

**MODELING SNOW CRYSTAL GROWTH II:
A mesoscopic lattice map with plausible dynamics**

JANKO GRAVNER
Mathematics Department
University of California
Davis, CA 95616
e-mail: gravner@math.ucdavis.edu

DAVID GRIFFEATH
Department of Mathematics
University of Wisconsin
Madison, WI 53706
e-mail: griffeat@math.wisc.edu

(Revised version, September 2007)

Abstract. We present a local lattice model for the evolution of snow crystals that combines diffusion-limited aggregation with anisotropic attachment kinetics and an idealized quasi-liquid layer. Despite a voluminous modeling literature, this is apparently the first approach that successfully captures the essential form of core and tip instabilities, branch faceting, and other aspects of real snow crystal growth. As parameters are varied, our nearest neighbor system reproduces the basic features of most of the observed varieties of planar snowflakes and offers new insights into their genesis.

2000 *Mathematics Subject Classification.* Primary 82C24. Secondary 35R35, 60K35.

Keywords: Coupled lattice map, crystal growth, diffusion-limited aggregation, Stefan problem.

Acknowledgments. Many thanks to Ken Libbrecht for his inspirational collection of photomicrographs, and for sharing with us his published and unpublished research on the physics of snowflakes. We also thank Cathy Loeb for editorial help, and Natasha Didenko for programming assistance.

Support. JG was partially supported by NSF grant DMS-0204376 and the Republic of Slovenia's Ministry of Science program P1-285. DG was partially supported by NSF grant DMS-0204018.

1 Introduction

Since the pioneering observations of Kepler [Kep], Descartes [Des], and Hooke [Hoo] four centuries ago, scientists have been fascinated by snow crystals, commonly known as *snowflakes*, with their mysterious mix of complexity and symmetry. Early on, their amazing diversity reminiscent of life forms was often attributed to a supreme being. Over the past century, physicists and mathematicians have devised and studied various theoretical models in an attempt to illuminate the mystery. Our goal in this paper is to introduce a new algorithm for the growth of planar snow crystals, based on physical principles, that captures several essential features of their development. As far as we are aware, this is the first computationally feasible approach that exhibits the complete range of plate and dendrite varieties occurring in nature.

To set the stage for our model, let us review briefly what is known about snowflake formation in terms of weather and chemistry (cf. [Stu], [Hob]) and summarize the empirical record and classification of their structure:

Meteorology

Three forms of water coexist in saturated clouds below 0°C: supercooled liquid water, vapor, and ice. Initially, liquid and vapor predominate. The transformation to ice typically begins at colder temperatures near the tops of clouds. Crystal nucleation starts with minuscule “dust” particles such as leaf or clay material consisting of as few as a thousand molecules. Due to heat of condensation, water droplets form around these particles and then gradually freeze to initiate ice crystal growth. Additional liquid water can freeze to the nucleating snowflake by contact, but under ideal conditions the liquid in the cloud first evaporates, due to a pressure gradient, so that growth occurs almost entirely by condensation of vapor directly onto the crystal surface. This process exhibits a complex dependence on humidity and temperature, with lesser effects from other physical parameters, yielding especially rapid formation of crystals in supersaturated environments at about -15°C . At first, snowflakes wander very slowly through the clouds. Later, as they gain mass, they tumble downward, eventually falling to earth.

Crystallography

The crystal structure of ordinary ice, I_h , may be viewed roughly as a stack of two-dimensional lattices of hexagonal cells. Reflecting this structure, an ice crystal that nucleates from an irregular seed first convexifies and regularizes, then typically grows to form a right prism on the order of a micron in size, with two hexagonal *basal* facets and six rectangular *prism* facets. Qualitative differentiation of form takes place as growth continues from micron scale to a final size typically a few millimeters across. Throughout its development, the molecular crystalline structure of any snowflake is the I_h lattice.

Photomicrography

Since the late 1800s, many thousands of natural snowflakes have been photographed and organized by type. The most extensive early collection was assembled by Vermont farmer Wilson Bentley [BH, Ben2]. The majority of these are beautiful, almost perfectly symmetric planar crystals of the kind we seek to model. Indeed, Bentley is widely credited with establishing idealized snowflake designs as icons of winter time. In his authoritative scientific study from the mid-1900s, Nakaya [Nak] provided a much more diverse collection featuring three-dimensional crystals such

as needles, exotic *capped columns* (or *tsuzumi*), and a great many irregular, highly asymmetric specimens. State-of-the-art photomicrographs of both two-dimensional and three-dimensional snowflakes may be found in the beautiful popular science book [LR] and on Libbrecht's superb Web site [Lib4].

Planar classification

As the database of snowflake specimens has developed, so have attempts to classify them. Early proposals of [She], [BH], and [Nak] were consolidated in the Magono-Lee classification [ML] of all snow crystals into 80 types (one of them called Miscellaneous). A new, simplified classification by Libbrecht [Lib5] reduces the number of types to 35. Standard planar crystals comprise 13 of the Magono-Lee types and 6 of the Libbrecht types, as shown in Fig. 1. Note that some of the former are hybrids of the latter. Such schemes are inevitably even less precise than plant taxonomy since there would seem to be intermediate examples between any two distinct forms. For instance, the dividing line between a fern-like crystal and a stellar dendrite is inevitably arbitrary.

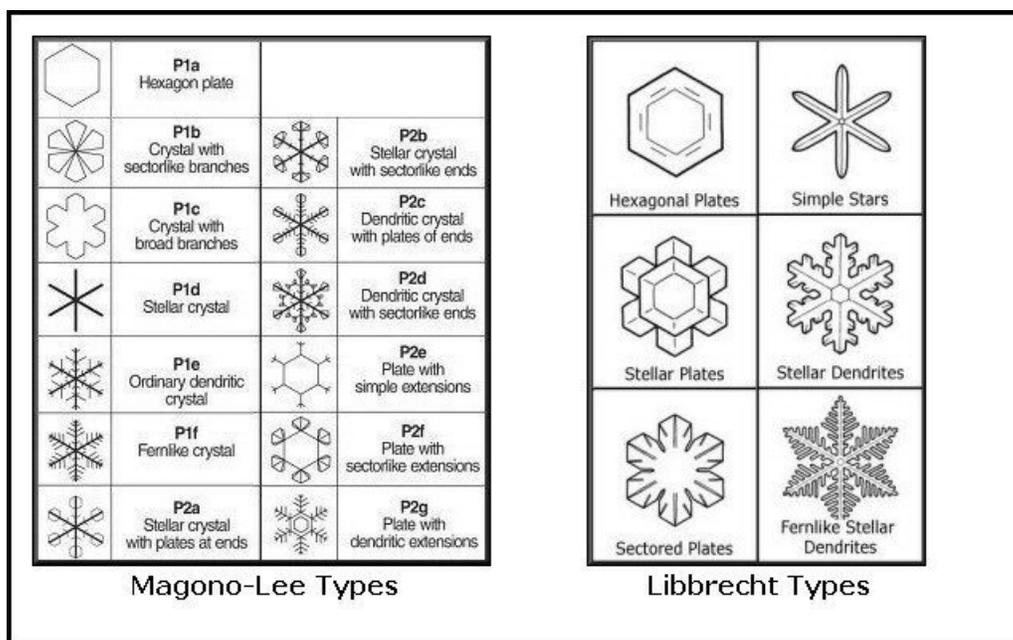


Figure 1. Classification of planar snow crystals ([ML],[Lib5])

Despite the abundance of final product, there is still scant knowledge of how snow crystals evolve. It is seemingly impossible to track their natural development, so scientists have attempted to grow them artificially. Nakaya [Nak] was the first to do so successfully, on the end of a rabbit hair. Since then various ingenious techniques have been developed, e.g., using cloud tunnels [FT] or replacing the rabbit hair with an ice needle [Lib4]. Despite such painstaking efforts, the more exotic types of snowflake are difficult to reproduce in the lab, and many aspects of the growth process remain poorly understood. Much of the research to date has attempted to decipher the elaborate markings that both Bentley [Ben1] and Nakaya [Nak] referred to as *hieroglyphs*, although *fossils* seems a more appropriate metaphor.

With the advent of high-speed computing, mathematical modeling is an increasingly promising method of studying snow crystal growth. Of course, there is a tradeoff between physical detail and efficient simulation, so designing a realistic yet tractable algorithm is a balancing act. This article introduces a reasonably simple, two-dimensional lattice map that captures several basic features of the thin, faceted, and dendritic crystals observed at temperatures between about -10°C and -20°C . Figure 2 shows two real snowflakes – a sectored plate and a stellar dendrite – and corresponding simulated crystals generated by our model. The photomicrographs by Ken Libbrecht are two of four chosen by the US Postal Service for a 2006 set of commemorative stamps. Our simulations here and throughout the paper are rendered in gradient palettes on and off the snowflakes, with darker shades corresponding to more mass. The shading in Libbrecht’s images results from color filtering. As one indication of our model’s verisimilitude, it evidently does a good job of recreating the characteristic *ridges* of planar snow crystals dubbed *axes of crystallization* a century ago in [She].

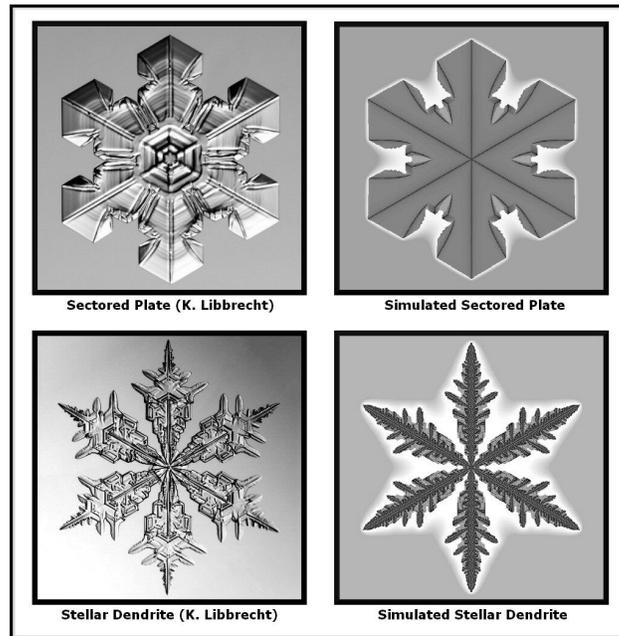


Figure 2. Comparison of natural and simulated crystals

The remainder of the paper is organized as follows: The next section describes the observed structure, or *morphology*, of planar snow crystals in some detail. This background is important for two reasons. First, it clarifies the underlying assumptions of our algorithm that determine its scope and limitations. Second, it illustrates some key features of real growth that the model captures rather well. Section 3 briefly discusses the role of randomness in snowflakes since a seeming paradox between symmetry and complex disorder lies at the heart of their mystery. Next, we review the prior modeling literature in Section 4. There have been dozens of attempts at snow crystal simulation over the past 25 years, from ad hoc mathematical constructions to Monte Carlo methods, partial differential equations (PDE), cellular automata, and coupled lattice maps. We discuss the pros and cons of the most successful models and indicate how our approach synthesizes their strong points.

Section 5 presents the details of our snowflake simulator. One main goal of the present project is to frame the aggregate dynamics of snow crystals in terms of *hydrodynamics of particle systems*, an area of contemporary research at the interface between stochastic processes, statistical physics, and PDE. Our model suggests that random processes on the microscopic (nanometer) scale become quasi-deterministic on the mesoscopic (micron) scale, leading to complex symmetric structure on the macroscopic (millimeter) scale subject to a small level of noise. We motivate the simulator by analogy with a much simpler growth model known as *internal diffusion-limited aggregation (IDLA)*, for which the hydrodynamic limit has been rigorously established. Drawing on the discussion of Section 2, we also explain how attachment kinetics at the crystal surface informs our simulator in a crucial way. Finally, we offer a glimpse of the role of system parameters in determining the morphology of our artificial snowflakes.

Section 6 offers additional case studies. They focus on the central issue of environmental variability, e.g., gradual or abrupt changes in weather, the fall to earth, and random fluctuations in the vapor. How necessary are such inhomogeneities for the full range of snowflake types? How much about a trajectory through the clouds can be deduced from the final pattern? We offer speculative replies to these questions that are somewhat at odds with conventional wisdom. Section 7 concludes the paper with a brief discussion of unresolved puzzles and some future directions suggested by our work.

2 Morphology of planar snow crystals

As mentioned earlier, nucleating snowflakes typically form hexagonal prisms once they reach a few microns in diameter. While this geometry reflects the underlying I_h lattice, recent research (e.g., [BBS], [Wet]) makes it clear that subtle molecular surface processes underlie the initial stages of growth and that early aberrations in form can have a lasting effect on a crystal's shape. Our model *assumes* a micron-scale regular prism as its starting point. The different local geometry of basal and prism facets leads to different growth rates in the corresponding lattice directions. For reasons still poorly understood, these two rates exhibit a sensitive, nonmonotone dependence on temperature and supersaturation. Just below freezing, and between -10°C and -20°C , the six prism facets advance much more rapidly than the two basal facets, leading to “classic” snowflakes that are effectively planar. At other subzero temperatures, the opposite happens, producing slender three-dimensional crystals such as columns and needles. To avoid the physical complications and computational demands of three dimensions, we restrict our attention to the planar regime and develop an essentially two-dimensional model with variable profiles of *crystal mass*, *diffusing mass*, and *boundary mass*; experimental evidence and an effect known as the Schwoebel-Ehrlich barrier support this simplification (cf. [Lib3]).

Although snowflakes formed between 0°C and -4°C represent many of the types of Fig. 1, three substantial complications arise in that temperature range:

- *Roughening transition.* Just below 0°C , the characteristic facets of planar crystals morph into dome-shaped boundaries; this is presumably a roughening transition due to significant melting.
- *Sublimation.* Close to 0°C , snowflakes are especially susceptible to sublimation, the process whereby ice returns directly to vapor. This conversion erodes and smoothes boundaries, significantly degrading the crystal design.

• *Riming and graupel.* Above -10°C , snowflakes are also quite vulnerable to riming, where supercooled water fails to vaporize before reaching the crystal boundary, but rather freezes on contact in the form of droplets. Steady accumulation of rime eventually produces snow pellets known as graupel.

To avoid such complications, thereby improving the prospects for faithful representation of real snowflakes, we further limit our investigation to temperatures between -10°C and -20°C . The celebrated *Nakaya diagram*, introduced in [Nak] and later refined by various researchers, plots the dominant type of planar snowflake as a function of temperature and supersaturation. Fig. 3 is a sketch of the portion relevant to our study.

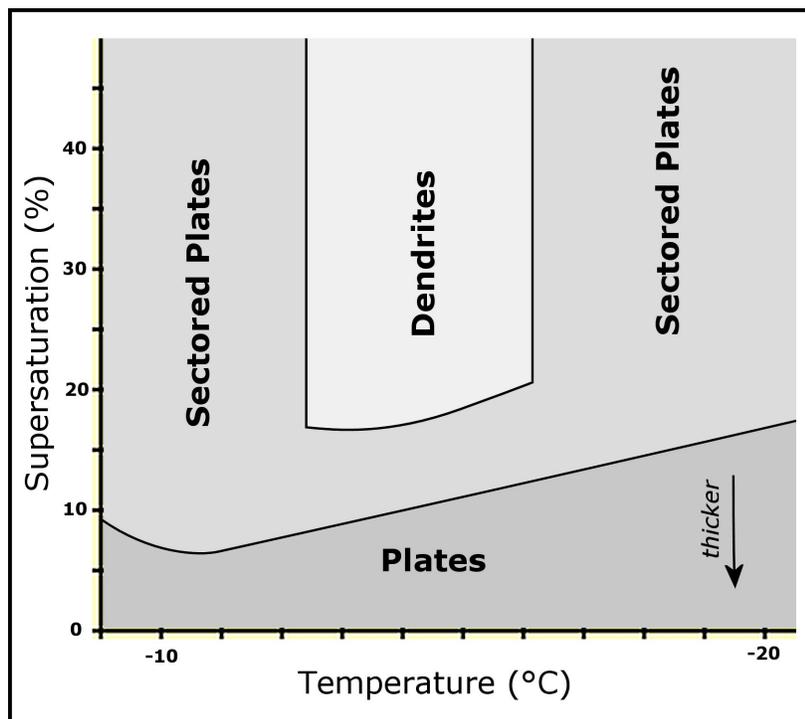


Figure 3. Detail of the Nakaya diagram (-10°C to -20°C)

The prototypical snowflakes we wish to model all take the form of relatively thin hexagonal prisms when they are a few microns across and grow much more quickly on their prism facets than on their basal facets. Research by physicists (cf. [YK], [Lib1], [Lib3] and extensive references therein) has reached consensus that three main effects govern the ensuing development:

- *Diffusion-limited solidification.* As vapor joins the crystal, a surrounding boundary layer is gradually depleted. Consequently, protrusions such as corners and edges are more accessible to the remaining vapor than indentations or the middles of faces. This initiates self-reinforcing feedback that promotes irregular growth and branching.
- *Anisotropic attachment kinetics.* The thermodynamics of how vapor attaches to a snow crystal is still poorly understood, but local geometry is important. Molecules have more difficulty bonding to a relatively flat portion of the boundary than to a rougher concave part. Note that

this effect tends to regularize shape and counteract the previous one, accounting for the facets observed in many full-grown snowflakes.

- *Quasi-liquid layer.* Very near the crystal surface, molecules are bound more tightly than in the surrounding vapor, but not completely constrained to the I_h lattice. The behavior is akin to a hybrid state with dynamic exchange between vapor and ice, and this layer is believed to play an important role in exotic morphology (cf. [Lib1]). Our algorithm includes a separate mass field, which we descriptively call quasi-liquid, to incorporate qualitatively different dynamics within a tight boundary layer around the crystal in a natural (albeit crude) way.

Surface defects known as *microsteps* induce the layer-by-layer faceted growth of many crystals. In planar snowflakes, the microsteps tend to bunch, creating micron-scale macrostep dislocations. It is believed that macrostep dynamics give rise to a series of dynamic instabilities that we will now describe, comparing photomicrograph illustrations with our model simulations (see, e.g., [Fra], [YK] and references therein).

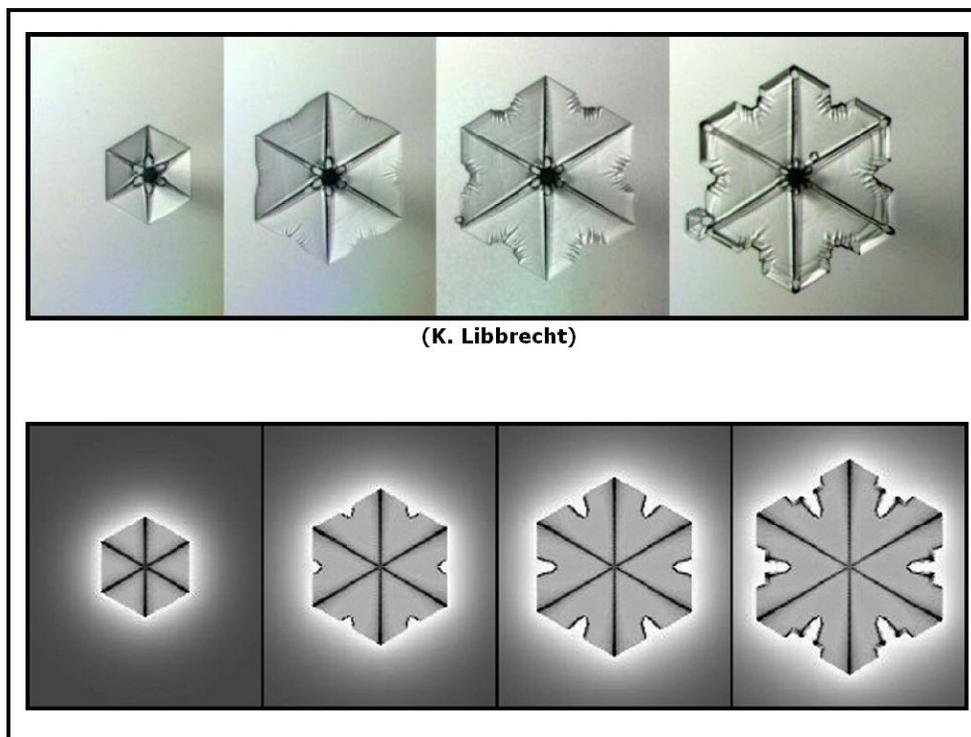


Figure 4. The first instability: in the lab (top) and simulated (bottom)

The first instability

Growth is effectively diffusion-limited once the crystal reaches a few microns in size. Corners of the thin hexagonal prism are most accessible to vapor, so macrosteps develop there first and “run” toward the middles of the six prism facets, depositing new layers of ice. Eventually, since the attachment rate is highest at the tips and lowest in the centers of the faces, an instability occurs when the macrostep chain stops just short of adding a complete layer. (From a theo-

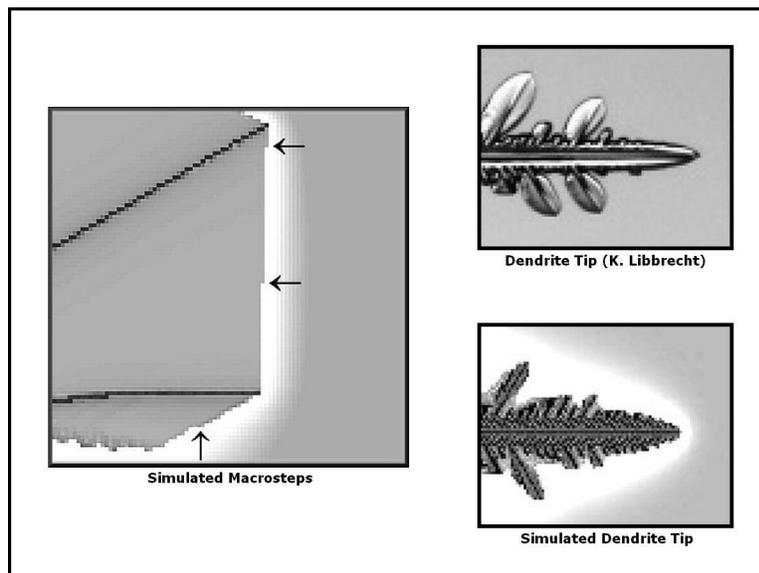


Figure 5. Macrosteps and ridges

retical point of view, this is a variant of the celebrated Mullins-Sekerka instability [MS]; some distinctions from the isotropic model are discussed in [YK].) Further growth reinforces the tips, signaling the emergence of the snowflake’s six main ridges. See actual and simulated stages of the process in Fig. 4. Onset of the first instability varies substantially, happening earliest in rapidly growing and branching dendrites and much later (or not at all) in slowly growing plates. Planar snow crystals generically retain hexagonal geometry at their tips, though this feature can be quite subtle in dendritic cases.

Macrosteps and ridges

Further development of planar snow crystals after the first instability continues to be driven by macrostep dynamics. The left panel of Fig. 5 shows a close-up of half a sectorized plate from the simulated crystal of Fig. 2. Notice three macrosteps marked by arrows, the top two moving downward, away from a principal axis, and the bottom one moving away from the lower corner. Macrosteps are most evident on faceted plate boundaries such as this, but also govern growth and branching of dendrites. The right panels of Fig. 5 compare typical real and simulated growing dendrite tips. Observe the ridges in both versions, on the main branch and also on side branches, as well as the characteristic narrow grooves to either side of the ridges. Note also the 60° angle of the side branches with respect to the main axis.

Tip instability

Some insight into the process of further branching is gained by observing macrostep dynamics near the tip of a sectorized plate such as the one in Fig. 2. Roughly, large-scale faceting lets us “zoom” to a degree not possible for dendrites with our lattice algorithm. The still frames of Fig. 6 compare the onset of a side branch in an artificial sectorized plate photographed in the lab [GN] with corresponding graphics from our simulation. Arrows in the first companion frames show the locations of the tip instability. As with the first instability, branching starts when a

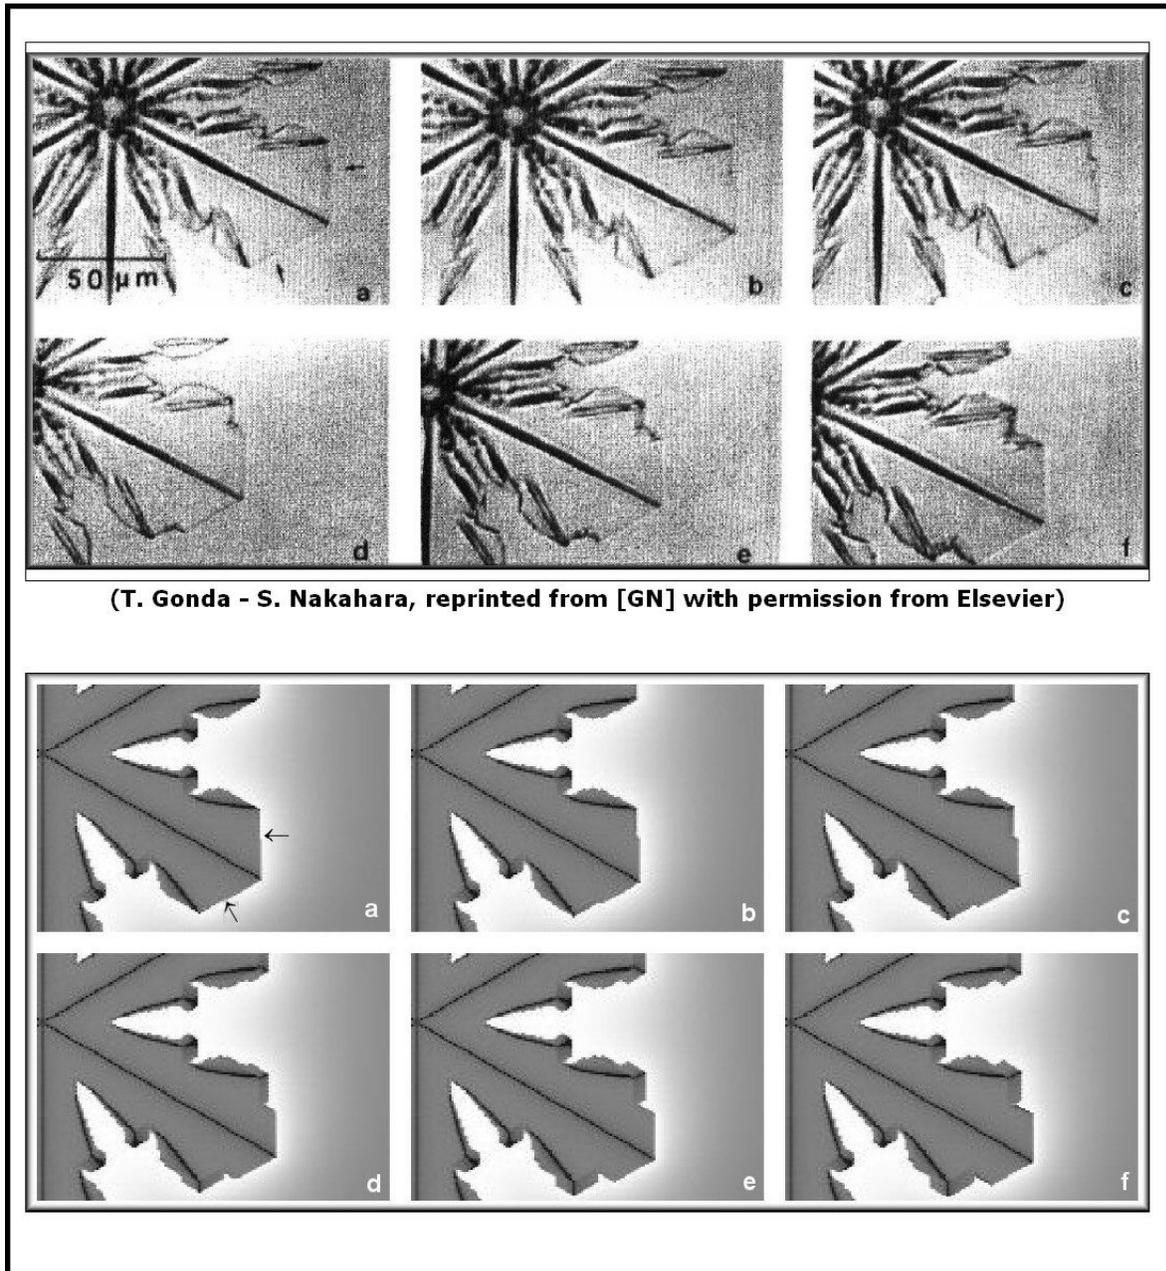


Figure 6. Tip instability: in the lab (top) and simulated (bottom)

macrostep wave stops short of adding a complete layer to the growing facet. Subsequent waves from the main tip reinforce the singularity and allow vapor near the other end of the facet to generate macrosteps emanating from that tip. By the final frames of both the lab and computer-generated crystals, it is clear that a side branch has developed. Although the aspect ratios of the two sequences are somewhat different, their qualitative agreement is striking. We consider this one of the best indicators that our local lattice map captures fundamental aspects of snow crystal growth.

Aftergrowth

Markings on snowflakes often give the impression that first a “backbone” formed, and interstices filled in later to create some of the faceted and web-like features. We have already identified ridge and branching structure created early on and have seen that macrosteps promote faceting, but careful inspection of certain cases strongly suggests an additional mechanism for aftergrowth, even when vapor has apparently been depleted locally by fast-developing portions of the crystal. The idea of multistage growth goes back to Shedd, who devised an elaborate explanation based on suddenly shifting environmental conditions [She]. Recently, an alternative explanation involving unique properties of ice has been proposed in [Lib2] (see also [Lib5]). Snowflakes such as the plate of Fig. 7 suggest that a sufficiently thin planar sheet undergoes a sudden, dramatic increase in growth rate at its edge, requiring very little vapor to spread rapidly. Libbrecht calls this the *knife-edge instability*. We have incorporated a speculative implementation of the effect into our algorithm, as evidenced by the thinner (lighter gray) portions of the simulations in Fig. 2. Ultimately, our rule simply implements the idea that there is some nonmonotonicity in the growth rate of the crystal vs. the density of surrounding vapor. This nonmonotonicity is counterintuitive and not well understood. Indeed, the knife-edge feature of our algorithm may amount to a two-dimensional “fix” for a genuinely three-dimensional phenomenon. But there is some experimental evidence for it, and it has appeared (albeit in a completely different form) in the modeling literature before (e.g., [XAR]).

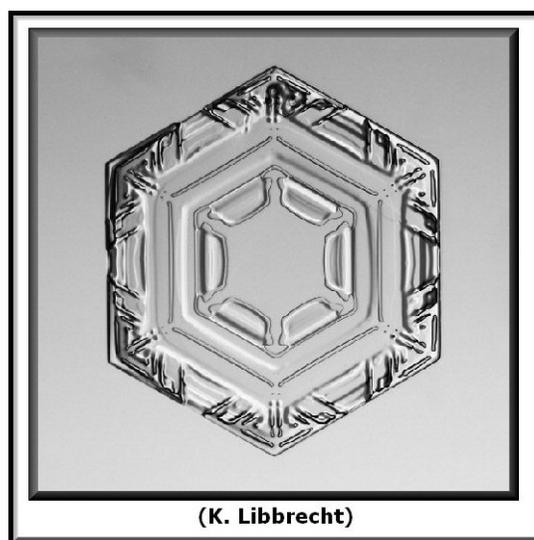


Figure 7. Evidence for aftergrowth

Even near -15°C , under ideal conditions for planar snowflakes, various spatial structures are frequently observed that relate directly to the thickness of the initially formed hexagonal prism (see, e.g., [Lib5]):

- *Double layers.* If there is sufficient distance between the two basal planes of a micron-scale prism, then since its 12 prism edges are most accessible to the surrounding vapor, further diffusion-limited growth may produce twin planar crystals joined by a short central column. Sometimes this process gives the impression of a 12-sided snowflake.
- *Vestigial centers and split crystals.* A related, rather common three-dimensional scenario produces planar spread of one basal facet at the expense of the other, leaving a stunted plate visible at the center of the larger crystal. The sectored plate of Fig. 8 is an example. Hybrid forms also occur, where a subset of the six radial axes grow in one basal plane and the remainder grow in the other, presumably due to subtle fluctuations in vapor supersaturation. (Such crystals sometimes break in two, producing apparently irregular snowflakes with fewer than six axes, and even the occasional symmetric three-sided example.)
- *Hollows and bubbles.* Other subtler three-dimensional artifacts of diffusion-limited growth include hollows, where the centers of the prism facets lag behind their more accessible edges, and bubbles, where such hollows later close up due to melting or other diffusive effects.

Since snow crystals are transparent, it is sometimes challenging to differentiate surface markings from three-dimensional effects of the kinds just described. When comparing our simulations to photomicrographs, the reader should bear in mind that our algorithm only attempts to emulate ideal planar growth, so the above effects are disregarded.

Melting and sublimation are also prevalent effects, even under perfect conditions. Nakaya and others have noted that most of the snowflakes in Bentley's collection [BH] are heavily sublimated. (Bentley also doctored his images in ways that limit their scientific value; see [Nak].) More recent collections such as those of [LR] and [Lib4] provide much more accurate and detailed records of ideal crystal growth, but rounded contours and surface diffusion in many instances indicate ongoing conversion of ice to liquid and vapor as the snowflake fell to earth. Our algorithm assumes that ice is permanent, i.e., solidification is not reversible.

Another complicating factor in modeling snow crystal growth is the role of motion and velocity. In its initial stages, a snowflake moves quite slowly, but as it gains mass, its velocity increases and begins to affect further development. Even in falling to earth, the effect of motion is similar to that of a somewhat higher supersaturation level on a motionless crystal, since the snowflake tends to tumble randomly, and since the diffusion rate in the surrounding vapor is high. Counteracting this tendency is the generally lower supersaturation at lower cloud levels. Still another factor is frequent upward vertical motion in the atmosphere capable of lifting a small crystal to colder regions with higher supersaturation that are more conducive to dendritic growth. Such influences on a snowflake's journey have led many researchers to attribute much of the observed structure to large-scale environmental transitions. [FT] includes an excellent discussion of velocity and temperature, as well as many other factors, in connection with cloud tunnel experiments designed to make snow crystals fall as slowly as possible. To simplify matters and explore which types can arise without changing conditions, our model assumes that a crystal grows from micron scale to its final size at a fixed location within an initially homogeneous field of vapor. Nevertheless, we will see that for a suitable fixed choice of parameters, a great

variety of planar crystal forms develops. By way of contrast, Section 7 includes a few “designer simulations” in which the supersaturation changes gradually or suddenly.

3 The puzzle of symmetric complexity

Fascination with snow crystals, for scientists and lay people alike, derives in large part from the remarkable similarity of their six main branches. Each is marked with a nearly identical mix of familiar geometrical forms and strange, unpredictable bumps and side branches. The blueprint seems familiar, yet the variation among specimens seems to be limitless. How do we account for dihedral symmetry of order 12 in the large combined with apparently random detail?

Early explanations dismissed the role of chance, often appealing to a higher being. Kepler invoked the “Creator’s design” and wrote, “I do not believe that even in a snowflake this ordered pattern exists at random.” [Kep]. Despite Nakayas’s extensive study both in nature and in the laboratory, he remained perplexed by the symmetry question:

The similarity of the form of the branches is itself a problem that is difficult to explain. There is apparently no reason why a similar twig must grow, in the course of the growth of the crystal, from one main branch when a corresponding twig happens to extend from another main branch. [...] In order to explain this phenomenon we must suppose the existence of some means which informs other branches of the occurrence of a twig on a point of one branch. [Nak]

As we shall see, there is no need whatsoever for a mysterious organizing force such that “one arm knows what the others are doing.”

The most widely accepted scientific explanation of symmetric complexity at present is based on environmental parameters that are locally homogeneous in space but variable in time. Libbrecht summarizes the argument:

Since the atmospheric conditions (e.g., temperature and humidity) are nearly constant across the small crystal, the six budding arms all grow out at roughly the same rate. While it grows, the crystal is blown to and fro inside the clouds, so the temperature it sees changes randomly with time. [...] And because all six arms see the same conditions at the same times, they all grow about the same way. The end result is a complex, branched structure that is also six-fold symmetric. And note also that since snow crystals all follow slightly different paths through the clouds, individual crystals all tend to look different. [Lib4]

Note that this perspective posits nearly deterministic growth in steady weather and attributes much of the disorder in a snowflake’s form to variable conditions at different locations in the clouds. Features such as dramatic branching or faceting are commonly interpreted as indications of sudden change in temperature or supersaturation along the crystal’s trajectory.

Another important part of the puzzle concerns *sample bias*. Simply put, most snow crystals are much less symmetric than people believe. Published collections of photomicrographs do not

reflect the relative statistical abundance of different types. As mentioned earlier, Bentley chose highly symmetric snowflakes almost exclusively. Nakaya’s photos offer a wider morphology, including many irregular types, but still tend to feature outstanding examples of any given variety. Libbrecht tells us (private communication) he rejects thousands of crystals for every one he considers beautiful enough to keep, where beauty tends to involve extraordinary symmetry.

Improved resolution of recent photomicrographs such as those in [LR], [Lib4], and [Lib5] also makes it clear that symmetry of a snow crystal’s six main branches is only approximate. Sectorial plates and other largely faceted snowflakes tend to be more symmetric than dendrites, as is evident in Fig. 2. Crystals with extensive side branching, e.g., fern-like examples, display the most obvious variation between main branches and deviation from mirror symmetry across each main ridge. The appealing order we detect is an overall visual impression enhanced by small-scale dynamic imperfections in which randomness surely plays a role.

Probability is sometimes invoked to explain why “no two snowflakes are alike.” The gist of the entropy argument is as follows: since a crystal typically has dozens of identifiable structural features selected from a sizable assortment of alternative forms, and any combination of those building blocks is possible, the total number of crystal designs is astronomically large. Implicit is the assumption that all these designs arise with comparable frequency, as if the main branches choose them at random. But then why should the six branches all make identical choices? Because the growth isn’t random, but rather the entropy choices are relegated to the local environment.

Even if a dendrite’s motion over the course of several minutes plays a significant role in its ultimate shape, the fine detail of its complex structure is generated on shorter time scales. Our snowflake simulator will make the case for a remarkably stable chaotic deterministic dynamics with a small level of random noise. Also, we will argue that an intrinsic space-time inhomogeneity caused by gradual depletion of vapor around the crystal plays a critical role in morphology.

4 A review of previous models

Researchers have attempted to model snowflakes and their evolution for more than a century. The first geometric pattern inspired by snow crystals, at about the time Bentley’s images gained wide exposure, was the Koch snowflake [Koc] introduced in 1904. Evocative of a sectorial plate with extensive branching, this early fractal curve was based on a simple substitution rule rather than physical principles. More recent constructions in the same vein include Gosper’s flowsnake [Gar], Vicsek’s snowflake [Vic], and the pentflake [Dix], symmetric and self-similar structures with beautiful mathematical properties, but not incorporating growth from a nucleus.

By the mid-20th century, snowflake motifs were firmly established as icons of winter time, despite the fact that very few people ever examine actual crystals firsthand. Many designs on holiday cards, in commercial art, or cut out by schoolchildren with paper and scissors more closely resemble mathematical idealizations than the real object. One frequently encounters symmetries never found in nature and occasionally even a fanciful animation of snowflake growth.

In 1981, Witten and Sander [WS] introduced a groundbreaking prototype for dendritic crystal growth known as *diffusion-limited aggregation* (DLA). Starting with a single “frozen” cell at the origin, particles perform simple random walk on the two-dimensional integer lattice \mathbb{Z}^2 and

freeze when they encounter a previously frozen particle at one of four neighboring sites. In the simplest version from a mathematical standpoint, the diffusing particles are released one by one “from infinity” (or the boundary of a lattice ball of radius R , where R is large). In a more physical variant, the nucleus is surrounded by a “vapor cloud” of particles independently distributed over the lattice with a prescribed density ρ . Simulations of the first version, or the second with small ρ , produce random dendritic growth with delicate fractal properties. See [Hal] for a survey of DLA and its connections to other fractal growth models such as Hele-Shaw flow and the Hastings-Levitov algorithm. A number of researchers have explored variants of DLA with features such as multiple occupancy and anisotropic freezing, but the resulting crystals retain highly stochastic geometry with only statistical symmetry and give little indication of polygonal faceting or the structural instabilities described in Section 2.

In 1984, Packard [Pac] proposed one of the more innovative early variants of DLA, a two-parameter deterministic lattice map with continuous mass at each site of \mathbb{Z}^2 and diffusion implemented by the discrete heat equation rather than random walks. In extreme cases, his model produces main branches with parabolic tips along the axes of the lattice, but the crystals are locally chaotic with little resemblance to real snowflakes. The same paper identified a class of solidifying cellular automata (CA) on the triangular lattice with the property that a site having exactly one frozen neighbor always becomes frozen at the next update. Since real snowflake growth favors the tips, Packard stated that exactly one frozen neighbor should cause freezing, but exactly two should not. Thus, there are 16 *Packard snowflakes*, corresponding to which counts from $\{3, 4, 5, 6\}$ also cause a site to freeze. One of these automata was popularized by Wolfram in a *Scientific American* article [Wol1] as an extremely simple local rule with dynamics ostensibly so complex that computer simulation would be the only effective way to understand system behavior. In fact, we give a rather complete rigorous mathematical analysis of the Packard snowflakes in a companion paper [GG2].

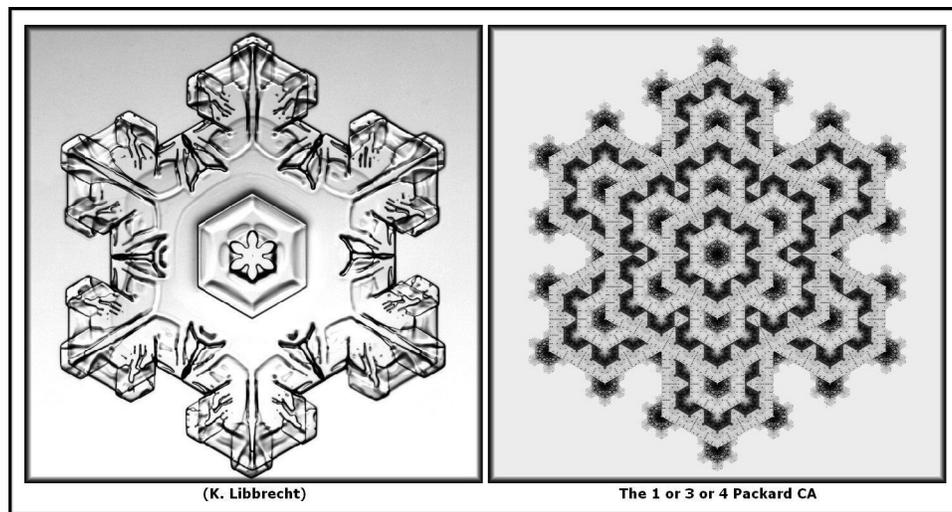


Figure 8. Comparison of a natural crystal (left) and a Packard CA (right)

Another claim of [Wol1], repeated often in the CA literature, asserts that Packard snowflakes capture the essential features of snow crystal growth. Fig. 8 shows a beautiful photo of a sectored

plate on the left and a still frame of the 1 or 3 or 4 Packard CA on the right (started from a single frozen cell and rendered in a periodic gradient palette reflecting the time a cell joins the crystal). The resemblance is certainly intriguing. According to one popular account of complex systems modeling [Lev], “An elementary schoolchild could look at any of the gorgeous pictures of computer screens in Packard’s collection and instantly identify it as a snowflake.” Moreover, the CA suggests how the real crystal of Figure 8 might have nucleated: switching back and forth between a full hexagon and branched shapes with boundary closely related to the Koch curve (see [GG2]), doubling in diameter each time it convexifies. Two decades later, Wolfram continues to claim that nature emulates simple cellular algorithms:

One expects that during the growth of a particular snowflake there should be alternation between treelike and faceted shapes, as new branches grow but then collide with each other. And if one looks at real snowflakes, there is every indication that this is exactly what happens. The simple cellular automaton seems remarkably successful at reproducing all sorts of obvious features of snowflake growth. [Wol2]

As we have seen in Section 2, however, actual nucleating snow crystals do not branch until they have grown to micron scale, so they certainly do not behave like the CA model at the outset. Also, from each hexagonal state Packard snowflakes branch at the tips, whereas the real first instability occurs in the middle of facets. What is more, diffusion limited growth typically causes main branches to progress more quickly than side branches, so even with aftergrowth there is very little evidence for repeated convexification on a dyadic scale. Like the Koch snowflake, Packard’s automata are beautiful mathematical structures in their own right, but all scientific evidence suggests that real sectored plates evolve in quite a different way. Exploring the disparity was our initial motivation for the research presented here.

Three papers from the late 1980s merit a mention in this review. Nittmann and Stanley [NS] combined ingredients of DLA with a surface tension parameter to obtain an admittedly ad hoc model with tunable parameters that captures some qualitative features of dendritic snowflakes. Family, Platt, and Vicsek [FPV] stressed the role of anisotropic surface tension in producing regular dendritic crystals. Their similarly artificial approach was featured in the best-selling book on chaos theory [Gle]. The most successful stochastic model to date is due to Xiao, Alexander, and Rosenberger [XAR]. They constructed a version of DLA with anisotropic attachment and surface diffusion, encoded by means of random walks and next-nearest neighbor freezing probabilities on the triangular lattice. To our knowledge, [XAR] is the first paper with compelling evidence of realistic instabilities and faceting. As with any Monte Carlo approach, however, deterministic behavior is veiled by fluctuations unless the system size and run time are enormous. Even with noise reduction techniques, computation of detailed crystal structure is not feasible.

Advances in the 1990s incorporated the theory of PDE and boundary value problems more directly. Yokoyama and Kuroda ([YK], [Yok]) analyzed a continuous state model featuring a quasi-liquid layer and computable by finite element methods. They did not obtain numerical results beyond the first tip instability, but they demonstrated some rudimentary faceting. Their scheme involves rather complicated physically motivated anisotropic attachment kinetics and stresses the role of macrosteps:

In particular, the surface kinetic process plays a very important role in the formation of such surface polyhedral features of snow crystals as hexagonal prisms and facets seen at the tips of dendrites. We emphasize that such surfaces are molecularly smooth and that they cannot grow without the lateral motion of steps. [YK]

Kobayashi [Kob] proposed a phase-field approach to dendritic crystal growth, whereby the state of the system in a boundary layer of small but positive thickness is represented as a mixture of vapor and ice. Many subsequent papers have attempted to model two- and three-dimensional dendrites using this paradigm. See, e.g., [KL] for an application to image processing. A more classical approach was taken by Schmidt [Sch], who studied an anisotropic two-dimensional Stefan problem and presented corresponding numerical methods. Whereas the transition from circular initial condition to a hexagon and then the onset of the first instability look promising, his deterministic simulations break symmetry before any side branching occurs, presumably due to either roundoff error or an intrinsic singularity in the system of PDE.

Most recently, Reiter [Rei] introduced a coupled lattice that does a remarkably good job of emulating stellar dendrites. A simpler variant of Packard's original continuous-state scheme, Reiter's model has two parameters: vapor density and an ad hoc constant shift in the mass field. Since the attachment rate is constant, his crystals lack the correct geometry for instabilities, macrosteps, and faceting, and display quite sensitive dependence on system parameters. The only sectorized plate candidates found within his phase portrait evolve more like a Packard CA than a real snow crystal. Nevertheless, the mesoscopic framework of Reiter's model strikes the best balance to date between speed of computation, on the one hand, and physical realism, on the other. The algorithm we present in the next section may be viewed as a refinement of Reiter's model that incorporates physically motivated features from previous work such as [XAR] and [YK].

5 Our algorithm for two-dimensional snow crystal growth

Based on empirical evidence described in Section 2, the model we will now present for ideal planar snowflake dynamics makes the following simplifying assumptions:

- A1. *Nucleation and the early stages of growth produce a thin hexagonal prism with basal facets a few microns in diameter and prism facets much shorter in height;*
- A2. *Subsequent growth is effectively two-dimensional, by attachment of vapor at the edge of the basal surfaces and with no other significant changes on the basal surfaces;*
- A3. *Crystal melting and sublimation effects are negligible;*
- A4. *The quasi-liquid layer of mixed ice and vapor near the edge of the basal surfaces stays at most a few microns in width throughout the crystal's formation (no roughening);*
- A5. *Crystal motion is either unimportant or can be interpreted as change in supersaturation.*

Whereas our assumptions surely ignore many details of the real story, Figs. 2 and 4-6 show that essential features are still captured surprisingly well.

In our mesoscopic two-dimensional model, each site represents a cell approximately $5\mu\text{m}$ across. Thus, the thin prism of (A1) corresponds to a single initially frozen cell at the origin surrounded by diffusing vapor. To reflect the underlying crystalline structure of I_h , we tessellate the plane as a honeycomb and choose the centers of the hexagonal cells for sites of our triangular lattice \mathbb{T} so that each site has six nearest neighbors. Keep in mind, though, that a single site of \mathbb{T} represents more than $(10^3)^3 = 10^9$ (nanoscale to micron scale in each of three dimensions) water molecules. At each discrete time $t = 0, 1, 2, \dots$ and with each site $x \in \mathbb{T}$, we associate a Boolean variable and three varieties of mass. Formally, the state of the system at time t at site x is $\xi_t(x) = (a_t(x), b_t(x), c_t(x), d_t(x))$ where the attachment flag

$$a_t(x) = \begin{cases} 1 & \text{if } x \text{ belongs to the crystal at time } t, \\ 0 & \text{otherwise;} \end{cases}$$

and

$$\begin{aligned} b_t(x) &= \text{the boundary mass at } x \text{ at time } t && (\textit{quasi-liquid}), \\ c_t(x) &= \text{the crystal mass at } x \text{ at time } t && (\textit{ice}), \\ d_t(x) &= \text{the diffusive mass at } x \text{ at time } t && (\textit{vapor}). \end{aligned}$$

Our dynamics will ensure that the snowflake consists entirely of crystal mass, all three mass types can coexist on the boundary of the snowflake, and only diffusive mass occurs elsewhere. Initially, there is crystal mass 1 at the origin and diffusive mass ρ everywhere else.

Before describing our simulator's details, let us briefly discuss a much simpler, mathematically tractable aggregation model as motivation for the mesoscopic setting. *Internal diffusion-limited aggregation* (IDLA) is a relative of the Witten-Sander model in which particles successively dropped at the origin at a constant (random) rate perform independent simple random walks until they reach the boundary of the crystal, at which time they freeze. Whereas DLA reinforces crystal tips, IDLA reduces them by filling indentations. [LBG] showed that IDLA has a circular asymptotic shape in the plane (and spherical shape in any dimension). Thinking of the random walks as diffusion on a microscopic scale, and letting the process evolve until the circular crystal has assumed a macroscopic size many orders of magnitude larger, one expects the relative proportions of walks still diffusing on mesoscopic chunks to assume a deterministic isotropic density profile that evolves over time. Indeed, [GQ] proved that, suitably rescaled, the space-time occupation statistics of planar IDLA and various generalizations converge to solutions of corresponding one-phase Stefan problems. Simply put, the stochastic particle system obeys a dynamical law of averages, or *hydrodynamic limit*, described by a classic PDE. As a corollary, [GQ] computed the speed of two-dimensional IDLA as the solution of a transcendental equation and found the limiting density profile explicitly.

The analogy with snowflake models is illustrated well by papers described in the previous section. [XAR] introduced a microscopic stochastic particle system akin to IDLA. Although the physical ingredients were carefully chosen, random fluctuations preclude a computationally feasible, realistic result, even with extensive averaging of Monte Carlo runs. On the macroscopic scale, the Stefan problems of [Sch] and [GQ] are analogous. While the latter is exactly solvable, anisotropic curvature and additional boundary conditions in the former present formidable challenges, both mathematically and for simulation. Moreover, it is not at all clear how to incorporate the structure of a quasi-liquid layer into a well-posed boundary value problem. For IDLA, [GQ] made use of deterministic mesoscopic dynamics to prove convergence of rescaled

stochastic particle systems to their hydrodynamic limits. Attachment rates of the random microscopic system were replaced by occupation thresholds in that mesoscopic scheme. The same mechanism provides a key ingredient in the simulator we will now specify.

The Snowflake Simulator

Recall that $\xi_t(x) = (a_t(x), b_t(x), c_t(x), d_t(x))$, with $a_0(\mathbf{0}) = c_0(\mathbf{0}) = 1$, $b_0(\mathbf{0}) = d_0(\mathbf{0}) = 0$; and for all $x \neq \mathbf{0}$, $a_0(x) = b_0(x) = c_0(x) = 0$, and $d_0(x) = \rho$. This initialization corresponds to a mesoscopic prism at the origin surrounded by homogeneous vapor with density ρ . A typical simulated crystal, with ρ between about .3 and .9, reaches a final diameter of 400–600 cells over 10,000–100,000 updates of the following nearest neighbor rule. In the deterministic case, Steps (i)–(iv) below are performed in order every discrete time unit. In the randomized version Step (v) is added. The physical processes represented by the steps naturally occur in parallel, but we order them for the sake of computation. Over tens of thousands of updates, the ordering of steps is insignificant, due to the diffusive hydrodynamic scaling.

In the sequel, denote $\mathcal{N}_x = \{x\} \cup \{y : y \text{ is a nearest neighbor of } x \text{ in } \mathbb{T}\}$, and set

$$\begin{aligned} A_t &= \{x : a_t(x) = 1\} = \text{the snowflake at time } t; \\ \partial A_t &= \{x \notin A_t : a_t(y) = 1 \text{ for some } y \in \mathcal{N}_x\} = \text{the boundary of the snowflake at time } t; \\ A_t^c &= \{x : a_t(x) = 0\} = \text{the sites not in } A_t; \\ \bar{A}_t^c &= (A_t \cup \partial A_t)^c = \text{the sites not in } A_t \text{ or } \partial A_t. \end{aligned}$$

Also, we use $^\circ$ (degree) and $'$ (prime) notation to denote amounts of mass before and after a step or substep is completed. This is necessary since some mass allocations may change more than once during a single cycle of the steps. At the end of each cycle the time t advances to $t + 1$.

Steps of the update rule:

i. Diffusion

Diffusive mass evolves on A_t^c by discrete diffusion with uniform weight $\frac{1}{7}$ on the center site and each of its neighbors. Reflecting boundary conditions are used at the edge of the crystal. In other words, for $x \in \bar{A}_t^c$,

$$(1) \quad d_t'(x) = \frac{1}{7} \sum_{y \in \mathcal{N}_x} d_t^\circ(y),$$

and for $x \in \partial A_t$ any term in the sum corresponding to $y \in A_t$ is replaced by $d_t^\circ(x)$. This is a standard implementation of the heat equation for vapor transport on the complement of the snowflake. One can tune the diffusion rate, which varies somewhat in nature, by modifying the weight of the center site in (1), but similar effects are achieved by changing the vapor density ρ so we use only the uniform distribution for simplicity.

ii. Freezing

Proportion κ of the diffusive mass at each boundary site crystallizes. The remainder (proportion $1 - \kappa$) becomes boundary mass. That is, for $x \in \partial A_t$,

$$(2) \quad b_t'(x) = b_t^\circ(x) + (1 - \kappa)d_t^\circ(x), \quad c_t'(x) = c_t^\circ(x) + \kappa d_t^\circ(x), \quad d_t'(x) = 0.$$

Parameter κ is one of three mass transfer rates that emulate the dynamics of the quasi-liquid layer on the boundary of the crystal. The other two, μ and γ , appear in step *iv* below. In the freezing step, we suppose that during each discrete unit of time, representing a few hundredths of a second, a small portion κ of the vapor at the crystal boundary freezes directly while the rest becomes quasi-liquid.

iii. Attachment

This key step in the algorithm decides when a boundary site joins the snowflake. We have seen that anisotropic attachment is crucial to realistic growth. In our context, this means that if $x \in \partial A_t$, the condition for setting $a_{t+1}(x) = 1$ depends on the values of ξ_t° on \mathcal{N}_x . Let $n_t^\circ(x) = \#\{y \in \mathcal{N}_x : a_t^\circ(y) = 1\}$ be the number of attached neighbors of x at time t . Our attachment rule divides into three cases depending on whether $n_t^\circ(x) = 1$ or 2 , $n_t^\circ(x) = 3$, or $n_t^\circ(x) \geq 4$. Motivated by the IDLA parallel described above, attachment will require variable thresholds of boundary mass. The higher $n_t^\circ(x)$, the easier it is to attach, so the lower the threshold.

A boundary site with 1 or 2 attached neighbors needs boundary mass at least β to join the crystal:

$$(3a) \quad \text{If } x \in \partial A_t^\circ, n_t^\circ(x) = 1 \text{ or } 2, \text{ and } b_t^\circ(x) \geq \beta, \text{ then } a_t'(x) = 1.$$

This is the case when the local mesoscopic geometry near x corresponds to a tip or flat spot of the crystal. (Distinguishing the two cases turns out to be of minor significance.) In our simulations, β is typically between about 1.05 and 3. We assume $\beta > 1$ since 1 is the basic threshold of the case to follow next.

A boundary site with 3 attached neighbors joins the crystal if either it has boundary mass ≥ 1 , or it has diffusive mass $< \theta$ in its neighborhood and it has boundary mass $\geq \alpha$:

$$(3b) \quad \text{If } x \in \partial A_t^\circ, n_t^\circ(x) \geq 3, \text{ and} \\ \text{either } b_t^\circ(x) \geq 1 \text{ or } \left(\sum_{y \in \mathcal{N}_x} d_t^\circ(y) < \theta \text{ and } b_t^\circ(x) \geq \alpha \right), \text{ then } a_t'(x) = 1.$$

In this case, when there is a local mesoscopic concavity of the crystal boundary near x , the attachment threshold of boundary mass is normalized to 1. There is an additional way to fill concavities, however: parameters α and θ implement the *knife-edge instability*, as we explain later in this section. We also note that α could be a complicated function of the diffusive mass, rather than having a single jump at θ . We have experimented with a variety of such rules, and (3b) consistently gave the most realistic results.

Finally, boundary sites with 4 or more attached neighbors join the crystal automatically:

$$(3c) \quad \text{If } x \in \partial A_t^\circ, n_t^\circ(x) \geq 4, \text{ then } a_t'(x) = 1.$$

Here ice is sufficient near x that the attachment threshold is 0, precluding the formation of occasional single-cell holes in the snowflake (mainly for aesthetic reasons).

Once a site is attached, its boundary mass becomes crystal mass:

$$(3d) \quad \text{If } x \in \partial A_t^\circ, \text{ and } a_t'(x) = 1, \text{ then } c_t'(x) = b_t^\circ(x) + c_t^\circ(x), \text{ and } b_t^\circ(x) = 0.$$

Attachment is permanent, and there are no further dynamics at attached sites.

iv. Melting

Proportion μ of the boundary mass and proportion γ of the crystal mass at each boundary site become diffusive mass. Thus, for $x \in \partial A_t$,

$$(4) \quad b'_t(x) = (1 - \mu)b_t^\circ(x), \quad c'_t(x) = (1 - \gamma)c_t^\circ(x), \quad d'_t(x) = d_t^\circ(x) + \mu b_t^\circ(x) + \gamma c_t^\circ(x).$$

Melting represents mass flow at the boundary from ice and quasi-liquid back to vapor, reverse effects from the freezing of step *iii*. Typically μ is small and γ extremely small.

v. Noise

The diffusive mass at each site undergoes an independent random perturbation of proportion σ :

$$(5) \quad d'_t(x) = (1 \pm \sigma)d_t^\circ(x) \quad \text{with probability } \frac{1}{2} \text{ each.}$$

The microscopic statistical physics of snow crystal growth only produces deterministic coarse-grained dynamics in the hydrodynamic limit. As mentioned previously, our mesoscopic cells contain on the order of 10^{10} molecules, so growth is almost deterministic but still subject to chance. Our simulator with only steps *i-iv* provides an ideal model, complex but nonrandom. Adding step *v* is perhaps the simplest way to get a feeling for the role of residual stochastic effects. Parameter σ signifies a tiny noise level in the diffusive flow of vapor due to environmental fluctuations.

Having completed the formal description of the algorithm, let us make a few remarks about computer implementation. Ideally, our model evolves on the infinite triangular lattice \mathbb{T} , but for effective visualization and efficient simulation, a finite rectangular grid is best. Fortunately, there is a graph isomorphism between \mathbb{T} and the homogeneous graph consisting of sites of the two-dimensional integers \mathbb{Z}^2 , with edges connecting each site to nearest neighbors in 6 of the 8 horizontal, vertical, and diagonal directions, say $\{N, S, E, W, NE, SW\}$. Call this graph \mathbb{T}' . All simulated crystals in this paper were generated on \mathbb{T}' and mapped to \mathbb{T} using a rotation of 45° , then a vertical rescaling by $1/\sqrt{3}$. For a bounded region of \mathbb{T}' it is most convenient to choose a square grid of diameter L centered at the origin, with periodic boundary conditions. Care must be taken to ensure that the snowflake is not influenced by boundary effects. For rapidly growing dendrites, it is safe as a rule of thumb to take L 50% larger than the final diameter of the crystal. For especially slow-growing stellar plates, L should be at least twice as large as the final diameter.

Without noise (no Step *v* above), note that total mass is conserved. Not only is this property appealing from a physical perspective; it also helps in debugging code and checking numerical stability. For computational efficiency, one can also exploit exact symmetry in the deterministic model by taking the finite lattice to be a hexagon of diameter L with patched wrap boundary conditions, in which case it suffices to compute the dynamics on $\frac{1}{12}$ of the space.

The remainder of this section is devoted to a preliminary discussion of the digital crystals grown by our simulator, which we like to call *snowflakes*. Figures 2 and 4–6 show some key features. The gray colors inside (resp., outside) the crystal represent the values of c_t (resp.,

d_t), with darker shades corresponding to higher values. A better understanding can be achieved by exploring the roles of the eight parameters: ρ , β , α , θ , κ , μ , γ , and σ . The first seven are addressed below; we defer a discussion of σ and the randomized simulator until Section 6. We also note that the values of all parameters used in the figures are given in the Appendix. But first, a few general remarks about our model are in order.

To begin with, we briefly address the issue of scale. As mathematical objects, our snowflakes can grow to arbitrary size. What is the typical morphology as $t \rightarrow \infty$? This is a very challenging question, but it seems that a crystal that does exhibit the first instability and side branching should resemble a fern once it grows large enough. In particular, the *primitive* case in which all parameters other than ρ and β are 0 seems to satisfy these conditions. We say that the growing crystal A_t has a (possibly nonconvex) *asymptotic shape* L if $t^{-1}A_t$ converges to L as $t \rightarrow \infty$ in the Hausdorff metric. Moreover, we say that A_t has *asymptotic density* ρ_∞ if ϵ^{-2} times the counting measure on the set $\epsilon \cdot (\cup_{t \geq 0} A_t)$ converges weakly to the constant ρ_∞ times the Lebesgue measure on \mathbb{R}^2 as $\epsilon \rightarrow 0$. The reader is referred to [GG1, GG2] for a detailed discussion of these concepts and their relevance to Packard snowflakes. In this paper, we emphasize modeling issues, so our attention is restricted to realistic time scales, but it seems that many of our snowflakes have both an asymptotic shape and an asymptotic density. Simple stars and simple plates should have these properties by definition, and we also suspect they should hold in the primitive case.

Of course even planar snowflakes are actually three-dimensional, and photomicrography reveals substantial variation in the thickness of various features. For instance, the ridges of main and side branches are much thicker than plates that form between them. In our model, ridges accumulate more crystal mass because of anisotropy: a higher threshold of boundary mass is needed at the crystal tip than in a concavity. From one perspective, this is a formal device capturing the hydrodynamic effect of variable random attachment rates, lower rates corresponding to higher thresholds. However one can imagine vapor particles concentrating in regions where it is harder to freeze and thereby “piling up” when they do attach. The remarkable message of Figs. 4–6, one that we do not entirely understand so far, is how well crystal mass in our simulator emulates the thickness of actual snowflakes.

Despite its relatively simple approach to a decidedly complex physical phenomenon, our algorithm is certainly overdetermined. Six of the eight parameters seek to capture attachment kinetics in terms of freezing and melting rates across the quasi-liquid layer at the crystal boundary. Existence of such rates is quite plausible, and it might even be possible to estimate them experimentally, but they are surely not independent variables. Rather, they should all be governed in a very complicated way by temperature, supersaturation, and other more basic physical parameters. Starting on a mesoscopic scale, we can offer little insight into outstanding mysteries such as the form of the Nakaya diagram. Moreover, some craft is needed tuning our parameter set in order to achieve plausible dynamics. For example, the morphology of the rather uninteresting crystal in the bottom right frame of Fig. 11 is exceedingly common among random choices of parameters. Our point of view is that a parameter set that produces a final set close to a physical crystal can be then used to provide insight into aspects of its possible evolution such as the order of appearance of its features or expansion velocities at different stages.

To make sense of an eight-dimensional phase portrait is a daunting task. Our strategy here is to choose a sampling of realistic examples somewhat arbitrarily and to vary a single parameter we wish to study with all others held fixed. For simplicity, the simulations in this section are

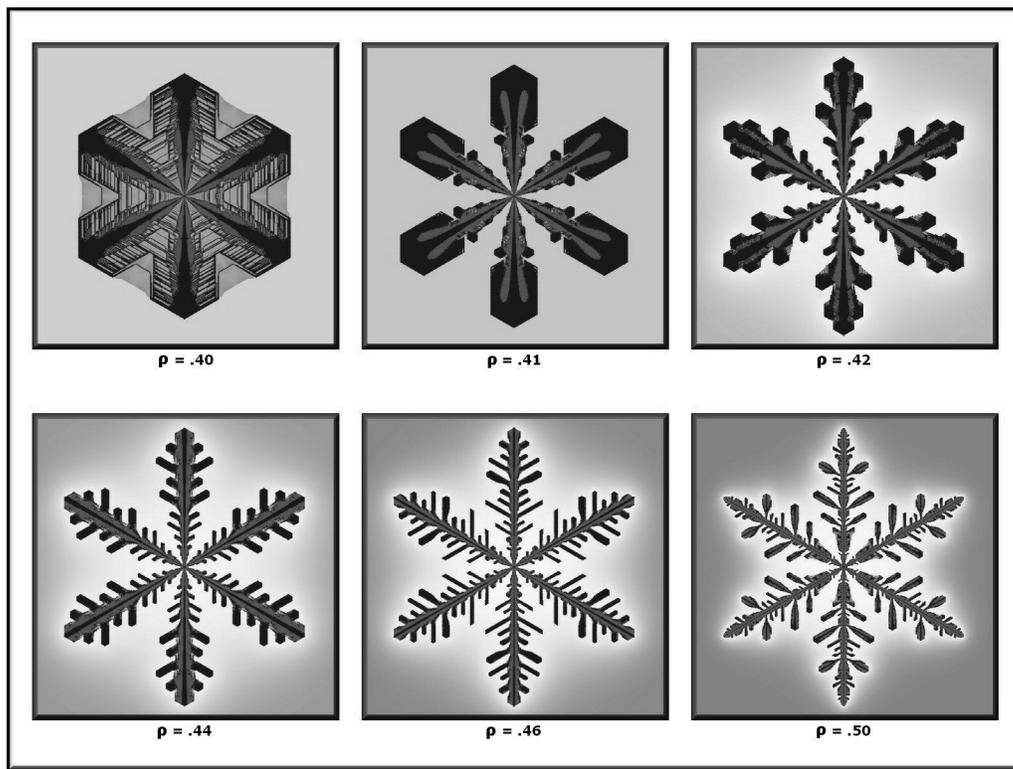


Figure 9. Dependence on the vapor density parameter ρ

deterministic ($\sigma = 0$, i.e., no step v). Let us begin by illustrating the role of vapor density ρ . Figure 9 shows a sequence of snowflakes grown to comparable size with vapor density varying from .4 to .5 and all other parameters identical. Naturally, the crystals grow more rapidly as ρ increases; observe the transition in morphology from plate to sectored plate to dendrite. The plate case exhibits delayed, thin filling of concavities due to the knife-edge instability. Vapor density is the easiest parameter to understand since it plays a role closely related to supersaturation.

Recall that parameter β controls the anisotropy of attachment — how much harder it is to freeze to a mesoscopically protruding or flat chunk of boundary than in a valley. The disparity is governed by how much β exceeds 1. The series of snowflakes in Fig. 10, with β increasing from 1.9 to 2.8, shows how anisotropy promotes faceting and a transition from fern to dendrite to sectored plate. Note three lessons of the sequence. First, increasing β delays the onset of the first instability, increasing the size of the central hexagonal plate before main branches develop. Second, the hexagonal form of branch tips grows in scale and definition, whereas the propensity for side branching is diminished. Finally, note the curious nonmonotonicity in knife-edge aftergrowth, increasingly present up to $\beta = 2.7$ and then absent for $\beta = 2.8$.

Our version of the knife-edge instability in terms of α and θ in substep (3b) of the algorithm is the most speculative aspect of the simulator. Recall there is evidence for a sudden change in the growth rate of a plate once its edge is sufficiently thin ([Lib2], [Lib5]). We achieve a similar effect indirectly by requiring a boundary mass threshold of only $\theta < 1$ for attachment

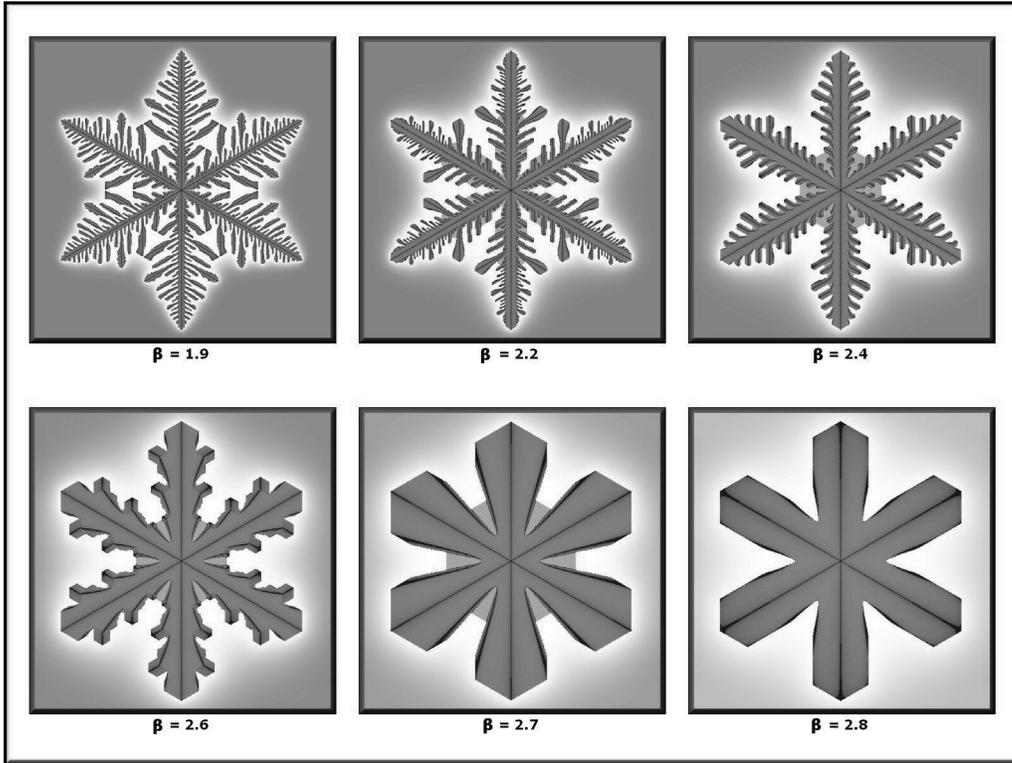


Figure 10. Dependence on the anisotropy parameter β

once the diffusive mass in a mesoscopic boundary patch dips below α . Intuitively, the edge should become thinner as it has less access to vapor. But then, since a very thin plate requires little vapor to advance, and due to this rather mysterious instability, more rapid aftergrowth occurs. Tuning α and θ controls both the strength and the geometry of our implementation. The closer θ is to 0, the more delay there is in attachment off the ridges. The closer α is to 0, the more the instability manifests itself in concavities between ridges rather than along sides of ridges. For example, in Fig. 2 the simulated stellar plate has a much lower value of θ than the simulated stellar dendrite. The intrinsic nonmonotonicity of the knife-edge instability and our rather crude attempt to capture it can lead to excessive burstiness of aftergrowth unless the parameters are adjusted carefully. But we are convinced that some ingredient such as this offers the only explanation for thin, faceted growth observed in the concavities of some planar snow crystals after most of the vapor has been depleted there.

The next two parameters we examine control the mixture of crystal, boundary, and diffusive mass present within our simulated boundary layer. According to the freezing step (2), vapor diffusing to the edge of the crystal enters the quasi-liquid state except for a small proportion κ that is deposited directly as ice and simply waits to attach to the snowflake. Such ice does not enter into the attachment step (3). Since vapor is most concentrated near tips of the main branches, increasing κ somewhat tends to deprive these tips of more boundary mass than the tips of side branches. Consequently, simulations with very small κ have less side branching than those with a somewhat larger value, as shown in Fig. 11. Increasing κ still further inhibits even

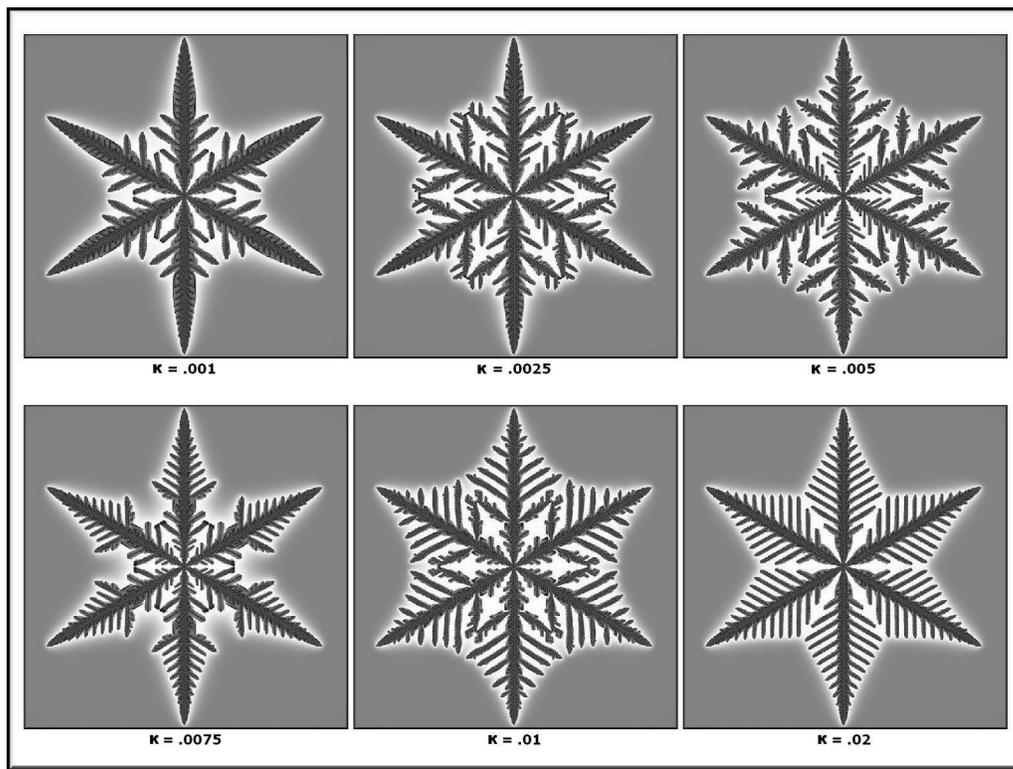


Figure 11. Dependence on the crystallization parameter κ

side branching, a nonmonotonicity we illustrate in the next section that indicates the subtlety of boundary layer dynamics.

The melting step (4), which releases a small proportion μ of boundary mass back to the vapor field, has an effect opposite to that of freezing. Thus, increasing μ promotes faceting. Fig. 12 starts from a delicate stellar dendrite generated by a parameter set with $\mu = .04$. Note the apparent parabolic tips for this snowflake and the second in the sequence, with $\mu = .05$. As melting increases further, though, hexagonal tips become evident first on the main branches, then on the larger and less frequent side branches. By the final frame, with $\mu = .09$, the morphology has clearly changed to a stellar plate. What is the type of the second-to-last crystal? This example nicely illustrates how snow crystal growth generates a continuum of possible forms, so that the boundary cases of any classification scheme are inevitably arbitrary.

Melting also releases a small proportion γ of crystal mass in the boundary layer back to the vapor field. The role of γ is very similar to that of μ but will typically be much smaller. We include the possibility of nonattached ice reverting to vapor mainly for completeness and thus omit an illustration of its effect.

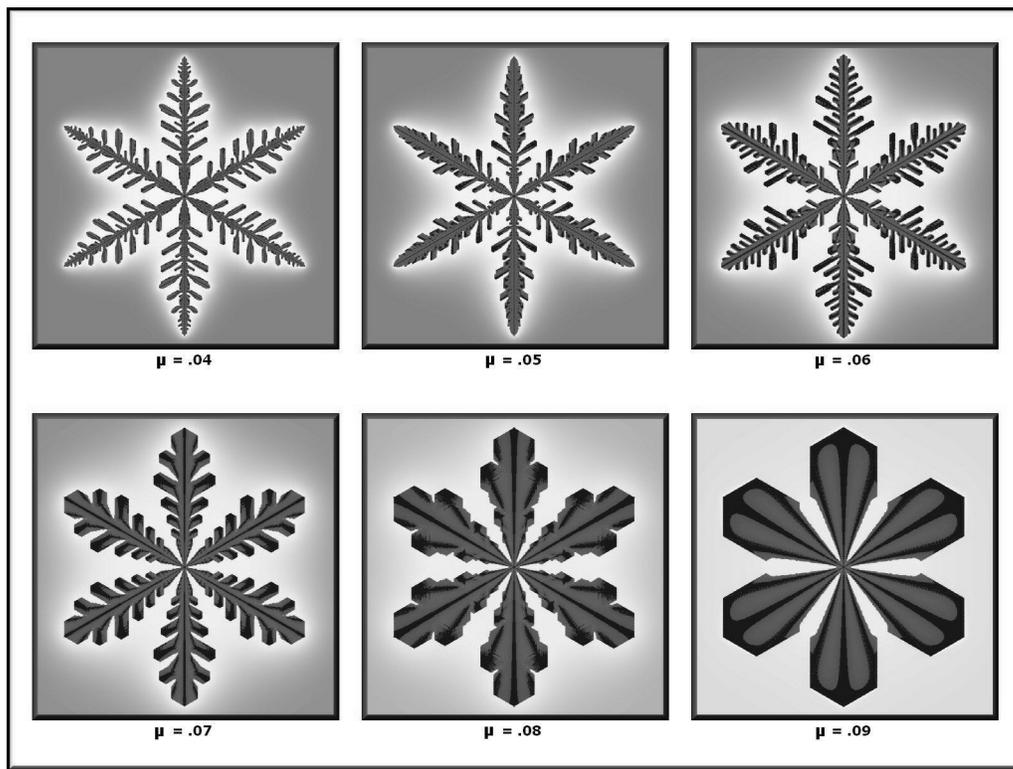


Figure 12. Dependence on the melting parameter μ

6 Additional case studies

Our investigation of individual parameters in the last section identified several different ways the phase portrait of our simulator interpolates between simple plates at one extreme and fern-like crystals at the other, with sectored plates and stellar dendrites in between. All other types of the Magono-Lee classification are also generated, but some rely on rather delicate interplay between the competing forces of diffusion-limited growth, anisotropic attachment, and boundary layer dynamics. We now present a collection of case studies to indicate the full range of snowflakes. The first few provide additional types that, somewhat surprisingly, arise from a homogeneous vapor cloud with suitable density ρ . Then we offer some speculation about the hieroglyphs observed on many simple plates. Third, we exhibit several crystals grown in variable environments, with either a gradual or a sudden change in the surrounding vapor density. Finally, the role of randomness is addressed, as captured by our simulator's noise parameter σ .

Our first case studies illustrate a fundamental feature of snowflake formation often overlooked by researchers: growth from micron scale to final size typically takes place far from equilibrium. Local loss of vapor due to solidification, most prevalent in the concavities between tips of main and side branches, often continues to develop over the entire evolution. Our graphics of simulated crystals render diffusive mass in a gradient from dark gray = density ρ to white = density 0, so the depletion is represented by a halo-like region surrounding each crystal. Even in an environment constant over space and time, diffusion-limited growth causes an intrinsic temporal

inhomogeneity near the boundary. There is no *need* for external forces to induce branching or other apparent variations in morphology.

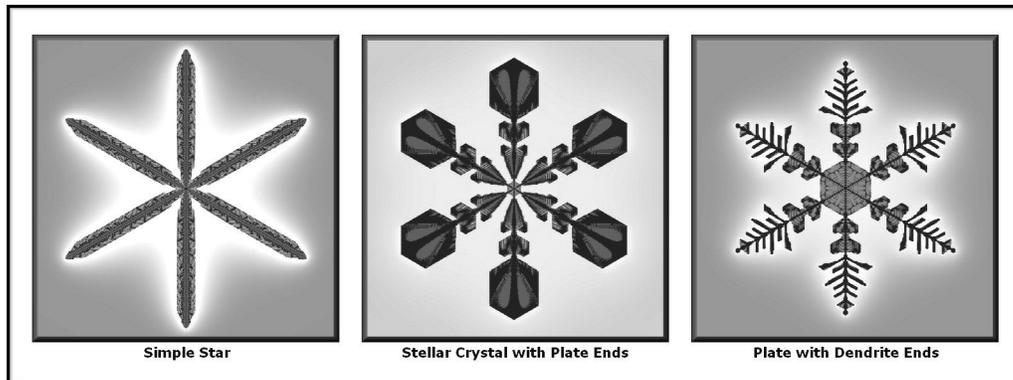


Figure 13. Deterministic crystals in a homogeneous environment

Three unusual crystals grown with our basic simulator are shown in Fig. 13. Let us analyze them individually:

A simple star

Rather high vapor density ($\rho = .65$) compared to the anisotropy index ($\beta = 1.75$) promotes early onset of the first instability and rapid advance of the six main tips. Very strong direct freezing ($\kappa = .15$) inhibits further branching. Suitable levels of melting ($\mu = .015$, $\gamma = .00001$) and aftergrowth ($\alpha = .026$, $\theta = .2$) promote faceted sides along the main branches without further development of a central plate. The result is an elegant geometric shape with elaborate internal markings after 15,000 updates.

A stellar crystal with plate ends

Here the vapor density is much lower ($\rho = .36$), but so is the anisotropy ($\beta = 1.09$), and direct freezing is negligible ($\kappa = .0001$). The result is much slower growth that still manages to exhibit an early first instability without appreciable side branching. A very high melting rate ($\mu = .14$, $\gamma = .00001$) coupled with moderate knife-edge aftergrowth ($\alpha = .01$, $\theta = .0745$) promotes a complex faceting process by repeatedly repairing tip instabilities. Vapor depletion reinforces the tendency to facet over time, resulting in very large, hexagonal plate ends with ridges and other complex hieroglyphs after more than 100,000 updates.

A plate with dendrite ends

Now suppose we slightly raise the vapor density ($\rho = .38$), lower the anisotropy ($\beta = 1.06$), and maintain a small rate of direct freezing ($\kappa = .001$). These changes retain an early first instability but tilt the balance toward dendritic growth. With virtually the same melting rates compared as the last example. ($\mu = .14$, $\gamma = .00006$) but quite different aftergrowth parameters ($\alpha = .35$, $\theta = .112$), we obtain the remarkable snowflake at the right of Fig. 13 after 20,000 updates. It is important to note that the central plate continues to spread long after the first instability, as the knife-edge effect repeatedly fills regions between the six main tips. Our next case study examines this process in more detail.

The explanation for many hieroglyphs on the central plates of snowflake photomicrographs remains elusive. As noted in Section 2, some apparent surface details are actually see-through views of three-dimensional features such as the vestigial center of a dual layer or encased bubbles. Ubiquitous main ridges and their offshoots are unquestionably formed by tip dynamics and branching instabilities. But concentric, irregularly spaced lateral bands known as *ribs* are also common, as in the left image of Fig. 14. Conceivably, these marks might be formed by diffusion on the plate surface, but there is little physical evidence for such ([Lib3]), and one would not expect the ribs to be so well defined and linear.

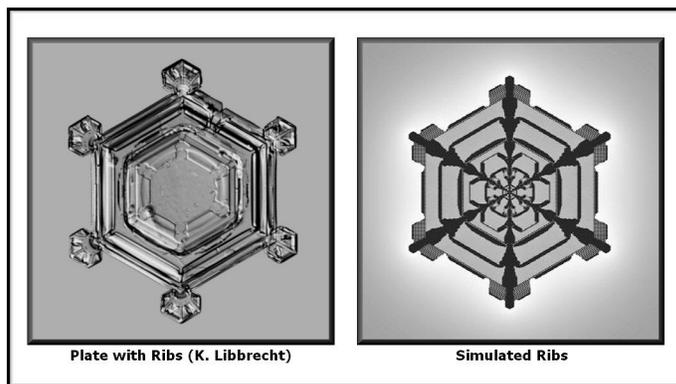


Figure 14. Natural and simulated ribs

Our simulator suggests a possible cause of ribs based on knife-edge effects. We suspect that, as proposed long ago in [She], diffusion-limited growth first forms the network of ridges (a kind of skeleton for the crystal), and then gaps fill after a slight delay. In many cases, our model continues growing a central plate long after the first instability, adding hexagonal bands that almost keep up with the six main tips. Nonmonotonicity causes the bands to stop growing from time to time, with regular spacing in some instances and irregular spacing in others. Each time a band stops, there is a local accumulation of mass along its edge before growth begins again. The shape of these bands is nearly hexagonal, reflecting the geometry of levels sets in the diffusion field. The right image of Fig. 14 shows a simulation that produced rather realistic ribs in this way. Often the crystal tips eventually establish enough of a lead that aftergrowth of the central plate stops altogether. At the very least, such examples suggest that after the first instability, natural snow crystals may undergo a more complex succession of branching and convexification than is suggested by Fig. 4 and that rib hieroglyphs may be formed in the process.

For the remainder of this section, we turn to the central issue of environmental variability. Hydrodynamic changes due to crystal motion, or external temperature and supersaturation fluctuations, should be well approximated by a time-dependent schedule for the density and attachment parameters of our deterministic simulator. A very small residual noise from the microscopic physics should also be present on the mesoscopic scale. Let us turn to these perturbations of our basic model.

Figure 15 shows six deterministic snowflakes grown in vapor fields with changing density ρ over time. In the first three examples, the external equilibrium changes gradually, as one would expect in most physical settings on the relevant space and time scales, cf. [Lib5].

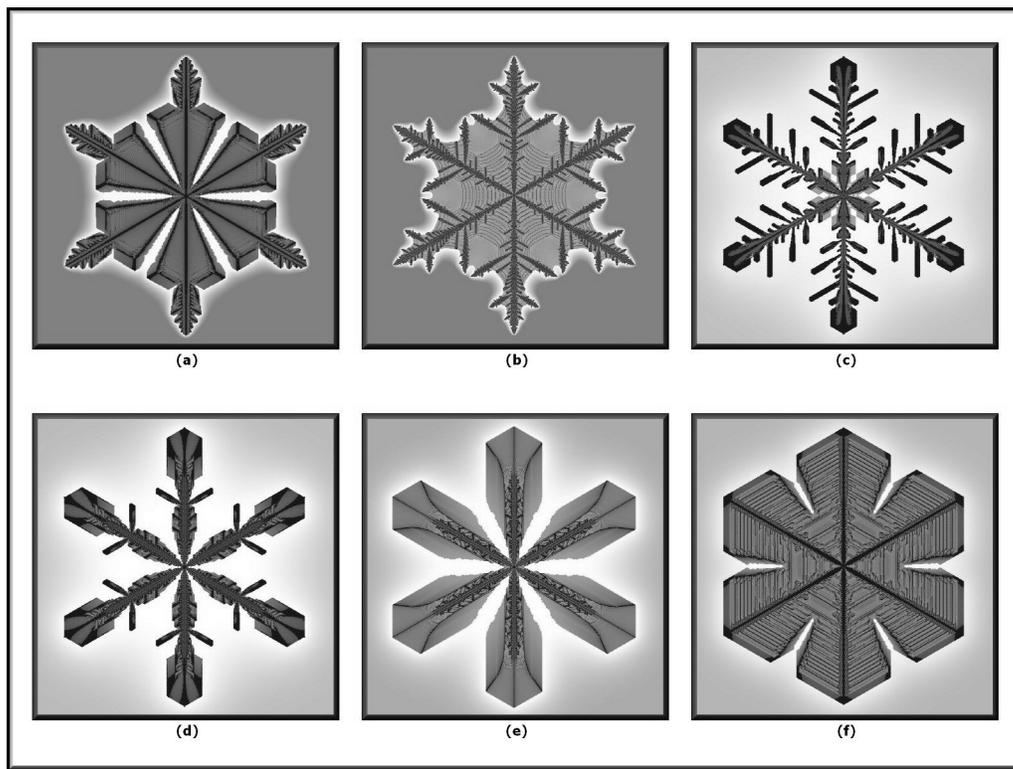


Figure 15. Deterministic crystals in varying environments

Cases (a) and (b) are meant to explore the possible role of velocity once a snowflake gains sufficient mass to start falling more quickly. As mentioned in Section 2, experiments have shown that this effect is qualitatively similar to a continuous boost in supersaturation. Thus, we modify our algorithm by increasing the diffusive mass everywhere outside the crystal by a small proportion at each time step, resulting in a dramatic change over more than 10,000 updates. The first example started from a phase point representing a sectored plate, the second from a simple plate. Not surprisingly, the tips in both cases develop side branches and become more dendritic as added exterior diffusive mass counteracts depletion at the crystal boundary. Note the visual impression of rather abrupt changes in morphology despite gradual changes in the environment. Incidentally, the sectored plate of Fig. 2 in the introduction was also generated with a slowly increasing vapor density. Case (c) differs in that ρ decreases gradually, as when a crystal drifts lower in the clouds and branching leads to faceting.

For comparison, cases (d)–(f) in Fig. 15 feature a sudden decrease in vapor density, as might occur when a snowflake leaves the cloud cover. The stellar dendrite of (d) develops broader hexagonal facets in a less saturated environment. As previously mentioned, environmental changes are expected to affect the attachment kinetics, which is illustrated by (e). Here, in addition to ρ , the crystallization parameter κ changes suddenly, resulting in outwardly curved ridges. Finally, (f) illustrates a common hieroglyph: just before falling to earth, hexagonal plates and broad sectored plates often develop thick edges called *rims* in [Lib5]. The rim’s inner

boundary corresponds to the transition time in this example, but it is less clear when the vapor level plummeted in the previous two.

A more systematic study of hydrodynamic variability in our simulator would include non-homogeneous schedules not only for the density but also for the six attachment parameters, since they should also be modified by fluctuating temperatures. We anticipate the same general conclusion, however. Whereas meteorological changes undoubtedly influence the development of most snow crystals, we believe that most of the major types can arise from diffusion-limited solidification in a homogeneous vapor cloud. Therefore, we are wary of interpreting specific structural features as signaling environmental changes. In any case, vapor depletion at the crystal boundary plays an essential role in the onset of branching and other morphological transitions.

Consider now the impact of our final parameter σ , which introduces tiny random fluctuations into the diffusing vapor density. Recall that this simple device encodes residual stochasticity on the mesoscopic scale. Figure 16 shows three representative simulations with $\sigma \approx 10^{-5}$. This is a plausible level of noise since, as noted earlier, our mesoscopic cells contain on the order of 10^{10} molecules. Here we will make do with a few intriguing observations about the randomized crystals:

- *Overall symmetry.* Just as in photomicrographs, the visual impression is one of remarkable symmetry even though closer inspection reveals many glitches. Especially in dendritic cases, it seems paradoxical that some seemingly random features exhibit symmetry breaking whereas others are cloned quite precisely from one main branch to the next. We argue that the latter are manifestations of complex hydrodynamics, not extraordinary serendipity.
- *Stability.* Anisotropic attachment apparently makes snowflake morphology stable under very small parameter changes not only in the deterministic case, but also in the randomized model with sufficiently small $\sigma > 0$. (See [GG1] for a simpler deterministic growth model exactly stable under small random perturbations.) Slow, faceted growth is inherently more stable, so hydrodynamic symmetry is better preserved in plates than in dendrites. Evidence of randomness is barely discernible within the aftergrowth of the sectorized branching example in Fig. 16.
- *Attractors.* Careful examination of the faceted dendrite in Fig. 16, and many other random simulations, reveals a remarkable property seemingly shared by some real snowflakes. Very small noise, causing the crystal to diverge slightly from deterministic dynamics, appears to select from a limited number of possible trajectories for branch development. In our faceted dendrite, three main branches are virtually identical, two others are also very similar but quite different from the first group, and the remaining branch differs considerably from any of the others. Such a delicate balance between pattern and randomness may offer a clue to the puzzle of symmetric complexity discussed in Section 2.

This paper has presented many simulation still frames in an attempt to communicate key features and implications of our model. The dynamics are conveyed much more effectively by time-lapse animations that display the entire evolution from a single cell to final form in 15 seconds or so. A dozen such sample movies are available for download, along with the source code for our simulator, from

<http://psoup.math.wisc.edu/Snowfakes.htm>

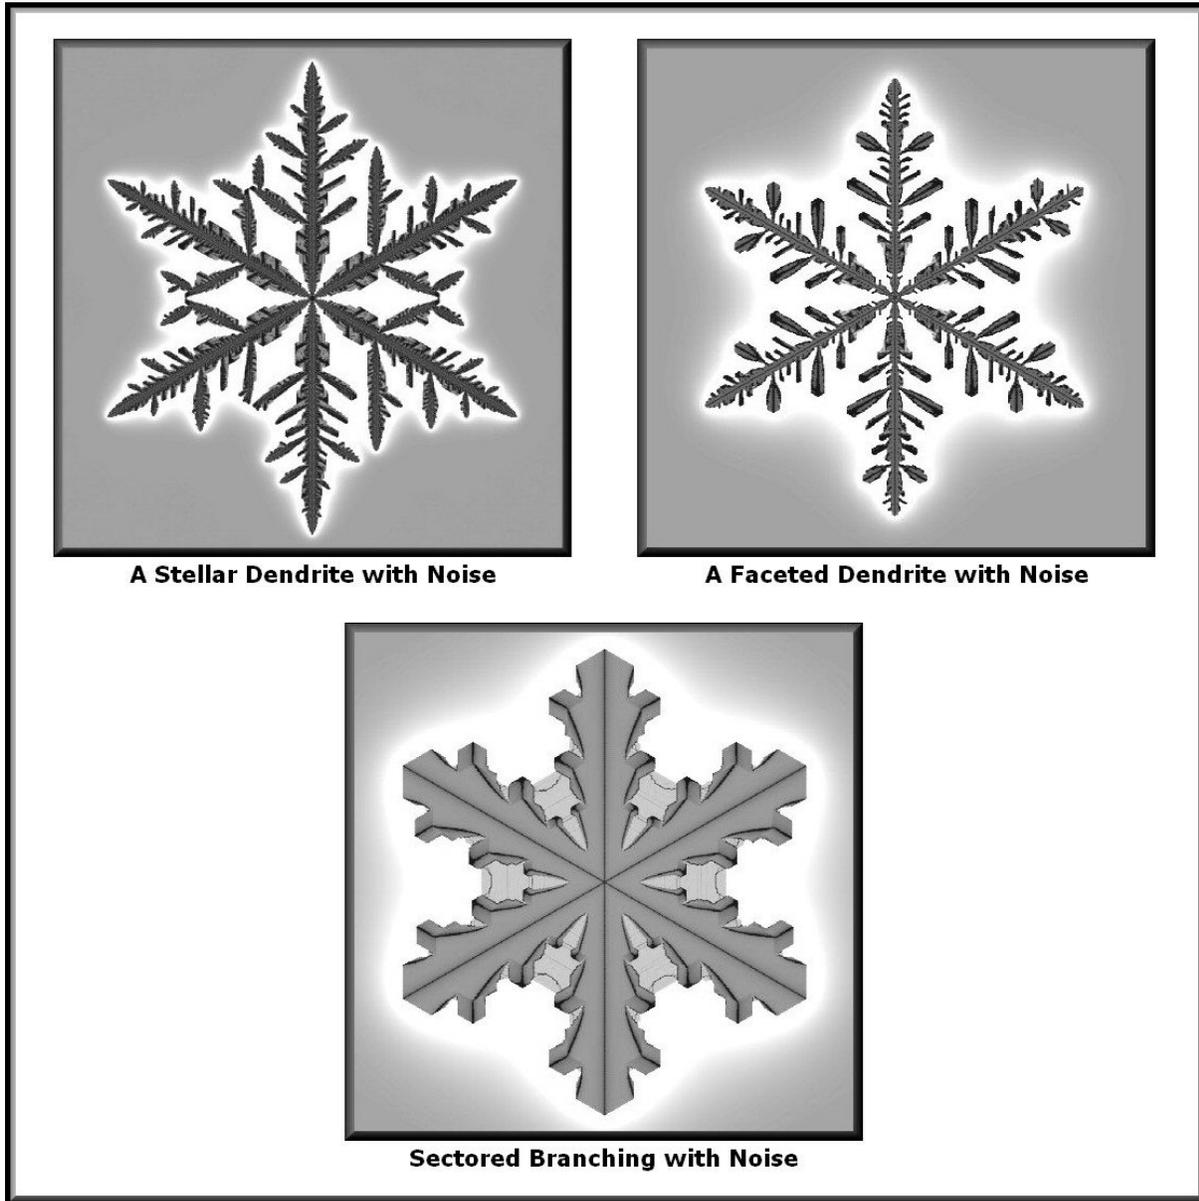


Figure 16. Crystals in an environment with small random noise

7 Remaining riddles and future directions

We conclude by mentioning several avenues for further study of our current simulator and extensions we intend to pursue:

Detailed analysis of the phase portrait

A more in-depth study of the parameter space promises additional insights about macrostep-driven crystal growth. Ideally, we would hope to reach a better understanding of the correspondence between the spatial and temporal units of our simulator and the actual physical scales. A word of caution, though. Even in a homogeneous spatio-temporal environment, snow crystal growth is characterized by very subtle diffusion-limited, finite-size effects. There may be no well-defined phase transitions distinguishing morphological types as parameters vary. For instance, the apparent change from parabolic tips of stellar dendrites to polygonal tips of sectorial plates may only be a matter of scale. It seems likely that classification of snow crystals will always remain to some degree qualitative.

Complexity of side branching

In many dendrite cases, the side branches tend to be of varying lengths. A long side branch is often succeeded by several short ones. Is there any way to predict when long branches will appear? Simulations often suggest a regularity that seems difficult to pin down. In other cases, the lengths of branches and the spacings between them exhibit simple patterns. As parameters vary, the behavior seems to alternate between these two scenarios. It would be illuminating to devise a dynamical system that would isolate and mimic these aspects of our model.

Sublimation

Our algorithm assumes permanent attachment, whereas some degree of sublimation appears in most snowflake photomicrographs. Conversion of ice back to vapor is most prevalent near the ground and during collection, but occurs to a lesser extent throughout the growth process. The effect is somewhat reminiscent of motion by mean curvature, smoothing and convexifying local structure and eroding crystal tips most rapidly. We hope to incorporate a tunable sublimation parameter into future refinements of the simulator. A successful implementation would yield even more realistic snowflakes.

Three-dimensional model

We are developing [GG3] a spatial version of our simulator in order to study mesoscopic growth for a broader range of snow crystals. The seed will be a small hexagonal prism with variable thickness, and we will attempt to reproduce observed features of true three-dimensional diffusion-limited growth that lead to double layers, hollows, and so forth. Of course, this project is more computation-intensive, but a deterministic algorithm that exploits symmetry is feasible. Effective visualization also poses new challenges.

Triangular and 12-sided snowflakes

Two relatively uncommon varieties of planar snowflake, 3-sided and 12-sided, lie outside the scope of this study. Equilateral triangular crystal plates, usually with snub tips, remain largely

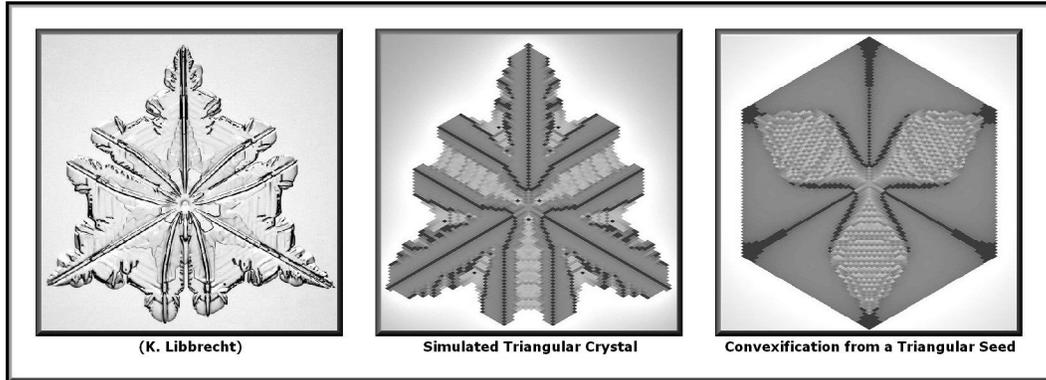


Figure 17. Snowflakes with dihedral symmetry of order 6

a mystery to science. Most likely the evolution of these snowflakes differs fundamentally on the microscopic scale from that of hexagonal prisms. Once they reach a few microns in size, however, some triangular snowflakes revert to more hexagonal forms. Our basic model – started from a very small, symmetric 3-sided star – does a good job of reproducing initial triangular growth that tends toward dihedral symmetry of order 12. (See Fig. 17 in which the third image illustrates the transition from triangular to hexagonal form with large β .) Such simulations shed some light on 3-sided snowflakes but also indicate how sensitive mesoscopic dynamics are to the early crystal geometry. Twelve-sided examples are somewhat better understood [Lib5]. Crystal *twinning* very early in the growth process causes the usual hexagonal prism to develop as two halves subject to a 30° rotation. Ideally, each half then forms an independent dendrite, giving the appearance of dihedral symmetry of order 24 if the prism is suitably thin. We could of course give the prism halves independent simulators at the proper orientation; a more ambitious approach would incorporate twinning into the proposed three-dimensional model.

Hydrodynamic description

External DLA processes such as ours, even in their deterministic mesoscopic incarnation, pose significant challenges when it comes to proving convergence to PDE. This is no surprise — external dynamics, as opposed to their internal counterparts, create irregular boundaries. One can nevertheless expect that the dynamics converges to a continuous limit when the lattice \mathbb{T}' is shrunk by a factor of ϵ , time is sped up by factor ϵ^{-2} , and $\epsilon \rightarrow 0$. The limiting evolution should be described by an evolving vapor density profile $\tilde{\rho} = \tilde{\rho}(x, t)$, $x \in \mathbb{R}^2$, $t \geq 0$ and an increasing family of crystals $\tilde{A}_t \subset \mathbb{R}^2$. Certainly, $\tilde{\rho}$ should satisfy the appropriate diffusion equation off \tilde{A}_t :

$$(6) \quad \tilde{\rho}_t = \frac{2}{7}(\tilde{\rho}_{xx} + \tilde{\rho}_{xy} + \tilde{\rho}_{yy}), \quad x \notin \tilde{A}_t.$$

The evolution of \tilde{A}_t is much more difficult to specify. A natural conjecture would be to suppose that \tilde{A}_t expands at every point of its boundary $\partial\tilde{A}_t$ in the direction of its external unit normal ν_x with velocity

$$(7) \quad w \left(\frac{\partial\tilde{\rho}}{\partial\nu_x}, \nu_x \right).$$

If one believes that the general Stefan problem given by (6) and (7) describes the growth of real snowflakes in a homogeneous environment, then macroscopic effects of the attachment kinetics

are embodied in the function $w = w(\lambda, u)$. Here, $\lambda \geq 0$ is a measure of local supersaturation, and the unit vector u gives the boundary incline. To date, one can only surmise properties of this function from empirical observations. Even in our relatively simple model, the dependence of w on the six parameters (β , κ , α , θ , μ , and γ) seems impossible to describe. For example, take the primitive case (when only ρ and β are nonzero), and u close to the vertical direction, say. As λ varies from large to small, w apparently undergoes a transition from a regime where growth is easier using the macroscopic version of (3a) to a regime where macrostep dynamics are more advantageous — this is the reason for ridges in the leftmost bottom frame of Fig. 4.

Despite these difficulties, or perhaps because of them, the rigorous hydrodynamic limit of the model described in this paper (or a corresponding DLA particle system) is a fascinating open problem. Even existence and uniqueness of the solution to the PDE given by (6) and (7) are bound to be quite an analytical challenge. Once established, this limit should give insight into many intriguing puzzles, of which we mention just one: in the primitive case, how does the time of first instability scale with large β ?

References

- [BBS] N. Bacon, M. Baker, B. Swanson, *Initial stages in the morphological evolution of vapour-grown ice crystals: a laboratory investigation*, Q. J. R. Meteorol. Soc. 129 (2003), 1903–1927.
- [Ben1] W. Bentley, *A study of snow crystals*, Appleton’s Popular Science Monthly 53 (1898), 75–82.
- [Ben2] The Bentley Collection, The Schwerdtfeger Library.
<http://library.ssec.wisc.edu/bentley/index.html>
- [BH] W. A. Bentley and W. J. Humphreys, “Snow Crystals,” McGraw-Hill, 1931 and Dover, 1962.
- [Des] R. Descartes, “Les Météores,” 1637; éd. Adam et Tannery, Paris, Vrin, t. IV, 1965.
- [Dix] R. Dixon, “Mathographics,” Dover, 1991.
- [FPV] F. Family, D. Platt, T. Vicsek, *Deterministic growth model of pattern formation in dendritic solidification*, J. Phys. A: Math. Gen. 20 (1987), L1177–L1183.
- [Fra] F. Frank, *Snow crystals*, Computational Physics 23 (1982), 3–22.
- [FT] N. Fukuta, T. Takahashi, *The growth of atmospheric ice crystals: A summary of findings in vertical supercooled cloud tunnel studies*, Journal of the Atmospheric Sciences 56 (1999), 1963–1979.
- [Gar] M. Gardner, *Mathematical games*, Scientific American 133 (December 1976), 124–128.
- [GG1] J. Gravner, D. Griffeath, *Random growth models with polygonal shapes*, Annals of Probability 34 (2006), 181–218.
- [GG2] J. Gravner, D. Griffeath, *Modeling snow crystal growth I: Rigorous results for Packard’s digital snowflakes*, Experimental Mathematics 15 (2006), to appear.
- [GG3] J. Gravner, D. Griffeath, *Modeling snow crystal growth III: 3d snowflakes*, in preparation.
- [Gle] J. Gleick, “Chaos: Making a New Science,” Penguin Books, 1987.
- [GN] T. Gonda and S. Nakahara, *Dendritic ice crystals with faceted tip growing from the vapor phase*, Journal of Crystal Growth 173 (1997), 189–193.
- [GQ] J. Gravner and J. Quastel, *Internal DLA and the Stefan problem*, Annals of Probability 28 (2000), 1528–1562.
- [Hal] T. Halsey, *Diffusion-limited aggregation: A model for pattern formation*, Physics Today 53 (2000), 36–41.
- [Hob] P. Hobbs, “Ice Physics,” Clarendon Press, Oxford, 1974.
- [Hoo] R. Hooke, “Micrographia,” 1665; Dover, 2003.

- [**Ke**] J. Kepler, “Strena Seu de Nive Sexangula,” 1611. Translated as “The Six-Cornered Snowflake,” trans. Colin Hardie, Clarendon Press, Oxford, 1966.
- [**KL**] T. Kim and M. Lin, *Visual simulation of ice crystal growth*, Eurographics/SIGGRAPH Symposium on Computer Animation, 2003, pp. 86–97.
- [**Kob**] R. Kobayashi, *Modeling and numerical simulations of dendritic crystal growth*, Physica D 63 (1993), 410–423.
- [**Koc**] H. von Koch, *Sur une courbe continue sans tangente, obtenue par une construction géométrique élémentaire*, Arkiv för Matematik, Astronomi och Fysik 1 (1904), 681–702.
- [**LBG**] G. Lawler, M. Bramson and D. Griffeath, *Internal diffusion limited aggregation*, Annals of Probability 20 (1992), 2117–2140.
- [**Lev**] S. Levy, “Artificial Life: A Report from the Frontier where Computers Meet Biology,” Penguin, 1992 and First Vintage Books, 1993.
- [**Lib1**] K. Libbrecht, *Morphogenesis on ice: The physics of snow crystals*, Engineering and Science 1 (2001), 10–19.
- [**Lib2**] K. Libbrecht, *Explaining the formation of thin ice crystal plates with structure-dependent attachment kinetics*, Journal of Crystal Growth 258 (2003), 168–175.
- [**Lib3**] K. Libbrecht, *The physics of snow crystals*, Reports on Progress in Physics 65 (2005), 855–895.
- [**Lib4**] K. Libbrecht, <http://www.its.caltech.edu/~atomic/snowcrystals/>
- [**Lib5**] K. Libbrecht, “Field Guide to Snowflakes,” In preparation, 2006.
- [**LR**] K. Libbrecht, P. Rasmussen, “The Snowflake: Winter’s Secret Beauty.” Voyageur Press, 2003.
- [**ML**] C. Magono and C. Lee, *Meteorological classification of natural snow crystal*, J. Fac. Sci. Hokkaido 2 (1966), 321–335.
- [**MS**] W. Mullins, R. Sekerka, *Morphological stability of a particle growing by diffusion or heat flow*, J. Appl. Phys. 34 (1963), 323–329.
- [**Nak**] U. Nakaya, “Snow Crystals: Natural and Artificial,” Harvard University Press, 1954.
- [**NS**] J. Nittmann, H. Stanley, *Non-deterministic approach to anisotropic growth patterns with continuously tunable morphology: the fractal properties of some real snowflakes*, J. Phys. A: Math. Gen. 20 (1987) L1185–L1191.
- [**Pac**] N. H. Packard, *Lattice models for solidification and aggregation*, Institute for Advanced Study preprint, 1984. Reprinted in “Theory and Application of Cellular Automata,” S. Wolfram, editor, World Scientific, 1986, pp. 305–310.
- [**Rei**] C. Reiter, *A local cellular model for snow crystal growth*, Chaos, Solitons & Fractals 23 (2005), 1111–1119.

- [Sch] A. Schmidt, *Approximation of crystalline dendrite growth in two space dimensions*, Acta Math. Univ. Comenianae 67 (1998), 57–68.
- [She] J. Shedd, *The evolution of the snow crystal*, Monthly Weather Review 47 (1919), 691–703.
- [Stu] R. Stull, “Meteorology for Scientists and Engineers,” Brooks Cole, 2000.
- [Vic] T. Vicsek, “Fractal Growth Phenomena,” Second Edition, World Scientific Publishing Co., 1992.
- [Wet] J. Wettlaufer, *Dynamics of Ice Surfaces*, Interface Science 9 (2001), 117–129.
- [WS] T. Witten, L. Sander, *Diffusion-limited aggregation, a kinetic critical phenomenon*, Phys. Rev. Lett. 47 (1981), 1400–1403.
- [Wol1] S. Wolfram, *Computer software in science and mathematics*, Scientific American 251 (September 1984), 188–203.
- [Wol2] S. Wolfram, “A New Kind of Science,” Wolfram Media, 2002.
- [XAR] R. Xiao, J. Alexander, F. Rosenberger, *Morphological evolution of growing crystals: a Monte Carlo simulation*, Physical Review A 38 (1988), 2447–2456.
- [Yok] E. Yokoyama, *Formation of patterns during growth of snow crystals*, Journal of Crystal Growth 128 (1993), 251–257.
- [YK] E. Yokoyama, T. Kuroda, *Pattern formation in growth of snow crystals occurring in the surface kinetic process and the diffusion process*, Physical Review A 41 (1990), 2038–2050.

8 Appendix: parameter values used in figures

Fig. ^a	ρ	β	α	θ	κ	μ	γ	σ
2tr	.8	2.9	.006	.004	.05	.015	.0001	.00002 ^b
2br	.64	1.6	.21	.0205	.07	.015	.00005	0
4	.58	3.2	0	0	0	0	0	0
6b	.8	2.6	0	0	.05	.015	.0001	.00005 ^b
9		1.3	.08	.025	.003	.07	.00005	0
10	.8		.004	.001	.05	.015	.0001	0
11	.635	1.6	.4	.025		.015	.0005	0
12	.5	1.4	.1	.005	.001		.001	0
13l	.65	1.75	.2	.026	.15	.015	.0001	0
13m	.36	1.09	.01	.0745	.0001	.14	.00001	0
13r	.38	1.06	.35	.112	.001	.14	.0006	0
14r	.37	1.09	.02	.09	.003	.12	.000001	0
15tl	.65	1.8	.6	.067	.001	.05	.0005	.00002 ^b
15tm	.35	1.4	.001	.015	.05	.015	.01	.00005 ^b
15tr	.5	1.3	.08	.025	.003	.07	.00005	-.00001 ^b
15bl	.66	1.6	.4	.025	.075	.015	.00005	0 ^c
15bm	.65	1.75	.2	.026	.15	.015	.00001	0 ^d
15br	.6	1.3	.3	.1	.0001	.1	.0001	0 ^e
16l	.66	1.6	.4	.025	.075	.015	.00005	.000006
16m	.65	1.6	.2	.0245	.1	.015	.00005	.00001
16r	.8	2.6	.006	.005	.05	.015	.0001	.00005
17m ^f	.58	2	.08	.011	.1	.01	.00005	0
17r ^f	.58	3	.04	.02	.1	.01	.00005	0

^aThe letters *t, b, l, r*, and *m* refer to the position of an image within the figure (top, bottom, left, right, and middle).

^bHere, this is the proportion of increase in diffusive mass at each time step, not the noise.

^cVapor decreases by 40% at time 4000.

^dAt time 10,000, vapor decreases by 30% and κ changes to .01.

^eVapor decreases by 40% at time 7500.

^fStarted from a Y-shaped crystal with 16 sites.