BML revisited: Statistical physics, computer simulation and probability

Raissa M. D'Souza
Department of Mechanical and Aeronautical Engineering
Center for Computational Science and Engineering
University of California, Davis, CA 95616
rmdsouza@ucdavis.edu

Abstract

Statistical physics, computer simulation and discrete mathematics are intimately related through the study of shared lattice models. These models lie at the foundation of all three fields, are studied extensively, and can be highly influential. Yet new computational and mathematical tools may challenge even well established beliefs. Consider the BML model, which is a paradigm for modeling self-organized patterns of traffic flow and first-order jamming transitions. Recent findings, on the existence of intermediate states, bring into question the standard understanding of the jamming transition. We review the results and show that the onset of full-jamming can be considerably delayed based on the geometry of the system. We also introduce an asynchronous version of BML, which lacks the self-organizing properties of BML, has none of the puzzling intermediate states, but has a sharp, discontinuous, transition to full jamming. We believe this asynchronous version will be more amenable to rigorous mathematical analysis than standard BML. We discuss additional models, such as bootstrap percolation, the honey-comb dimer model and the rotor-router, all of which exemplify the interplay between the three fields, while also providing cautionary tales. Finally, we synthesize implications for how results from one field may relate to the other, and also implications specific to computer implementations.

1 Introduction

Fundamental to the study of statistical physics is the formulation and analysis of discrete microscopic models, especially discrete lattice models. Understanding phase transitions is grounded in analysis of the Ising and Potts models, as well as percolation on various lattices. Understanding kinetic "critical" phenomena is grounded in analysis of models such as Diffusion Limited Aggregation, and Ballistic Deposition. For a review of these models and their roles in the development of contemporary statistical mechanics, see for instance Ref. [1]. Common amongst all these models is that though simple to formulate, rigorous mathematical results, if any, have been difficult to obtain. This holds particularly true for kinetic models, which are typically non-linear dynamical systems. Rather than mathematical analysis, much progress is made in a complementary way via computer simulation.

Computer simulation is a form of experimental mathematics. The data generated, and often visualization of that data, can provide key insights. New simulation techniques and increasing computational power continually bring new results, which may challenge widely held existing beliefs. Several examples currently illustrate this well, and we expect continued progress along these lines. The models considered can be of broad theoretical interest, while also paradigms for understanding applied physical systems. Herein we present some new results for one such model, the BML model discussed in detail below. BML is a theoretical underpinning for the study of jamming phenomena and congestion patterns of car traffic on roadways.

Though a powerful tool, the lengthscale and timescale attained during simulation may differ greatly from the lengthscale and timescale associated with a physical system or with the near infinite ones typically utilized in rigorous mathematical treatment. Without a clear understanding of distinct regimes and their correspondence, results from one form of analysis may be misleading when extrapolated to another. Recent developments for the bootstrap percolation model, also discussed herein, illustrate this well. Also interspersed throughout this article is discussion of other models which currently illustrate the interplay between statistical physics, computer simulation, and discrete math. (For the latter, we could equivalently use the term combinatorics or probability theory). Of course many other equally interesting models are not discussed, however we believe the ones chosen here highlight strengths and pitfalls, while also the necessity of all three perspectives. In addition, they are all very recent developments. Table 1, found in the concluding section, summarizes the models discussed and their significance.

2 The BML model of traffic

2.1 Self-organization and phase transitions

Several models from statistical physics are used extensively to study traffic: traffic of cars on roadways, or traffic of data packets in computer network queues. In either case, the traffic is composed of discrete entities, and we are interested in predicting flow, congestion patterns and jamming phenomena. One of the most cited examples is the Biham, Middleton, and Levine model (BML) of two-dimensional traffic flow [2]. It is a simple Cellular Automata (CA) model which shows a range of self-organized behaviors. The standard understanding is that the BML model has a first-order phase transition, from free-flowing traffic to fully jammed traffic, as a function of traffic density. Almost all our understanding has come from computer simulation. In fact, rigorous analysis of the BML model is even given as an "unsolved puzzle" in [3]. More importantly BML has become a theoretical underpinning for traffic modeling. Currently there are over two-hundred citations in the scientific literature referencing BML and its first-order phase transition. For recent reviews see Refs. [4, 5, 6]. The influence of BML has even spread to fields beyond traffic modeling, such as management, where it has been suggested as a loose metaphor for information processing in a business organization, capturing lateral flow and upward flow of information.

The BML model consists of two species of "cars" (called "red" and "blue") moving on a two-dimensional square lattice with periodic boundary conditions. Red cars want to move east-ward. Blue cars want to move north-ward. And they alternate attempts to do so. First all the red cars in synchrony attempt to advance one lattice site to the east. Any car succeeds so long as the site it wants to occupy is currently empty (no red or blue car is already occupying it). Then all the blue cars in synchrony attempt the corresponding advance north-ward. Cars that advance are said to have velocity $v_i = 1$. Stationary

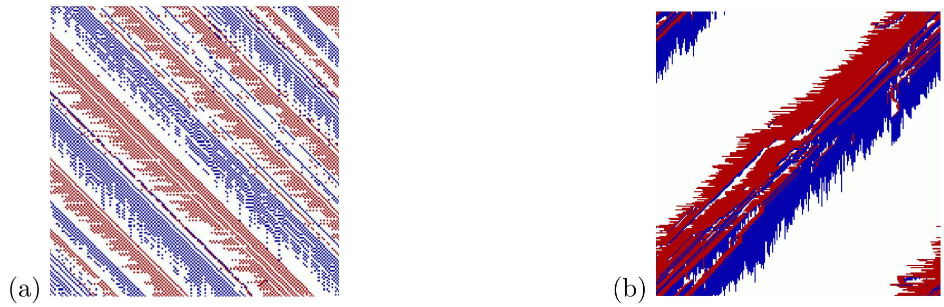

Figure 1: Typical configurations for the BML model on an $L \times L$ lattice, with $L = 256$. (a) The free-flow phase reached by any random initial configuration with small density, ρ. Self-organized bands of separated east- and north-bound cars allow all particles to advance each update, thus $v = 1$. (b) The global jam phase reached by any random initial configuration for larger ρ. Here all cars are immobilized, thus $v = 0$. The conventional belief is that there is a first-order phase transition between these two behaviors as a function of ρ, with critical density $\rho_c \approx 0.35$.

(i.e., "blocked") cars have $v_i = 0$. (We will use the notation v_i to refer to the velocity of an individual car, and the notation v to refer to the average velocity over all cars present on the lattice). We can think of the BML lattice as having a traffic light at each site, with all lights synchronously timed to alternate between east-ward and north-ward flow. The dynamics is fully deterministic. The only randomness is in the initial condition, when an empty lattice is populated uniformly at random with an equal fraction of cars of each species, with total car density ρ.

The model is simple to formulate, yet the behaviors it displays, quite complex. Starting from a random initial condition with small ρ, tiny jams of cars initially form but quickly dissipate. Cars "phase-lock" and separate into non-interacting bands of red cars and bands of blue cars, allowing all cars to move freely at each time step, and a steady-state is reached where each car has $v_i = 1$ (equivalently $v = 1$). An example configuration is shown in Fig. 1(a). Starting from a random initial condition with larger ρ, small jams initially form and fluctuate in size, until eventually one jam dominates and grows to encompass all cars, and a steady-state of complete gridlock is reached where no car can ever move again ($v = 0$). Figure 1(b) shows such an example configuration. An abrupt, discontinuous, transition between these two behaviors has been repeatedly observed, and numerical evidence suggests the critical density for the transition is $\rho_c \approx 0.35$. The evidence also suggests that the value of ρ_c tends to decrease with increasing lattice size, however the behavior of BML on an infinite size lattice is not known.

2.2 Existence of intermediate states

We recently studied the BML model via computer simulation [7]. Surprisingly we found a whole new range of behaviors, that had gone unnoticed up until now. Instead of a phase transition, we found bifurcation points and discovered stable co-existing intermediate states. These states have highly structured geometric patterns of wave fronts of jams moving through otherwise freely flowing traffic. An example is shown in Fig. 2. Furthermore, the BML system is extremely sensitive to the *aspect ratio* of the underlying lattice — the ratio of the number of lattice sites in the horizontal direction, L, to the number in the vertical, L'. The location of the bifurcation points, the range of the windows for phase coexistence, the

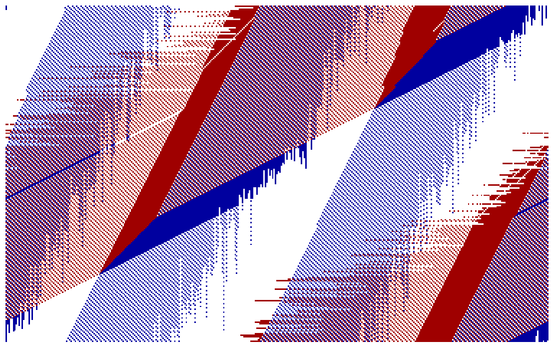

(a)

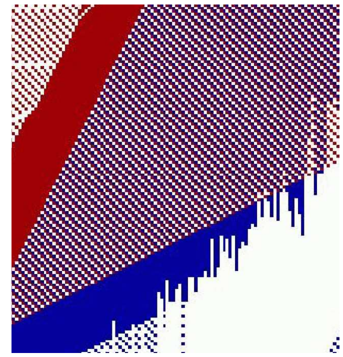

(b)

Figure 2: (a) Typical intermediate configuration on a lattice of width $L = 377$ by height $L' = 233$ (two successive Fibonacci numbers). Bands of free-flowing traffic intersect at jammed wave fronts. (b) Closeup of the region just outside a jam.

exact geometry of the intermediate states, and the number of coexisting phases, depend on the aspect ratio. Typically, when simulating models from statistical physics, we assume that any lattice effects have been averaged out well before we reach the full size of the system. With BML this is not the case, precise details of the underlying lattice manifest themselves at the macroscopic scale. Other researchers have previously presented evidence that BML is sensitive to boundary conditions, without quantifying the dependence (see Ref. [7] for details).

We first discovered the intermediate states on standard lattices with $L = L'$. Here their order is not as precise as in Fig. 2. Bands of free-flow intersect at jams as in Fig. 2, but small chains of particles dispersed at random, move throughout the background space. A natural conjecture would be that these intermediate states are transient, with the expectation that eventually the system will converge to one of the two standard fixed points of free-flow or fully jammed. We simulated over one-hundred independent realizations for a range of densities ρ, and ran each one until convergence to free-flow or fully-jammed, or up to a stopping time of $\tau \sim 10^8$ (which required about a month of computer time). More than half of the realizations were in the latter category. They were still in the intermediate phase at the stopping time, making this phase at least meta-stable since it persists for as long as we can reasonably observe. Moreover, every realization in the intermediate phase had an average velocity over all cars, of $v \sim 2/3$, independent of ρ, making all the velocities observed quantized to either $v = 1$, $v = 0$, or $v \sim 2/3$. For details see [7].

Recent results from the probability community were soon brought to our attention, namely that the honey-comb dimer model shows sensitivity to the lattice aspect ratio [8]. In particular, at the critical point, the partition function depends on the shape of the domain on which it is calculated. At the urging of one of the authors of [8] (D. Wilson), we considered lattices where L and L' are two successive Fibonacci numbers (thus the aspect ratio of the space converges to the golden mean). Implementing such a system, a crisp regular geometry repeatedly emerged; as shown in Fig. 2. The exact microscopic patterns for each independent realization were so precisely ordered, it lead us to test for, and show, that "intermediate" states on such lattices are not transient, they have converged to a fixed-point steady-state behavior which is periodic! The exact microscopic configurations recur every τ timesteps, where τ is on the order of the number of lattice sites ($L \times L'$). Furthermore, the convergence time to the periodic limiting behavior is not very large, about 50,000 timesteps for a system of size 5000 lattice sites.

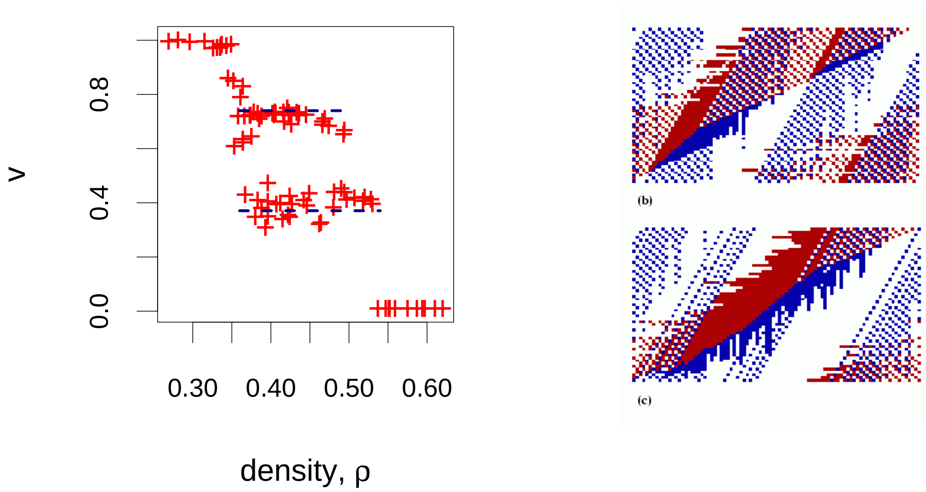

(a)

Figure 3: BML on an $L = 89$ by $L' = 55$ lattice. (a) The average velocity for each individual realization, v, versus the density, ρ, for that realization. Note the bifurcation points, with two distinct coexisting phases (each one is a fixed-point, periodic limit cycle). The dashed lines are the predicted velocities from [7]. The geometries of each phase is distinct, with a typical higher velocity configuration illustrated in (b), and a lower velocity one in (c).

Figure 3 shows the numerical results for simulations of BML on lattices of size $L = 89$ by $L' = 55$. (In [7] results for a limited range of ρ were presented, here we show the results for the full range). The horizontal axis denotes the overall density of cars ρ. The vertical axes denotes the average velocity over all cars in that particular realization, once the system has converged. For very low density, we see the expected behavior of free-flow ($v = 1$). For very high density we see the expected behavior of fully-jammed ($v = 0$). Yet, at intermediate values the system goes to one of two fixed-point behaviors, each of which is a periodic limit cycle. The numerical evidence suggests that the basin of attraction for each fixed-point is quite large. We simulated 70 independent realizations of systems initialized at random with ρ between 0.35 and 0.53. Every one went to one of these two fixed-points, with the more standard states (of free-flow or jam) not once seen. They do not appear to be accessible to systems initialized at random with these densities. Despite extensive testing, we could not identify any characteristic of an initial random configuration that was correlated with distinguishing between the two attractors. Also we do not know the stability of flipping from one attractor to the other. In [7] we were able to formulate a self-consistent set of equations, based on the macroscopic constraints imposed by the underlying lattice, which allowed us to calculate the asymptotic velocities of the observed intermediate states. The predictions from this theory are shown as the dashed lines in Fig. 3(a).

All the simulations reported in Ref. [7] were done on a special purpose cellular automata (CA) computer, the CAM8 [9]. Since publication of [7], we have replicated all results on standard personal computers using two different open source CA programs, Mirek's Cellebration [10], and Trend/jTrend [11, 12]. (The results presented herein were obtained with jTrend.) Results have since also been replicated by other researchers [13]. Further information and dynamic images can be found in [14].

2.3 Synchronous clocks and global coordination

With any Cellular Automata, it is natural to question the role of synchrony and global coordination. Full synchronization requires that the entire system move in lock-step, to the beat of one global clock. Though reasonable in some situations, typically we expect local coordination, not global. Recall that the BML lattice can be regarded as having a traffic light at each site, with all lights initialized in the same east-flow phase, and timed to synchronously oscillate between east- and north-flow. We first test removal of of full synchrony. Then we restore synchrony, but test removal of global coordination during the initialization of the system. Either of these changes to the BML model destroys the self-organization that leads to the free-flow state shown in Fig. 1(a). Instead, small jams continually form and dissipate, but persist indefinitely, leading to a steady-state velocity $v < 1$. The characteristic size of these jams increases smoothly with density for low density, but we then see a sharp, discontinuous, transition to fully jammed. This abrupt onset of jamming occurs at a much lower density than for the conventional BML model. In addition, with either of these changes, the intermediate states discussed in Sec. 2.2 are no longer observed. Thus the self-organizing structures and properties of the model are lost.

We first remove full synchrony. Traffic lights still oscillate in unison, but now, whether a car advances or not during the appropriate phase is determined by the flip of a random, unbiased, coin. (Time is still a discrete variable, but cars no longer all advance synchronously.) The number of moves attempted by any car should now be a Poisson distributed random variable. As an example, consider the situation where a red and a blue car are both poised to move into the same unoccupied site. Previously, under the conventional BML dynamics, the red (east-bound) car would move first, taking over the empty site, and thus blocking the blue (north-bound) car. With the random update, it is equally likely that the red car does not advance during its phase, thus the blue car could move first and block the red.

Figure 4 shows results for simulations of this asynchronous Poisson updating on lattices of $L = 89$ by $L' = 55$, allowing direct comparison to results presented in Fig. 3. The horizontal axis denotes the overall density of cars ρ. The vertical axes denotes the average velocity over all cars in that particular realization, once the system has converged. For small densities, we see that $v < 1$ and decreases smoothly with increasing ρ. The behavior seems to be described approximately by the expression $v \sim (1 - \rho^{1/2})$. Then there is an abrupt onset of full jamming ($v = 0$), which occurs approximately at the value $\rho'_c = 0.15$. For all higher densities, $v = 0$. This value at onset, ρ'_c, appears to have a length dependence, and to decrease with increasing system size (also seen for simulations of standard BML). On larger systems (with $L = 144$, $L' = 89$), $\rho'_c \approx 0.13$. For smaller systems (with $L = 55$, $L' = 34$), $\rho'_c \approx 0.155$.

We then restore full synchrony (the global clock), but we remove global coordination. Instead of initializing all traffic lights in the same east-flow state, we initialize based on the flip of an unbiased coin at each site. If heads, we initialize in east-flow. If tails, we initialize in north-flow. All lights still oscillate between the two states in synchrony. Here we find almost identical results. We observe a smooth decrease in v with ρ for small ρ and then the abrupt onset of full jamming. Results for $L = 89$, $L' = 55$, are shown in Fig. 4. Onset of jamming occurs approximately at $\rho'_c = 0.125$.

We did find one perturbation to the BML dynamics which preserves the free-flow states, yet destroys the intermediate states. We implement the standard BML model but introduce an extremely small probability for a car to flip from east- to north-bound (and vice versa), at each update. For instance, we make the probability $p = 1/(LL')$. For small ρ, spatially separated bands of east- and north-bound cars still self-organize. There is some disorder but at a course-grained level the same stripes as shown in

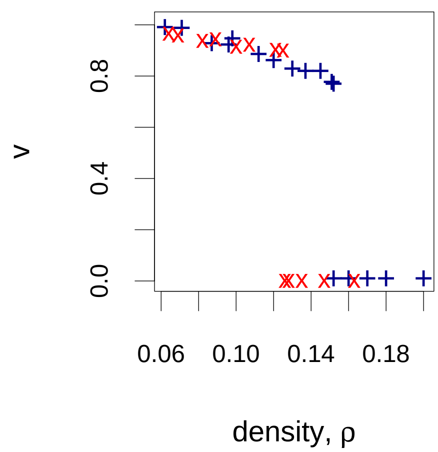

Figure 4: Variants of BML. The "+" are results for asynchronous BML, with random coin-flip (i.e., Poisson) updating. The "x" are results for BML with the traffic lights initialized at random. Both variants lack the self-organized structures observed for BML, yet show an abrupt onset of full jamming, with $v = 0$. For Poisson updating, $\rho_c' \approx 0.15$, for random initialization, $\rho_c' \approx 0.125$. Note, as in Fig. 3, results are for lattices of size 89×55.

Fig. 1(a) appear. Yet the intermediate states are not observed, even on the "Fibonacci"-lattices (lattices with two subsequent Fibonacci numbers as the vertical and horizontal lengths).

Lack of synchrony has previously been shown to destroy self-organized patterns. One of the more prominent instances involves the spatial prisoner's dilemma [15]. With synchrony, complex fractal structures form, which were thought to be inherent to the model, and to have implications for the evolution of cooperation of distinct species in an ecosystem. However shortly after introduction of the model, it was shown that without synchrony, no structure forms [16]. In the BML model, synchrony and global coordination are fundamental for getting to the free-flow state exemplified in Fig. 1(a) and also the periodic limit cycle exemplified in Fig. 2. Whether synchrony and global coordination are reasonable assumptions certainly depends on the physical scenario in mind. In the case of modeling traffic flow in an actual city, perhaps some in-between, local coordination, would be most reasonable.

Yet there are systems where complex structures persist even without a global clock. And, from a purely mathematical perspective, replacing a global clock with a probabilistic one can radically change the level of difficulty associated with analyzing a model. For instance consider the case of iterated games, like checkers or chess, where players alternate turns. The deterministic updating can be replaced by a random one, with a coin tossed each update to decide which player gets the next move. It was recently shown that some such games, for instance Hex, can be famously difficult to analyze if played in the strict alternating turn format. Yet, if replaced with the random-turn updating, a simple optimal strategy can be provably shown to exist [17].

2.4 First rigorous mathematical results

Since the time of publication of [7], the first rigorous result has been obtained for BML [13]. Previous mathematical results relied on partly non-rigorous methods, such as mean field theory. The result in [13] can be paraphrased as follows. Consider the BML model implemented on an infinite lattice or on an $L \times L'$ lattice, with periodic boundary conditions (i.e., an $L \times L'$ torus), where $L, L' \to \infty$, while the ratio L/L' remains finite. There exists $\rho_j < 1$ (which is close to, but strictly less than, one), such that for any $\rho > \rho_j$ almost surely the system jams. (Technically, no car moves infinitely often, and the overall configuration is eventually constant). This proof considers the chain of successive cars blocking any one car (called a blocking path). If $\rho = 1$ there are no empty spaces on the lattice, so there are an infinite number of blocking paths. If there is a sufficiently low proportion of empty spaces, many of the blocking paths survive, and in fact form a network of paths, which block any given car. The proof uses techniques from percolation theory, considering the super-critical regime for oriented percolation. For precise formulation and details of the proof see [13]. An interesting aspect is that the proof is robust to asynchrony – it holds even if random coin-flip updating is used (similar to that in Sec. 2.3).

Note that ρ_j is very close to 1, thus much larger than the value at the conjectured phase transition of $\rho_c \approx 0.35$. Nonetheless, it gives us some hope that further rigorous results will be obtained for the BML model. The proof holds for BML on an infinite torus, while simulations are done on finite tori, and the correspondence between the two is not known. In the infinite BML system we suspect that we would not see the extent of self-organization leading to full free-flow or to the ordered intermediate states. Given the lack of self-organized structures observed for the asynchronous variants of BML, their smooth decay in steady-state velocity with ρ for small ρ, and that the proof discussed above holds for random coin-flip updating, we believe the asynchronous, Poisson update version presented in Sec. 2.3 may be more amenable to rigorous mathematical treatment the the conventional BML model, and leave this as an open challenge.

Most recently, the following results have been shown [18]. Consider an $L \times L$ torus (hence, the aspect ratio is a square). If there are fewer than $L/2$ cars, all possible initial configurations will "self-organize" (i.e., attain $v = 1$) in finite time. If we consider instead choosing the initial condition uniformly at random, if there are fewer than $N \log N$ cars, it can be shown that the system almost surely does not jam (i.e., attain $v = 0$). The proofs rely on considering the occupancy of adjacent diagonals along the lattice, for details see Ref. [18]. Note that $L/2$ cars (or even $N \log N$ cars) corresponds to a density $\rho = 1/(2L)$ (or, respectively, $\rho = \log N/N$), which is much smaller than densities typically of interest. Yet, as with Ref. [13], such clever new proofs give us hope that more results for the BML model will be attained.

3 Bootstrap percolation

The tale of bootstrap percolation uniquely illustrates other aspects of the interplay between statistical physics models, computer simulation and discrete mathematics. The most striking being the "notorious" lack of agreement between the extensive body of work via computer simulation and recently attained rigorous mathematical results (for a review see, for instance [19], and references therein). Bootstrap percolation [20] is a model for nucleation and growth, which is used extensively to model general jamming phenomena and dynamical arrest. The BML model, discussed so far, also models granular flow and jamming, but only of a simple, specific type. Jamming phenomena and dynamical arrest are of much broader general interest, as they are thought to lead to the formation of complex condensed states of matter, such as glasses and gels, and other long-lived non-equilibrium solids.

The essence of the bootstrap percolation process is that an empty lattice is initially populated at random with some specified density of particles, ρ. Then a dynamical spreading process is considered — if an empty site has c neighbors who are populated, the empty site itself becomes populated. This process is iterated until a final steady-state configuration is reached. We can ask basic questions once the process completes, such as is the entire lattice populated? More relevant to complex materials, we might ask, is there a resulting connected path of populated sites that spans (i.e., percolates) across the lattice? Both behaviors seem to exhibit a sharp onset as a function of initial density ρ, with the exact nature depending on the dimension and size of the lattice and the value of c.

We focus on the case of a square lattice in two-dimensions, with $c = 2$, and whether or not the full lattice is populated at the conclusion of the bootstrap process. Here the onset is reminiscent of a first-order process, with a sharp discontinuous transition. This onset had been studied extensively via numerical simulations, and a general consensus reached on the relation between system size and critical initial occupation density. Recently, rigorous mathematical results were obtained [21]. The puzzling aspect is that they are in great discrepancy with the values obtained via computer simulation. Even on lattices of sizes up to $L = 128,000$ values from simulation differ from analytic results [22]. This generated great topical interest, and a reexamination of the fundamental issues. Some form of resolution, and new techniques for simulating bootstrap percolation, have since elegantly been presented in Refs. [19] and [22]. Resolution lies in understanding that theoretical results are derived in the asymptotic limit, as the system size approaches infinite. It is estimated that simulations would have to be carried out on a lattice of length $L \sim 10^{47}$, for precise correspondence between simulation and theory to occur. A lattice of such $L \times L$ size is clearly out of the realm of any computer. But moreover, this is likely a lengthscale beyond the realm of any physical system of interest! Thus whether results from theory or simulation apply, depends upon the regime of interest, bearing in mind that neither accurately captures the typical lengthscales in physical systems: rigorous results are for lengthscales beyond the realm of physical systems, and those from current simulations are below that realm.

Also introduced in [19] and [22] is a new importance-sampling procedure, with bootstrap percolation formulated in terms of connected and disconnected "holes". To determine this topological feature of a hole requires examining a lengthscale much smaller than the system size. In addition, existence of holes are rare events. Thus, instead of simulating the full compute-intensive bootstrap process on a lattice, one can directly sample holes, allowing simulation of systems far larger than could have been accessed previously. They also develop a system of closed-form equations, similar to a path integral formulation, which agree with the rigorous results in the asymptotic realm, and with computer simulations in finite realms. It is surprising in retrospect, that bootstrap percolation appears to be the first well-known instance where such large discrepancies between simulation and theory have been encountered.

4 Rotor-router

Another interesting model to consider is the rotor-router, invented by J. Propp, and introduced in Ref. [23]. It is essentially a deterministic version of Internal Diffusion Limited Aggregation (IDLA) [24, 25]. Elucidating the connections between stochastic statistical physics models and their deterministic variants

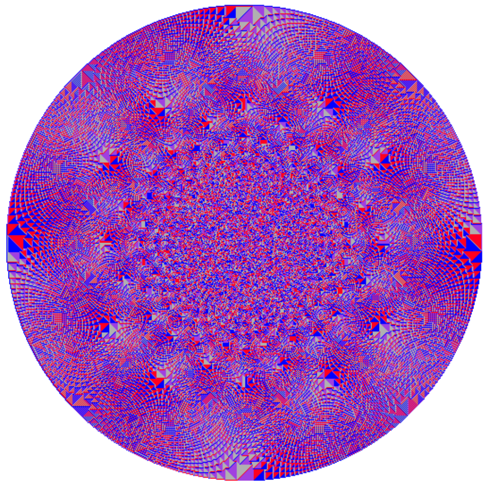

Figure 5: The rotor-router after aggregation of $N = 200,000$ particles. (Image courtesy of L. Levine.)

can be extremely useful. In some situations (if the deterministic version is also reversible) it allows us to explore the underlying, far from equilibrium, thermodynamics of the dynamical system. See, for instance, the Reversible DLA model [26]. And, as discussed in Sec. 2.3, probabilistic and deterministic versions of dynamics, such as iterated games, can have radically different levels of difficulty for analysis.

The rotor-router is a model of particles executing deterministic walks on a two-dimensional square lattice. At each lattice site, there is a binary variable corresponding to whether or not that site is occupied. All occupied sites also have a rotor, which can point in one of the four lattice directions (north, south, east or west). All lattice sites are initialized in the unoccupied state. Particles are introduced at the origin and execute the following "rotor-walk". If a particle arrives at an "unoccupied" site, it occupies it (i.e., it stops there and remains forever) and introduces a rotor at the site initialized pointing north. If, on the other hand, it arrives at an "occupied" site, the following happens. The state of the rotor at that site gets rotated by 90-degrees clockwise, then the particle exits in the newly indicated direction, entering the corresponding neighbor site, where the process is iterated. A particle thus walks on the lattice, directed by the states of the rotors, until reaching the first unoccupied site and sticks there. For instance, the first particle injected occupies the origin and introduces a rotor there pointing north. The second particle injected rotates that rotor at the origin from north to east, then moves into the east neighbor site, where it stops while concurrently introducing a north-rotor. The state of the rotors after the aggregation of 200,000 particles is shown in Fig. 5. There are four colors, indicating the four directions that any rotor can point. Note the dynamics of IDLA are almost identical, except that particles leave an occupied site in a randomly chosen direction, rather than being directed by a rotor.

There are many intriguing properties about the rotor-router. First, it is abelian. Consider the state of the rotors after N particles have completed their walks. We could have introduced these particles one-at-a-time, requiring that the predecessor particle complete its walk before the next particle is injected. Alternately, we could have injected all N particles at the same time and allowed them to walk in parallel (called a "rotor-router swarm" in [23]). In the latter case, whenever a site is occupied by more than one particle, one of those particles is chosen uniformly at random and updated. Either way, once all the dynamics conclude, the end state of the routers will be completely identical.

The abelian properties of the model can be shown rigorously, yet far more interesting properties can not. It has been observed through simulation that the rotor-router shape, as shown in Fig. 5, is almost a perfect circular disk. Recently, considerable progress towards a rigorous proof on the shape of the rotor-router has been made (as discussed below), yet proving perfect circularity remains an outstanding challenge. Far more complex is the fine-structure internal to this disk, which appears to have fractal and self-similar properties. In fact, it has been shown that during the evolution to the final state a rotor-router swarm can realize configurations from the abelian sandpile model [23]. (Of course, this intermediate behavior says nothing about the final structure.) In addition, if we simulate a large rotor-router instance and record the angle between each two subsequent stopping sites, we find the distribution of angles is strongly peaked at 90-degrees [27]. None of these properties have any explanation.

There is a shape theorem for IDLA which allows us to define inner and outer radius terms such that, in the limit $n \to \infty$, a rescaled IDLA shape is contained between two disks, one of radius $n - \delta_I(n)$, and the other of radius $n + \delta_O(n)$ [25]. Thus the IDLA shape is circular, with boundary fluctuations that are on the order of $\delta_I(n) + \delta_O(n)$. The best provable statement known is that these fluctuations are on the order of $n^{1/3}$ [28], while the best simulation estimates suggest they are logarithmic in n [29].

The likely first step towards a corresponding shape theorem for the rotor-router has just been made, which comes after considerable recent interest and effort. In [30], the authors prove that in the limit where $n \to \infty$, the shape of the d-dimensional rotor-router, properly rescaled, converges to the d-dimensional unit sphere. Technically the volume of the symmetric difference between the two shapes tends to zero. The proof is not as precise as the corresponding one for IDLA, as it does not address boundary fluctuations. It does not preclude the existence of internal holes or of long tendrils hanging off the circular boundary, provided they are of negligible volume when compared to the bulk. We know from simulation that either structure is extremely unlikely, as simulations indicate the rotor-router shape is a perfect circular disk. More precisely, after the aggregation of one million rotor-walkers, the largest difference between inner and outer radius observed is less than two *lattice sites* in size [31] — the rotor-router shape is about as close to a perfect circle as one can get on a discrete lattice.

The rotor-router has generated much interest, yet more questions than results exist. Several open questions about the basic shape of the rotor-router are simple to state, but seemingly very hard to answer. Simpler than proving circularity, we might ask, is the difference between the inner radius and the outer radius bounded? The proof in [30] does not even allow us to bound the *ratio* of the two, due to the possibility of the tendrils [32].

Unlike the BML and bootstrap percolation models, which originated in the statistical physics community, the rotor-router originated in the discrete math community. Given its many interesting properties, it is intriguing to attempt to find applications for modeling physical phenomena. We know that derandomizing dynamics can reduce error in simulation, for instance, with certain additional constraints, the rotor-router can simulate a random walk instance with less error (i.e., closer to the average behavior) than any single random instance of a random walk [23, 33].

Model Discussed	Significance
BML [2]	Dynamical jamming transition and traffic flow.
Honey-comb dimer model [8]	Importance of aspect ratio.
Spatial Prisoner's Dilemma [15, 16]	Role of synchrony.
Random-turn games [17]	Randomizing a deterministic system.
Bootstrap Percolation [21, 22, 19]	Importance of lenthscales and timescales.
Rotor-router [23]	To be discovered?

Table 1: Summary of models discussed herein, and the reason why the model is discussed.

5 Discussion and Implications

Statistical physics, computer simulation and discrete mathematics are all intimately related through the study of lattice models. Progress in each field hinges on the interplay and dialogue between all three. But to make the dialogue fruitful, many caveats must be kept in mind. First, the lengthscale and timescale of interest in physical systems, in computer simulations, and in mathematical treatments can differ. We need to understand whether or not there is correspondence between them, or results from one will be misleading if applied to the other. Second, we have to understand the role of synchrony and global coordination, and in which regimes they are necessary and reasonable assumptions. Both appear to be essential components for certain types of self-organization, as shown in Sec. 2.3. More generally, it seems synchrony tends to simplify the formulation and computer simulation of a model (such as allowing for highly efficient implementation on parallel computers), yet to complicate rigorous analysis.

Implications specific for computer simulation have also emerged from recent developments. Foremost, as others have recently been urging [34], we need to replicate the highly influential models. Consider the BML model, which has been studied extensively since its introduction in the early 1990's. Why were the complex intermediate states not noticed before? The basins of attraction of these states comprise a significant portion of the phase-space, for an extensive range of ρ, even on the standard lattices traditionally studied (with square aspect ratios). Surely modern computing tools, such as visualization, give us an advantage over our predecessors. Yet deeper more ingrained biases and traditions may have prevented their detection. Studying lattices with different aspects ratios, may be an important component in understanding the behaviors of models from statistical physics, such as BML. There are also implications particular to the BML model when applied to the study of traffic flow. We show there is substantial flow, with large average velocity, in a regime were previously it was believed there was no flow and zero throughput. In addition, the onset of full-jamming can be delayed based on the geometry of the situation.

Other fundamental aspects include, of course, random number generation. This is especially important for stochastic dynamical systems, not so much for the BML model or for bootstrap percolation, which are deterministic once a random initial condition is specified. In the stochastic case, the dynamics of a pseudo-random number generator (PRNG) can couple in unidentifiable ways to the dynamics being simulated, producing results that may not correspond the to real phenomena of interest. The only way to identify this emergent coupling is through statistical testing with various distinct PRNGs, that each individually pass the battery of standard statistical tests. For an example, consider simulations of Ballistic Deposition discussed in [35].

6 Acknowledgements

Thanks to Lionel Levine for providing Fig. 5 as well as discussions and feedback on the rotor-router, to Melanie Mitchell for bringing to the author's attention the role of asynchrony in the Spatial Prisoner's Dilemma and to the anonymous referee.

References

[1] L. P. Kadanoff. *Statistical Physics, Statics, Dynamics and Renormalization.* World-Scientific, 2000.

[2] O. Biham, A. A. Middleton, and D. Levine. Self organization and a dynamical transition in traffic flow models. *Phys. Rev. A*, 46:R6124, 1992.

[3] P. Winkler. *Mathematical puzzles: a connoisseur's collection.* A. K. Peters, Ltd, Natick, MA, 2004.

[4] T. Nagatani. The physics of traffic jams. *Rep. Prog. Phys.*, 65(9):1331–1386, 2002.

[5] A. Schadschneider. Traffic flow: a statistical physics point of view. *Physica A*, 313(1-2):153–187, 2002.

[6] D. Chowdhury, L. Santen, and A. Schadschneider. Statistical physics of vehicular traffic and some related systems. *Phys. Rep.*, 329(4-6):199–329, 2000.

[7] R. M. D'Souza. Coexisting phases and lattice dependence of a cellular automata model for traffic flow. *Phys. Rev. E*, 71, 2005.

[8] R. W. Kenyon and D. B. Wilson. Critical resonance in the non-intersecting lattice path model. *Probability Theory and Related Fields*, 130(3):289–318, 2004.

[9] N. H. Margolus. CAM-8: A computer architecture based on cellular automata. In *Pattern Formation and Lattice-Gas Automata.* American Mathematical Society, 1996.

[10] M. Wojtowicz. MCell, version 4.2. *http://www.mirekw.com/ca/index.html.*

[11] H.-H. Chou, W. Huang, M. Burman, and R. K. Hansen. jTrend, version 1.1. *http://www.complex.iastate.edu/download/Trend/index.html.*

[12] H.-H. Chou, W. Huang, and J. A. Reggia. The Trend Cellular Automata Programming Environment. *Simulation: Transactions of The Society for Modeling and Simulation International*, 78(2):59–75, 2002.

[13] O. Angel, A. E. Holroyd, and J. B. Martin. The jammed phase of the Biham-Middleton-Levine traffic model. *Elec. Commun. in Probability*, 10(17):167–178, 2005.

[14] For further information on open questions and movies of the BML model see *http://mae.ucdavis.edu/dsouza/bml.html*.

[15] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.

[16] B. A. Huberman and N. S. Glance. Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci.*, 90:7716–7718, 1993.

[17] Y. Peres, O. Schramm, S. Sheffield, and D. B. Wilson. Random-turn Hex and other selection games. *arXiv:math.PR/0508580*, 2005.

[18] T. D. Austin and I. Benjamini. For what number of cars must self organization occur in the Biham-Middleton-Levine traffic model for any possible starting configuration? *arxiv.org/abs/math.CO/0607759*, 2006.

[19] P. De Gregorio, A. Lawlor, P. Bradley, and K. A. Dawson. Exact solution of a jamming transition: Closed equations for a bootstrap percolation problem. *Proc. Natl. Acad. Sci.*, 102(16):5669–5673, 2005.

[20] J. Adler. Bootstrap percolation. *Physica A*, 171:453–470, 1991.

[21] A. E. Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probability Theory and Related Fields*, 125(2):195–224, 2003.

[22] P. De Gregorio, A. Lawlor, P. Bradley, and K. A. Dawson. Clarification of the bootstrap percolation paradox. *Phys. Rev. Lett.*, 93(025501), 2004.

[23] M. Kleber. Goldbug variations. *The Mathematical Intelligencer*, 27(1):55–63, 2005.

[24] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Semin. Mat. Univ. Pol. Torino*, 49(1):95–119, 1991.

[25] G. Lawler, M. Bramson, and D. Griffeath. Internal diffusion limited aggregation. *Ann. Probab.*, 20(4):2117–2140, 1992.

[26] R. M. D'Souza and N. H. Margolus. Thermodynamically reversible generalization of diffusion limited aggregation. *Phys. Rev. E*, 60(1):264–274, 1999.

[27] L. Levine. Personal communication.

[28] G. Lawler. Subdiffusive fluctuations for internal diffusion limited aggregation. *Ann. Probab.*, 23(1):71–86, 1995.

[29] C. Moore and J. Machta. Internal diffusion-limited aggregation: Parallel algorithms and complexity. *J. Stat. Phys.*, 99:661–690, 2000.

[30] L. Levine and Y. Peres. Spherical asymptotics for the rotor-router model in $\mathbf{Z}^d$. *arXiv:math.PR/0503251*, 2005.

[31] L. Levine and Y. Peres. The rotor-router shape is spherical. *The Mathematical Intelligencer*, to appear.

[32] New results restricting the existence of tendrils and establishing a shape theorem for the rotor-router exist, but have yet to appear in print. Personal communication, L. Levine.

[33] J. Cooper and J. Spencer. Simulating a random walk with constant error. *Combinatorics, Probability and Computing*, to appear.

[34] M. Mitchell. The prospects and perils of complex systems modeling. *UC Davis, Hydrologic Sciences Graduate Group Colloquium*, Nov. 10, 2005.

[35] R. M. D'Souza, Y. Bar-Yam, and M. Kardar. Sensitivity of Ballistic Deposition to pseudo-random number generators. *Phys. Rev. E*, 57(5), 1998.