shape of crystals. ‘By simulating them molecule by molecule?’ he asked. I will never forget the are-you-crazy look on his face when I said yes. This is clearly a hopeless enterprise, but for the striking fact that a small DLA cluster and a large one look more or less the same (cf. figures 1, 2 and 3)! This feature, (approximate) scale-invariance, is what caught everyone’s attention, and

which remains a puzzle to this day. It also meant that we were not as crazy as we looked.

6.1. Scale invariance and the correlation dimension

The most obvious way to quantify the scale-invariance of a DLA cluster is to use the fractal geometry of Mandelbrot [39]. The word ‘fractal’ is defined in many ways by different authors. For the purposes of this article I will use the term to mean a geometric object in which the part is like the whole. This amounts to saying that if we take a subbranch from a cluster and blow up the picture, it will be statistically the same as a main branch. An equivalent statement is that the correlations of the matter in a cluster have no characteristic scale. Thus the correlation functions must be power laws in distance.

A quantitative way to put this is to define a generalized dimension (the correlation dimension) by assuming that the number of particles within a distance, r , of any particle on the cluster obeys:

$$\langle M(r) \rangle = A r^{D_2} , \quad (12)$$

where A is some constant. Then the two-point correlation function for the matter density, $\rho(r)$, is also a power law:

$$C(r) = \langle \rho(\mathbf{r} + \mathbf{s})\rho(\mathbf{s}) \rangle / \langle \rho(\mathbf{s}) \rangle \propto r^{d-D_2} , \quad (13)$$

where d is the dimension of space. For a non-trivial fractal D_2 is not an integer.

The correlation dimension has been measured many times for DLA clusters. Most often equation (12) is used in the simple form of looking at the *mass dimension*, i.e. figuring out how many particles (the total mass) lie within r of the origin, see figure 8. For DLA in two dimensions the data fits very well to $D_2 = 1.71$. Some workers have looked at DLA in dimensions up to eight [40].

Still another kind of dimension is the *box-counting* fractal dimension, D_0 . This is defined by covering the fractal with boxes of size ϵ . Then set $\lambda = \epsilon/R$, where R is the overall size of the cluster. Now count how many boxes there are in the covering and call this N_λ . The box-counting dimension is defined by

$$N_\lambda = B\lambda^{-D_0} , \quad (14)$$

where B is some constant. For DLA in two dimensions this quantity is close to 1.7. For the relationship between D_0 and D_2 see the next section.

The case of $d = 3$ is of particular physical interest. (The reason for looking at $d > 3$ is mainly theoretical; see below.) The best current value of D_2 is 2.49. In my opinion, too little attention has been paid to the three-dimensional case. As we will see in the next section, there is reason to

believe that DLA in $d = 3$ is better behaved than $d = 2$. Unfortunately, the beautiful pictures that caught everyone's attention are harder to produce for dimensions higher than 2. The intricate substructure of the cluster is hidden.

6.2. Multifractal scaling and the growth probability

The probability of growth of a DLA cluster is *very* inhomogeneous. Tips of the cluster are much more likely to grow than points inside the fiords, and the inhomogeneity grows with cluster size. This fact has led many workers to apply the theory of *multifractals* to DLA [41]. This theory arose in the theory of turbulence [42] and has also been applied in dynamical systems theory. However, the growth probability in DLA is one of the most successful and striking applications.

The growth probability is a measure defined on the surface of the cluster, namely a function, μ , such that $0 \leq \mu \leq 1$. For DLA $d\mu = \partial u / \partial n$. Another interesting measure is the uniform (or mass) measure, which is constant on the cluster.

We divide the surface of the cluster into boxes, as above. The probability for the i th box is $p_i = \int_{\text{box } i} d\mu$. For example, for the harmonic measure, p_i is the probability that a random walker will attach inside a box. If the probability varies wildly, the integral will scale with the box size. A measure is multifractal if it has two properties: (i) $p_i \sim \lambda^\alpha$; that is, p has singularities (e.g. near a sharp tip of the cluster). Now count how many boxes have the same singularity strength, and call this $N(\alpha)$. The second property is (ii): $N(\alpha) \sim \lambda^{-f(\alpha)g(\alpha)}$, where g is a smooth function. If $f(\alpha) > 0$ we can interpret it as a fractal dimension of the set on which singularity α occurs: compare equation (14). Now a plot of f versus α is a histogram of the information about the singularities.

There is another way to parametrize a multifractal measure, by defining generalized dimensions:

$$p_i^q \sim \lambda^{(q-1)D_q}. \quad (15)$$

The D_q and $f(\alpha)$ contain the same information: they are related by a Legendre transform [41].

Now consider the uniform measure. It is easy to see that the D_0 and D_2 defined from equation (15) are the same as those in the previous paragraph. A cluster such that all the D_q are the same under the uniform measure is an ordinary (mono-) fractal. DLA is, at least approximately, a fractal in this sense. However, some authors [43] have suggested that there are deviations from strict scaling, and that DLA clusters are multifractal with respect to the uniform measure. In my opinion the evidence for this idea is not completely convincing.

For the growth probability, things are much more interesting [41]. There is quite a lot of evidence that for

positive α (the tips, and in general, the outside of the cluster) there is well-defined multifractal scaling. Correspondingly, for $q \geq 0$ the D_q are different and well defined. However, for negative α (or negative q) we are dealing with the small probabilities which live in the fiords. Here, at least in two dimensions, the situation is strange. Basically, these small probabilities do not scale with a power law in system size, but more like an exponential. The harmonic measure varies *even more wildly* than can be accommodated in the multifractal formalism. However, in three dimensions the situation is nicer [44]: the harmonic measure seems to be multifractal. This is reasonable: in two dimensions fiords are completely screened because lines of connected particles completely block the entrance to the dead regions. In three dimensions the local topology is still linear, so particles can 'go around' the block and enter more deeply.

There are interesting relationships for the D_q for DLA. For the case $q = 1$ (which is, in fact, defined from equation (15) by the limiting process $q \rightarrow 1$) there is an exact relationship: $D_1 = 1$ [45]. Halsey has shown that we should expect $2D_3 = D_0$. In a recent paper Halsey and collaborators have suggested that there are enormous fluctuations for the $D(q)$ for $q > 3$ [46].

6.3. DLA really self-similar?

In our original work on the DLA model we assumed that we were producing self-similar fractals in the sense mentioned above, namely that the part was like the whole. The evidence available at that time was really insufficient to verify that assumption, but it seemed to be approximately true. For example, plotting the mass-radius relationship always gave (and gives) an excellent fit to a power law. However, there are other features of the simulations that do not seem to satisfy the hypothesis. This is one of the disturbing features of DLA which make it a continuing mystery.

An example which appeared very early was an analysis of the growth zone — the region of space over which new particles are added — by Plischke and Racz [47]. Their numerical evidence seemed to say that the growing region was a fraction of the radius which decreased with N . That is, the width of the growth zone was given by $N^{\nu'}$ whereas the radius grows as $N^{(1/D)}$. If, as seemed to be the case, $\nu' < 1/D$, we get a sharper and sharper growth zone. However, for a self-similar object we would expect that the growth zone would be proportional to the radius, say always being a certain percentage of it.

Later, more careful work for larger clusters [48] found that ν' increased with N . This is an example of a *slow crossover* to self-similarity, a phenomenon which plagues this subject, and which is not understood at all. Even more recently, another group has gone to even larger sizes

$N \approx 10^6$ [38]. They claim that the data are consistent with the possibility that DLA is never self-similar, but exhibits ‘infinite drift’.

There are many other examples of this kind of phenomenon. In my opinion, all of them are crossovers, and that the asymptotic state of DLA is self-similar. This, however, is hard to prove, and more important, crossovers that persist for million particle clusters are in themselves interesting and cry out for explanation. We have no theory which can treat these unusual phenomena, although, for some very recent work, see below.

7. Computer simulations

My first DLA program was hopelessly inefficient, and, in retrospect, embarrassingly naive. The worst thing that I did, initially, was to start the random walker far from the aggregate to represent the real physical situation. Unfortunately, the program spent most of its time aimlessly propagating the particle and not finding a place to stick.

However, I was discussing this over dinner, and my companion asked me why I did not start the particle at a random angle on a circle near the cluster. ‘Because its arrival probability at the cluster would not be random if it started close by’, I said. ‘How do you start it far away?’ she asked. ‘At random’, I said. ‘How much more random do you think it will be when it gets to the circle?’ she said. Whoops. I fixed the program using the trick of starting on a circle—and we quickly had nice clusters to look at.

A large number of people have worked out techniques to make DLA simulations tractable for quite large sizes [30,40,49]. The current record is on the order of 30 000 000 particles in two dimensions. I will outline in this section some of the methods used, including a very interesting algorithm originally introduced by Ball and Brady [30].

7.1. Some simple tricks

One simple trick is the one mentioned above: the random walker need not actually start far away from the aggregate. It can start at a random point on a circle of size R_c which just encloses the cluster since its probability of arrival on this circle is random. (This observation is due to M. E. Sander.)

The walker may wander away from the aggregate, outside of R_c . In that case, it is necessary to allow a free walk until it is quite far away (many aggregate radii). This is not a serious problem since it is possible to allow the walker to take large steps when it is outside. It cannot encounter any matter, so it can take a step as large as the distance to the enclosing circle, but in a random direction. (This trick was invented by P. Meakin.)

However there is a better way to do this. Rather than taking computer time to let the particle walk around outside the circle, we can instead find the probability density on the boundary and bring back the particle in one jump. This amounts to finding the Green’s function, $G(\mathbf{r}, \mathbf{r}_o)$, for equation (1) for a particle outside an absorbing disc of radius R_c . Then the probability to land on the disc is given by $\partial G / \partial r$ at $r = R_c$. The solution is easy to find using the method of images [50], and can be adapted to different boundary conditions quite easily.

The final result for a circular geometry is this: if the particle is at position $\mathbf{r}_o = (x_o, y_o)$ outside of R_c , then we calculate the position on circle as follows. Pick a random number, ξ , and put the particle at $\mathbf{r} = (x, y)$:

$$\begin{aligned} x &= \frac{R_c(1-V^2)x_o - 2Vy_o}{1+V^2}, \\ y &= \frac{R_c(1-V^2)y_o + 2Vx_o}{1+V^2}, \\ V &= \frac{r_o - R_c}{r_o + R_c} \tan(\pi\xi). \end{aligned} \quad (16)$$

7.2. The method of hierarchical maps

The Green’s function method elegantly fixes the problem of finding the cluster after the particle has wandered away. However, a big fractal cluster has big holes, and for large clusters most of the computer time is spent walking within the holes without finding a place to stick. If we knew how close the nearest point on the cluster was we could, as above, take random walk steps of this distance. However, we need to know, at any time, how far away the particle is from the irregular cluster. A naive search would not be efficient.

The Brady–Ball method [30] is an extremely efficient way to give a quick lower bound for the distance to the cluster, and, as a bonus provides a fast mechanism to locate the nearby particles. One modern implementation was made by my student, Ellak Somfai [51]. Here is how he explains what he did.

The cluster is put on an adaptively refined square mesh. That is, the whole cluster is covered with a square—which we call a map. The square is subdivided into four smaller squares and each is further divided, but only if the cluster is ‘sufficiently close’ to it. The rule is as follows: *each map is subdivided if and only if the cluster is closer to it than half of the side of the square.* The subdivision continues only up to a predefined maximum depth, k_{\max} . In our simulations the smallest maps are five to ten particle diameters. All particles of the cluster will be in one of the smallest (deepest) maps: a list of the particles is attached to these maps.

As the cluster grows the maps are updated. Each time a particle is added to a previously empty smallest map, the neighbouring maps (on all levels) are checked

to see whether they satisfy the rule. If not, they are subdivided until they do. When a walker lands somewhere, we find the deepest map containing the point. If this map is not at k_{\max} , then the particle is far away from any matter, and half the side of the map is a lower estimate of the walker's distance from the cluster. If, on the other hand, the particle lands in a map of depth k_{\max} , then it is close to the cluster. The particle lists of the map and of the neighbouring smallest size maps can be checked to calculate the exact distance from the cluster. Either way, the particle is enclosed in an empty circle of known radius and can be brought to the perimeter in one step.

Empirically, the efficiency of this scheme is quite impressive: the computational time, T , for an N particle cluster obeys $T \sim N^{1.1}$, and the memory is linear in N . The cluster of figure 3 was made using this algorithm.

8. Some theory

The DLA model is almost unique among well-studied statistical models in that it has not yielded to analysis. We have no completely acceptable 'theory of DLA': that is, the only satisfactory way to find the fractal dimension, the multifractal scaling spectrum, the crossovers and the overall shape of the cluster is to do a simulation, although some progress has been made on all of these questions. This challenge has intrigued the community of statistical physicists, and very considerable efforts have been made. In this section I will invoke an author's privilege to talk about my favourite approaches, and not even try to do justice to all of the very interesting techniques that have been suggested.

8.1. Estimates and bounds

There are a number of rather simple estimates we can make for the DLA process which put constraints on an eventual theory. They also make it clear why the problem is so hard.

8.1.1. *A lower bound on the fractal dimension.* As a first example, we can rather easily show that the fractal dimension of DLA never becomes independent of the dimension of space, d . This is in contrast to the usual situation in field theory, or in more conventional (equilibrium) macromolecule models where there is an 'upper critical dimension' above which the fractal dimension is independent of the spatial dimension. For example, a self-avoiding random walk is a two-dimensional fractal for all dimensions above four. A simple way to see this is to point out that if a random walker could penetrate a DLA cluster, it would 'fill up' and become dense. However, if $D \ll d$ this is exactly what would happen because then for an aggregate

whose scale is R , the effective number density $\rho \propto R^D/R^d$ would be small—the aggregate would be very wispy and tenuous, and walkers would penetrate and increase ρ instead of sticking on the outside and increasing R .

To be precise, we can estimate the number of intersections of the track of the random walk with a cluster by noting that the number of steps the walker takes inside the aggregate is $\propto R^2$. The number of intersections of this track with the fractal aggregate is then of order ρR^2 . If this is to be large enough to keep the walker on the edge, we need $R^2 \rho / R^d \propto R^{D+2-d}$ to grow with R . Therefore, a self-consistent D must satisfy $D+2-d \geq 0$, that is $D \geq d-2$. It is not difficult [52] to make this bound tighter and show that $D \geq d-1$. This bound on the fractal dimension has been checked up to $d=6$ [50]. Thus the favourite technique in field theory and phase transition theory, an expansion about an upper critical dimension, simply will not work for DLA.

8.1.2. *Singularities at the tips.* We pointed out above that the growth probability on the surface of a DLA aggregate is a wildly varying function which has strong singularities near all of the many sharp tips of the cluster. A theoretical framework for this array of singularities was given by Turkevich and Scher [53]. They pointed out that the solution of the Laplace equation near a sharp point naturally gives rise to a singularity in $\partial u / \partial n_s$: the fact that sharp points on grounded conductors have large electric fields is the essential point of the ordinary theory of the lightning rod. A DLA cluster has many sharp points. For example, if we consider the leading tip (which grows fastest) on a DLA cluster to be more or less a wedge of included angle β , then in two dimensions standard electrostatic theory gives:

$$\begin{aligned} u &\propto X^{\alpha_m - 1}, \\ \alpha_m &= \pi / [2\pi - \beta]. \end{aligned} \quad (17)$$

Here X is the distance of the observation point from the tip. For example, for a flat conductor, $\beta = 0$, we have an inverse square-root singularity. Now suppose we integrate equation (17) over some small distance near the tip. The total probability, which gives the tip velocity, is then: $v \propto (a/R)^{\alpha_m}$, where we have made the natural assumption that the probability must be made dimensionless by the only available length in the problem, the radius, R . Suppose we imagine adding particles at some constant rate. Then

$$\begin{aligned} dM/dt &= (dM/dR)(dR/dt) \\ &\propto R^{D-1}v \\ &\propto R^{D-1+\alpha_m}. \end{aligned} \quad (18)$$

Since dM/dt is constant we conclude that the singularity at the fastest growing tip obeys

$$D = \alpha_m + 1. \tag{19}$$

Thus an arbitrary ‘wedge angle’ gives rise naturally to a non-integer dimension for a DLA cluster. Unfortunately, no one has yet given a geometrical argument for the average angle at the leading tip. Working backwards from equation (20) using $D = 1.7$ gives a β about equal to the interior angle of a regular pentagon. Further, α_m should be the top of the $f(\alpha)$ curve mentioned above. This relation seems to be correct [41]. A systematic way to find D itself and the rest of the spectrum is not evident proceeding in this way.

8.2. Real-space methods and fixed-scale transformations

A natural way to try to renormalize the DLA process is to use the analogue to real-space renormalization techniques from phase transition theory. One of the most interesting attempts along this line is due to Barker and Ball [54]. In this technique the central object is the renormalized effective noise amplitude. They found that it slowly drifts to a fixed point, and they investigated the role of anisotropy.

Pietronero and collaborators developed a related scheme called the method of fixed-scale transformations. It is a real-space method where a small system at one scale is solved essentially exactly, and the behaviour at the next

coarse-grained scale estimated by assuming that there is a scale-invariant dynamics and estimating the parameters from the fixed-scale solution. This method has been reviewed extensively recently [55], and I refer the reader to this article.

8.3. Branched-growth theory

Halsey and Leibig [56] have given a theory which focuses on tip-splitting and the subsequent competition of the daughter tips as the key to the DLA process. In the theory a tip-splitting process is assumed to occur in such a way that the two daughters are slightly different in their initial numbers and probability of growth. These numbers are taken as random initial conditions and averaged over — the idea is that they depend on microscopic details which are decoupled from the subsequent scaling. Then the branch develops so that one or the other of the daughters wins and totally screens the other.

The two variables in the treatment are the probability of growth of the first daughter branch normalized by the total probability for the branch to grow, $x = p_1/p_b$, and the number in the first divided by the total number, $y = n_1/n_b$. All of the ‘initial conditions’ of the last paragraph are near $x = 0.5$, $y = 0.5$ since the branches start out exactly equivalent and microscopic noise perturbs them a bit. As n_b increases with growth, the noise is neglected, and the system flows in the phase plane (x, y) toward the final points where one daughter wins the competition and the other dies, $(0, 0)$ or $(1, 1)$. In the language of phase plane analysis, the




Figure 9. The Hastings–Levitov mapping. The map Z_N takes the unit circle in the w plane to an N -particle cluster in the z plane. The map ϕ transforms the circle to a circle with a bump. Thus $Z_{N+1} = N(\phi(w))$ is the next step in the mapping. The size of the bump must be adjusted as explained in the text.




Figure 10. A cluster generated by the Hastings–Levitov method. The outline is the image of the boundary of the unit circle under the map Z , and the positions where particles were added are shown.

daughter branches flow along the unstable manifold of the saddle point at $(1/2, 1/2)$.

The main work in the theory is to figure out an equation of motion for the flow, i.e. $dx/d(\ln n_b)$, and thus the form of the unstable manifold. This is done using a series of clever approximations for the development of the probability of growth on the surface and, finally, the numerical solution of a nasty nonlinear equation.

There are some very nice features in this approach, and it gives a reasonable value for the fractal dimension ≈ 1.66 . The scaling law $2D_3 = D_0$ and the Turkevich–Scher scaling law equation (19) are built in exactly. Halsey and collaborators have shown how to get multifractal dimensions and their fluctuations [46] using as input the form of the unstable manifold. However, there seem to me to be several problems, the main ones being that the interaction of different branches is not treated, and that this is a ‘one-shot’ approximation with no systematic way to improve it.

8.4. Conformal maps

The final attempt at a theory that I will discuss is much less well developed than the two others mentioned above, but it seems to me to be very promising, and to provide remarkable insights into the growth process which could lead to an analytic theory. It was invented by Hastings and Levitov [57] and developed in several recent papers [58–60].

The conformal map method uses the Laplacian growth version of DLA. It exploits the fact that a very convenient way to solve the Laplace equation is to use the classic method of conformal mapping [61]. Take the cluster to be a grounded conductor in the complex z plane, with a probability to grow at a point on its surface proportional to the charge there: $|\nabla u|$, where u is the potential with boundary conditions of unit flux at infinity and $u = 0$ on the surface. We can construct a complex potential such that $\text{Re}[\Psi(z)] = u$. If the cluster were merely a two-dimensional disc of radius r_0 then, from elementary electrostatics in two dimensions we have $\Psi(z) = \ln(z/r_0)$. It is convenient in what follows to define a new function, $h(z) = e^\Psi$. Clearly h is linear for large z for any cluster (since any shape looks like a disc from far away). Also $|h| = 1$ on the cluster, because $\text{Re}(\Psi) = 0$ there. Now we define a new (mathematical) plane, w such that $w = h(z)$; h is a *conformal map* from the exterior of the cluster to the exterior of the unit circle. Its inverse function $Z(w)$ is a conformal map which takes the exterior of the unit circle in the w plane to the exterior of the DLA cluster in the z plane, and obeys $Z \sim r_0 w$, for large $|w|$. We will focus attention on Z .

Now we can solve the Laplace equation in the w plane trivially, and transform back to the z plane. On the unit circle $|\nabla u| \propto |d\Psi/dz| = |dh/dz| = 1/|dZ/dw|$ so that the

growth probability is $1/|Z'|$. Intervals $d\theta$ on the unit circle in the w plane correspond to intervals of arc length ds with equal growth probability in the z plane. Thus the image of the growing tips is most of the circle and the fiords occupy a tiny region. The most useful characterization of the map is the Laurent series:

$$Z(w) = r_0 w + \sum_{k=0}^{\infty} A_k / w^k. \quad (20)$$

The coefficients, A_k , contain lots of very interesting information.

All of this would be of no real interest if we could not find the map. However we have two independent ways to do this. The first is the original Hastings–Levitov recursive method. Suppose we know the map at a certain stage of growth: Z_{N-1} . Now we add a ‘particle’ by adding to the outline a bump of area λ_0 in the z plane. Z_{N-1} will transform the unit circle to a circle with a bump whose area is transformed: $\lambda_N = \lambda_0 / |dZ_{N-1}/dw|^2$; see figure 9. The bump can be taken to be of the same fixed shape and to occur at an angle θ_N which is uniformly distributed in the mathematical plane (cf. the previous paragraph). Suppose we denote by $\phi_{\lambda_N, \theta_N}(w)$ the function that maps the unit circle to the circle with a bump. The recursion relation for $Z_N(w)$ is given by (see figure 9):

$$Z_N(w) = Z_{N-1}(\phi_{\lambda_N, \theta_N}(w)). \quad (21)$$

In practice, we choose a random angle θ_N , find λ_N , and generate any level of map. This is a nice method, but slow: it takes of order N^2 steps to generate an N -particle cluster (see figure 10).

The other way to get the map [60] is to use the original DLA model to grow the aggregate. Then we stop at stage N and record where M random walkers would attach to the cluster. This gives a set of points z_m . These are at angle $\theta_m \approx 2\pi(m/M)$ in the w plane, since we are sampling the probability, and equal increments of probability correspond to equal increments of θ . The A_k are the Fourier coefficients of the function $z(\theta_m)$. We have used this method to find the map up to $N \approx 10^6$.

So far we have a numerical procedure. However, it gives rise to many insights and leads to analytical progress. For example, in very recent work [60] we have shown that the growth zone of Pleshke and Racz [47] must scale with the radius using only some weak assumptions about the multifractal spectrum. We used the fact that the growth zone can be related to the amplitude, $|A_k|$, of the Laurent coefficients. The crossovers that are so puzzling can be related to the slow crossover of the lowest multipole moments of the probability to their asymptotic scaling behaviour. It is very interesting to note that the phase of the A_k ’s encode the multifractal spectrum itself. This area is in rapid evolution.

9. Summary

In this informal review I have tried to give the reader some flavour for the progress that has been made in studies of the DLA model. There is an enormous literature on the subject, but, oddly enough, many of the most useful insights have not penetrated to the physics community at large. This is unfortunate. Many of the insights and methods (particularly the numerical techniques) should be useful outside of the DLA area, but are all too often not recognized, and have to be rediscovered. My purpose in writing this is to try, in part, to correct that situation.

Acknowledgements

I have had useful discussions about the subjects described here with hundreds of people, and it would be impractical to list them. Nevertheless, I am very grateful for their insights. My research in this area has been supported by several sources over the years, and at present by DOE grant DEFG-02-95ER-45546.

References

- [1] Witten, T. A., and Sander, L. M., 1981, *Phys. Rev. Lett.*, 47, 1400.
- [2] Forrest, S. R., and Witten, T. A., 1979, *J. Phys. A Math. Gen.*, 12, L109.
- [3] Brady, R. M., and Ball, R. C., 1984, *Nature*, 309, 225.
- [4] Meakin, P., 1998, *Fractals, Scaling, and Growth Far From Equilibrium* (Cambridge: Cambridge University Press).
- [5] Witten, T. A., and Sander, L. M., 1983, *Phys. Rev. B*, 27, 2586.
- [6] Mullins, W. W., and Sekerka, R. F., 1963, *J. appl. Phys.*, 34, 323.
- [7] Langer, J., 1980, *Rev. mod. Phys.*, 52, 1.
- [8] Bales, G. S., and Zangwill, A., 1990, *Phys. Rev. B condens. Matter*, 41, 5500.
- [9] Tsui, F., Wellman, J., Uher, C., and Clarke, R., 1996, *Phys. Rev. Lett.*, 76, 3164.
- [10] Kunkel, R., Poelsma, B., Verheij, L. L., and Comsa, G., 1990, *Phys. Rev. Lett.*, 65, 733.
- [11] Ehrlich, G., and Hudda, F. G., 1966, *J. chem. Phys.*, 44, 1039; Schwoebel, R. L., and Shipsey, E. J., 1966, *J. appl. Phys.*, 37, 3682.
- [12] Saffman, P. G., and Taylor, G. I., 1958, *Proc. Roy. Soc. A*, 245, 312.
- [13] Paterson, L., 1984, *Phys. Rev. Lett.*, 52, 1621.
- [14] Matsushita, M., Sano, M., Hayakawa, Y., Honjo, H., and Sawada, Y., 1984, *Phys. Rev. Lett.*, 53, 286.
- [15] Sawada, Y., Dougherty, A., and Gollub, J. P., 1986, *Phys. Rev. Lett.*, 56, 1260.
- [16] Grier, D., Ben-Jacob, E., Clarke, R., and Sander, L. M., 1986, *Phys. Rev. Lett.*, 56, 1264; Grier, D., Kessler, D., and Sander, L. M., 1987, *Phys. Rev. Lett.*, 59, 2315.
- [17] For example: Garik, P., Barkley, D., Ben-Jacob, E., Bochner, E., Broxholm, N., Miller, B., Orr, B., and Zamir, R., 1989, *Phys. Rev. Lett.*, 62, 2703; Fleury, V., Chazalviel, J.-N., Rosso, M., and Sapoval, B., 1990, *J. Electroanal. Chem.* 290, 249; Argoul, F., Arneodo, A., Elezgaray, J., and Kuhn, A., 1997, *Fractals*, 5, 75.
- [18] Fujikawa, H., and Matsushita, M., 1989, *J. Phys. Soc. Jpn.*, 58, 3875.
- [19] Ben-Jacob, E., 1997, *Contemp. Phys.*, 38, 205.
- [20] The diffusive instability in this case is unexpectedly subtle, and has led to an interesting development in mathematical physics. See Kessler, D. A., Ner, Z., and Sander, L. M., 1988, *Phys. Rev. E*, 58, 107; Kessler, D. A., and Levine, H., 1998, *Nature*, 394, 556.
- [21] Family, F., Masters, B. R., and Platt, D. E., 1989, *Physica D*, 38, 98.
- [22] Fleury, V., and Schwartz, L., 1999, *Fractals*, 7, 33.
- [23] Kaandorp, J. A., 1993, *Growth Patterns in Physical Sciences and Biology*, edited by J. M. Garcia-Ruiz, E. Louis, P. Meakin and L. M. Sander (New York: Plenum), pp. 331–340.
- [24] Ben-Jacob, E., 1993, *Contemp. Phys.*, 34, 247.
- [25] Pelce, P. (ed.), 1988, *Dynamics of Curved Fronts* (Boston: Academic Press).
- [26] Couder, Y., Argoul, F., Arneodo, A., Maurer, J., and Rabaud, M., 1990, *Phys. Rev. A*, 42, 3499.
- [27] Tu, Y., and Levine, H., 1995, *Phys. Rev. E*, 52, 5134.
- [28] Ball, R., Nauenberg, M., and Witten, T. A., 1984, *Phys. Rev. A*, 29, 2017; Nauenberg, M., and Sander, L. M., 1984, *Physica*, 123A, 360.
- [29] Dawson, S. P., and Mineev-Weinstein, M., 1998, *Phys. Rev. E*, 57, 3063.
- [30] Ball, R. C., and Brady, R. M., 1985, *J. Phys. A Math. Gen.*, 18, L809.
- [31] Family, F., and Hentschel, G., 1987, *Faraday Discuss. chem. Soc.*, 83, 139.
- [32] Ball, R. C., Brady, R. M., Rossi, G., and Thompson, B. R., 1985, *Phys. Rev. Lett.*, 55, 1406.
- [33] Niemeyer, L., Pietronero, L., and Wiesmann, H. J., 1984, *Phys. Rev. Lett.*, 52, 1033.
- [34] Kadanoff, L. P., 1985, *J. stat. Phys.*, 39, 267.
- [35] Szep, J., Cserti, J., and Kertesz, J., 1985, *J. Phys. A Math. Gen.*, 18, L413.
- [36] Kolb, M., Botet, R., and Jullien, R., 1983, *Phys. Rev. Lett.*, 51, 1123; Meakin, P., 1983, *Phys. Rev. Lett.*, 51, 1119.
- [37] Kolb, M., 1985, *J. Phys. Lett.*, 46, L631.
- [38] Mandelbrot, B. B., Kaufman, H., Vespignani, A., Yekutieli, I., and Lam, C. H., 1995, *Europhys. Lett.*, 29, 599; Mandelbrot, B. B., Vespignani, A., and Kaufman, H., 1995, *Europhys. Lett.*, 32, 199.
- [39] Mandelbrot, B. B., 1982, *The Fractal Geometry of Nature* (W. H. Freeman).
- [40] Tolman, S., and Meakin, P., 1989, *Phys. Rev. A*, 40, 428.
- [41] Halsey, T. C., Jensen, M. H., Kadanoff, L. P., Procaccia, I., and Shraiman, B. I., 1986, *Phys. Rev. A*, 33, 1141; Hayakawa, Y., Sato, S., and Matsushita, M., 1987, *Phys. Rev. A*, 36, 1963; Ball, R. C., and Rath Spivack, O., 1990, *J. Phys. A*, 23, 5295.
- [42] Mandelbrot, B. B., 1974, *J. fluid Mech.*, 62, 331.
- [43] Vicsek, T., Family, F., and Meakin, P., 1990, *Europhys. Lett.*, 12, 217; Vicsek, T., 1990, *Physica A*, 168, 490.
- [44] Schwarzer, S., Wolf, M., Havlin, S., Meakin, P., and Stanley, H. E., 1992, *Phys. Rev. A*, 46, R3016.
- [45] Makarov, N. G., 1985, *Proc. Lond. math. Soc.*, 51, 369.
- [46] Halsey, T. C., Duplantier, B., and Honda, K., 1997, *Phys. Rev. Lett.*, 78, 1719.
- [47] Plischke, M., and Racz, Z., 1984, *Phys. Rev. Lett.*, 53, 415.
- [48] Li, G., Sander, L. M., and Meakin, P., 1989, *Phys. Rev. Lett.*, 63, 1322.
- [49] Meakin, P., Ball, R. C., Ramanlal, P., and Sander, L. M., 1987, *Phys. Rev. A*, 35, 5233.
- [50] Sander, E., Sander, L. M., and Ziff, R., 1994, *Comput. Phys.*, 8, 420.
- [51] Somfai, E., private communication.
- [52] Ball, R. C., and Witten, T. A., 1984, *Phys. Rev. A*, 29, 2966.
- [53] Turkevich, L., and Scher, H., 1985, *Phys. Rev. Lett.*, 55, 1026.
- [54] Barker, P. W., and Ball, R. C., 1990, *Phys. Rev. A*, 42, 6289.
- [55] Erzan, A., Pietronero, L., and Vespignani, A., 1995, *Rev. mod. Phys.*, 67, 545.
- [56] Halsey, T. C., and Leibig, M., 1992, *Phys. Rev. A*, 46, 7793; Halsey, T., 1994, *Phys. Rev. Lett.*, 72, 1228.
- [57] Hastings, M. B., and Levitov, L. S., 1998, *Physica D*, 116, 244.
- [58] Hastings, M. B., 1997, *Phys. Rev. E*, 55, 135.
- [59] Davidovitch, B., Hentschel, H. G. E., Olami, Z., Procaccia, I., Sander, L. M., and Somfai, E., 1999, *Phys. Rev. E*, 59, 1368.

- [60] Somfai, E., Sander, L. M., and Ball, R. C., 1999, *Phys. Rev. Lett.*, 83, 5523.
- [61] Churchill, R. V., 1990, *Complex Variables and Applications* (New York: McGraw-Hill).

Leonard M. Sander is Professor of Physics at the University of Michigan. He did his undergraduate work at Washington University in St. Louis, and his graduate work at the University of California, Berkeley where he was trained as a solid state theorist. He remained in solid state and condensed matter until 1981 when he stumbled upon the topic of this article. Since then he has been working in statistical physics and pattern formation.