Efficient Generation of Model Bulk Heterojunction Morphologies for Organic Photovoltaic Device Modeling

Michael C. Heiber[1,2,*] and Ali Dhinojwala[1,†]

[1]*Department of Polymer Science, The University of Akron, Goodyear Polymer Center,
Akron, Ohio 44325, USA*
[2]*Experimental Physics VI, Julius-Maximilians-University of Würzburg,
Am Hubland, D-97074 Würzburg, Germany*

Kinetic Monte Carlo (KMC) simulations have been previously used to model and understand a wide range of behaviors in bulk heterojunction (BHJ) organic photovoltaic devices, from fundamental mechanisms to full device performance. One particularly unique and valuable aspect of this type of modeling technique is the ability to explicitly implement models for the bicontinuous nanostructured morphology present in these devices. For this purpose, an Ising-based method for creating model BHJ morphologies has become prevalent. However, this technique can be computationally expensive, and a detailed characterization of this method has not yet been published. Here, we perform a thorough characterization of this method and describe how to efficiently generate controlled model BHJ morphologies. We show how the interaction energy affects the tortuosity of the interconnected domains and the resulting charge-transport behavior in KMC simulations. We also demonstrate how to dramatically reduce calculation time by several orders of magnitude without detrimentally affecting the resulting morphologies. In the end, we propose standard conditions for generating model morphologies and introduce an open-source software tool. These developments to the Ising method provide a strong foundation for future simulation and modeling of BHJ organic photovoltaic devices that will lead to a more detailed understanding of the important link between morphological features and device performance.

I. INTRODUCTION

Organic photovoltaics (OPVs) have received a great deal of attention over the last decade. In this time, research efforts covering a wide range of challenges have pushed the power-conversion efficiency of these devices from approximately 3% to approximately 11% [1]. Among these efforts, modeling and simulation have been important for testing our understanding of the fundamental physics of device operation and directing experimental efforts towards new and improved devices. Within the wide range of methods available, kinetic Monte Carlo (KMC) simulations are unique in their ability to incorporate nanoscale details while maintaining the ability to simulate a complete device. The complex nanoscale morphologies present in bulk heterojunction (BHJ) solar cells have been repeatedly shown to have a significant impact on device performance for a number of different donor-acceptor combinations, including polymer-fullerene blends [2], polymer-polymer blends [3], and small-molecule blends [4]. As a result, retaining nanoscale detail in KMC simulations is particularly critical for incorporating morphological features that

can be used to help understand how morphology affects device performance in greater detail.

In an attempt to generate model morphologies for small-molecule blends, Peumans *et al.* [5] introduced a method that utilizes the Kawasaki spin-exchange Ising model [6], which had previously been used to simulate phase separation in binary alloys [7,8]. This concept was then later simplified and applied to KMC simulations by Watkins *et al.* [9]. Since these pioneering studies, Ising-based morphologies have been used in KMC simulations to study a wide range of important OPV topics, from detailed studies on exciton diffusion and dissociation [3,10,11], charge separation and geminate recombination [12–15], bimolecular recombination [16], surface recombination [17], and charge injection [18,19], to broader studies on overall photocurrent generation [20–22] and complete current-voltage curve modeling [19,23–26]. Ising-based morphologies have also been used in master equation device modeling [27].

Several additional morphology models have also been used in KMC simulations, including a chain reptation model [28,29] and a Cahn-Hilliard model [30–32]. In addition, the Ising model has been adapted to produce morphologies similar to those measured by neutron reflectivity and neutron-scattering experiments [33]. While the Ising model may not accurately capture all morphological

---
*michael.heiber@physik.uni-wuerzburg.de
†ali4@uakron.edu

014008-1

features in all donor-acceptor blends used in OPVs, qualitatively, it produces a nanoscale bicontinuous morphology typical of many blends. As a result, this model has served as a reasonable approximation of BHJ morphologies and has become the dominant morphology model in the field.

Nonetheless, a rigorous characterization of this morphology generation technique has not been published. Understanding the details of this method, developing standard procedures, and making the technique openly available will ensure that accurate benchmarks and meaningful comparisons are being made. In the Ising model, the main parameter that can be used to control the phase separation process is the interaction energy. However, it is still unclear exactly how changing this parameter impacts the morphologies generated and the resulting simulated device performance.

In addition, making this technique more computationally efficient has the potential to significantly reduce calculation time and make the method widely accessible. Currently, the main computational challenge is that creating domain sizes typical of optimized devices (> 10 nm) while retaining a high resolution (1 nm) can take a considerable amount of calculation time, especially when a large lattice size is needed. This challenge has, so far, greatly limited the ability to systematically study the effects of morphological features on OPV device simulations. To address these issues, we present a thorough characterization of the morphologies generated using the Ising method, introduce methods to dramatically reduce the calculation time, and show how these changes impact KMC device simulations.

II. METHODS

A. Ising phase separation algorithm

When generating a morphology using the Ising model, a three-dimensional lattice is created, and the sites are randomly assigned as either donor or acceptor sites. Here, a 50:50 blend is implemented, and each site is defined to represent 1 nm^3. Periodic boundary conditions are used in the x and y directions, and hard boundaries are used in the z direction to represent a thin film. Next, a simulated phase separation process is executed in which the total energy of the system is allowed to relax over a series of iterations by allowing adjacent sites to be swapped.

To execute the phase separation process, a pair of adjacent sites with differing types is randomly selected from the lattice. Then, the total change in energy of the system that will result from swapping them $\Delta\epsilon$ is calculated and used to determine the probability of the swapping event,

$$P(\Delta\epsilon) = \frac{\exp[-\Delta\epsilon/(kT)]}{1 + \exp[-\Delta\epsilon/(kT)]}. \qquad (1)$$

Traditionally, the change in energy is calculated by first determining the energy of each site using the Ising

Hamiltonian [9]. However, here we develop a mathematically equivalent description that results in a much more computationally efficient algorithm that we name the bond formation algorithm. In this algorithm, the swapping process is thought of as the breaking of the bonds present in the initial state and the formation of new bonds in the final state. In this framework, the change in energy caused by swapping two sites is the difference between the total energy of the initial bonds and the total energy of the final bonds. To calculate this difference, all that needs to be known is the change in the number of each type of bond between the initial and final states. Since the method implemented by Watkins *et al.* includes only interactions between the first and second nearest neighbors, the total change in energy is calculated

$$\Delta\epsilon = -\Delta N_1 J - \Delta N_2 \frac{J}{\sqrt{2}}, \qquad (2)$$

where J is the interaction energy, ΔN_1 is the change in the number of first-nearest-neighbor bonds, and ΔN_2 is the change in the number of second-nearest-neighbor bonds. A more detailed description of the bond formation algorithm and a comparison to the algorithm based on the traditional energy calculation method is presented in Sec. I of the Supplemental Material [34].

Once the probability of the swapping event is calculated, a random number generator is used to determine whether the sites are swapped or not. To continue the phase separation, another suitable pair of sites is randomly chosen, and the process is repeated. Whether the sites are swapped or not, each iteration is counted, and the evolution of the system is measured by counting the number of Monte Carlo (MC) steps that have occurred. The number of MC steps is defined as the total number of iterations divided by the total number of sites in the lattice [9]. This allows the evolution of the phase separation process to be characterized with a parameter that is independent of the lattice size.

B. Smoothing algorithm

To modify the morphology, we also implement a smoothing algorithm that removes island sites and smooths rough domain interfaces. During smoothing, the lattice is scanned one site at a time, and for each site, a roughness factor is calculated. The roughness factor of a site is calculated by determining the fraction of the 26 total first, second, and third nearest neighbors that are not the same type as the target site. Island sites and sites at rough domain interfaces are surrounded by mostly sites of the opposite type and will have a large roughness factor. To smooth the domains, any site that has a roughness factor above a given threshold is switched to the opposite type. The lattice is continually scanned until all sites are found to have a roughness factor that is below the threshold.

We find that a smoothing threshold of 0.52 performs best by reducing the interfacial area without significantly affecting the domain size. With a 50:50 blend, this smoothing process has an equal probability of smoothing out donor or acceptor sites, and as a result, the blend concentration is not affected. However, if uneven blend ratios are used, this algorithm does slightly reduce the concentration of the minority component. A more detailed analysis of the smoothing threshold is presented in Sec. II of the Supplemental Material [34].

C. Morphology characterization

Once a morphology is generated, to characterize the average size of the domains, the pair-pair correlation method previously described by Lyons *et al.* is used [32]. The pair-pair correlation function is calculated for each donor site and each acceptor site. Then, two averages are calculated, one for all donor sites and one for all acceptor sites. In our pair-pair correlation function algorithm, a resolution of 0.5 nm is implemented. As a result, the distance between sites is rounded to the nearest 0.5 nm. The average domain size is determined by calculating when the correlation function first crosses over the bulk concentration (0.5). To calculate the crossover point, a linear interpolation process is used between the two points on either side of the crossover point.

We note here, that in numerous previous studies, the domain size has been estimated using the relationship, $d = \frac{3V}{A}$, where V is the volume and A is the interfacial area [10,12,13,16,18,20,23,35]. This relationship is only strictly valid when the domains are spherical, and since the Ising model produces highly nonspherical domains, this approximation severely overestimates the domain size. We estimate that these studies have likely overestimated the domain size present in their morphologies by about 75%. A more detailed comparison between these methods in presented in Sec. III of the Supplemental Material [34].

To characterize the shape and connectivity of the domains, the interfacial-area-to-volume ratio and tortuosity is calculated. The interfacial-area-to-volume ratio is calculated by counting the number of cubic site faces between a donor and an acceptor site and then dividing the total count by the total number of sites in the lattice. The tortuosity is defined for an individual site as the length of the shortest available path from the given site through the same domain type to the collecting electrode divided by the length of the corresponding shortest straight path [36]. To calculate this, a three-dimensional, breadth-first search, graph traversal method is used to determine the shortest path from all donor and acceptor sites to their respective collecting electrode. The tortuosity is then calculated for all donor sites at the cathode interface and all acceptor sites at the anode interface to give a data set that is representative of the charge-transport paths through the entire thickness of the film. Since a 50:50 blend is studied here, the donor

and acceptor paths should be statistically equal and are averaged together.

To determine how the morphological changes impact simulated device performance, two simple KMC simulation benchmarks are used. A detailed description of our KMC simulation methods can be found in our previous work [37]. In all KMC simulations, an uncorrelated Gaussian density of states is implemented with a standard deviation of 75 meV, and a temperature of 300 K is used. First, an exciton-quenching efficiency test is performed by generating excitons with a diffusion length of 10 nm and a lifetime of 500 ps [38]. In this test, all excitons that reach the donor-acceptor interface within their lifetime are dissociated, and those that do not reach the interface relax to the ground state. The interaction distance for exciton dissociation is set to 2 nm, in accordance with our previous studies [11]. Exciton-quenching tests are performed for 1000 excitons on nine different energetic disorder configurations for each morphology. Second, the effect of the morphology on charge transport is probed using a simulated thin film time-of-flight (TOF) experiment. In this test, a hole is created at a randomly selected donor site at one surface of the lattice and allowed to undergo standard hopping behavior under an applied electric field. When it reaches the opposite surface, it is removed from the lattice, the transit time is recorded, and the entire process is repeated. With an applied field of 10^7 V/m, TOF simulations are performed for 1000 carriers on nine different energetic disorder configurations for each morphology.

Additional computational details and calculation time benchmarks for the morphology sets are provided in Sec. IV of the Supplemental Material [34]. An open-source software package for supercomputer use and a more simple Web-based morphology generation tool to create morphologies using the methods described here are available online [39,40].

III. RESULTS AND DISCUSSION

A. Effect of the interaction energy

The first and most important behavior to understand is the effect of the interaction energy J on the generated morphologies. Previous studies have often used an interaction energy of $1kT$ [9,10,17,23,25] or have neglected to specify the interaction energy used without discussing the effect of changing the interaction energy. Our preliminary tests indicated that domain growth is much faster when using smaller interaction energies [11]. This finding prompts an investigation as to whether or not a smaller interaction energy can be used to more efficiently generate model morphologies with large domain sizes. To characterize the effect of the interaction energy in greater detail, 24 independent morphologies are generated on a 50 by 50 by 50 lattice for $J = 0.4$, 0.6, 0.8, and 1.0 kT, varying the number of iterations (MC steps) to create domain sizes in

the range of 5 to 10 nm. The domain size, interfacial-area-to-volume ratio, and tortuosity are calculated for each morphology. In the following sections, the data points in the figures indicate the mean of each data set, and the error bars represent one standard deviation.

Figure 1(a) shows how the domain size grows during the simulated phase separation process as a function of the number of MC steps for different magnitudes of interaction energy. In all cases, the domain growth is fast initially and then slows down over time. This slowing of domain growth over time is particularly pronounced at higher interaction energies, as described previously by Binder and Stauffer [41]. However, the number of MC steps required to reach a specific domain size varies dramatically, with about 1.5 orders of magnitude difference between interaction energies of 0.4 and 1.0 kT. Most significantly, this leads to

much longer computational time. This behavior can be explained by considering the effective mobility of the sites in the lattice. When the interaction energy is large, it is much less likely for sites to go through the energetically unfavorable intermediate states that are required for site rearrangement and eventual domain growth.

In addition, the variability of the domain size obtained for a specific number of MC steps increases as the domains grow in size. This general trend is present for all interaction energies tested but appears to be reduced slightly when using a smaller interaction energy. However, with an interaction energy of $0.8kT$, the domain sizes appear highly varied once an average domain size of 9 nm is reached. This suggests that the lower interaction energies produce morphologies with domains that are more uniform in size.

To compare the generated morphologies in more detail, the interfacial-area-to-volume ratio is calculated and is shown in Fig. 1(b). If the domains are shaped differently, the interfacial-area-to-volume ratio should be affected. For example, if the domains tend to be more spherical in shape, the interfacial-area-to-volume ratio will be lower than if the domains tend to be more cylindrical. It is very clear that the morphologies generated with an interaction energy of $0.4kT$ have a much larger interfacial-area-to-volume ratio. To visualize this difference, Fig. 2(a) shows a cross-sectional image of a morphology generated with an interaction energy of $0.4kT$. It is clear that this morphology has quite a few island sites and very rough interfaces that contribute to the large interfacial-area-to-volume ratio. However, it is still unclear if the underlying domain shape is significantly different.

In an attempt to investigate the potential domain shape differences apart from the effects of island sites and rough domain interfaces, the smoothing algorithm described in Sec. II is applied. Figure 2 shows how the smoothing algorithm modifies the morphology when using an interaction energy of $0.4kT$. From these cross-sectional images, it is clear that the smoothing algorithm successfully removes all island sites and smooths rough domain interfaces without significantly changing the size or shape of the domains.

To analyze and characterize the effect of smoothing, all previously generated morphology sets are smoothed, and the domain size, interfacial-area-to-volume ratio, and tortuosity are recalculated. The resulting data shown in Fig. 3(a) indicate that once smoothing is applied, the high interfacial-area-to-volume ratio of morphologies generated with low interaction energies is greatly reduced, approaching the values originally attained with higher interaction energies. The smoothing process has very little impact on the morphologies generated with $J = 1.0kT$. After smoothing, there is much less difference in the interfacial-area-to-volume ratio between the morphologies generated with different interaction energies at any given domain size. However, significant differences do start to arise when the
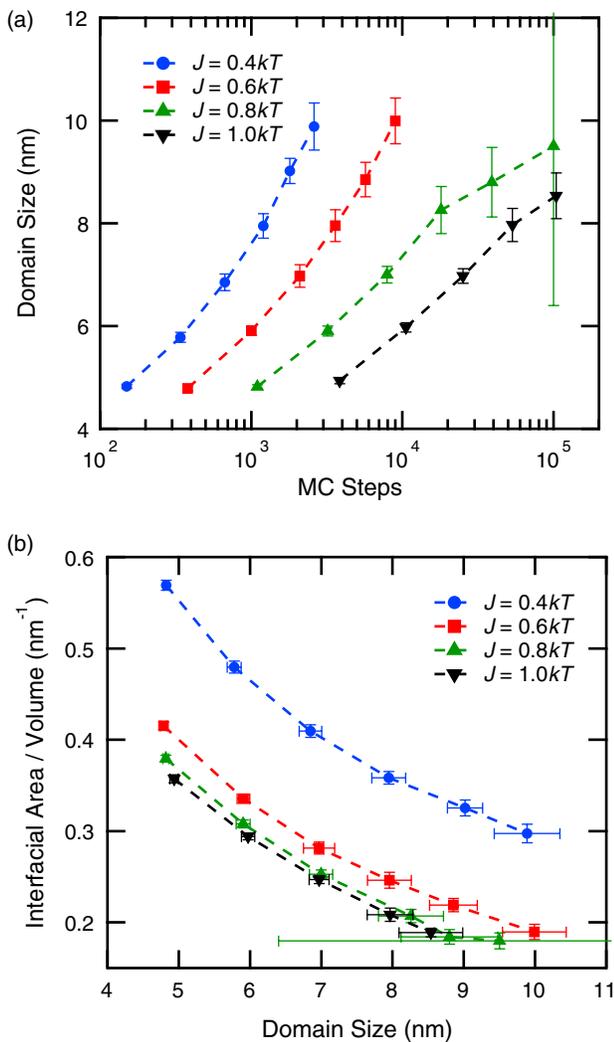
FIG. 1. The effect of the interaction energy J on morphology generation. (a) Growth of domains with increasing MC steps and (b) resulting interfacial-area-to-volume ratio for $J = 0.4kT$ (blue circles), $J = 0.6kT$ (red squares), $J = 0.8kT$ (green triangles), and $J = 1.0kT$ (black inverted triangles).
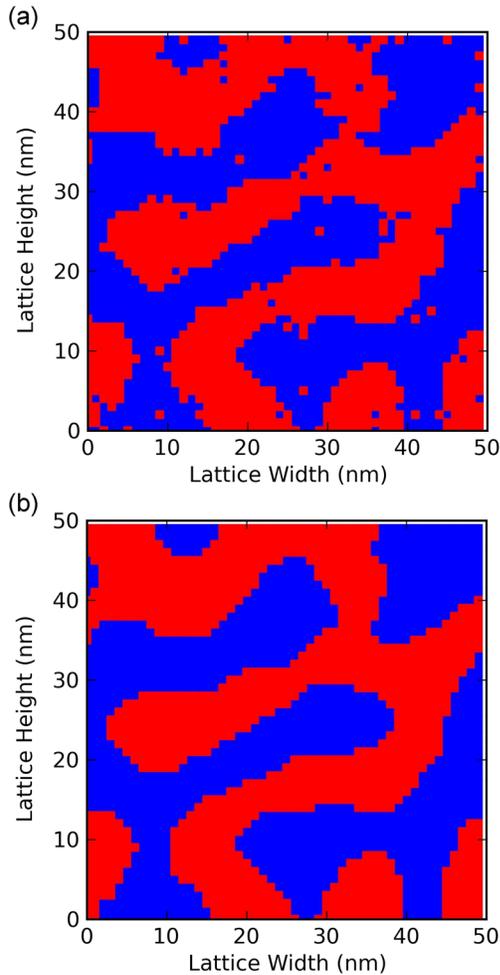
(a)

(b)

FIG. 2. The effect of smoothing on the generated morphologies. Cross-sectional images for $J = 0.4kT$ (7.0-nm domains after 750 MC steps) (a) without smoothing and (b) after smoothing.
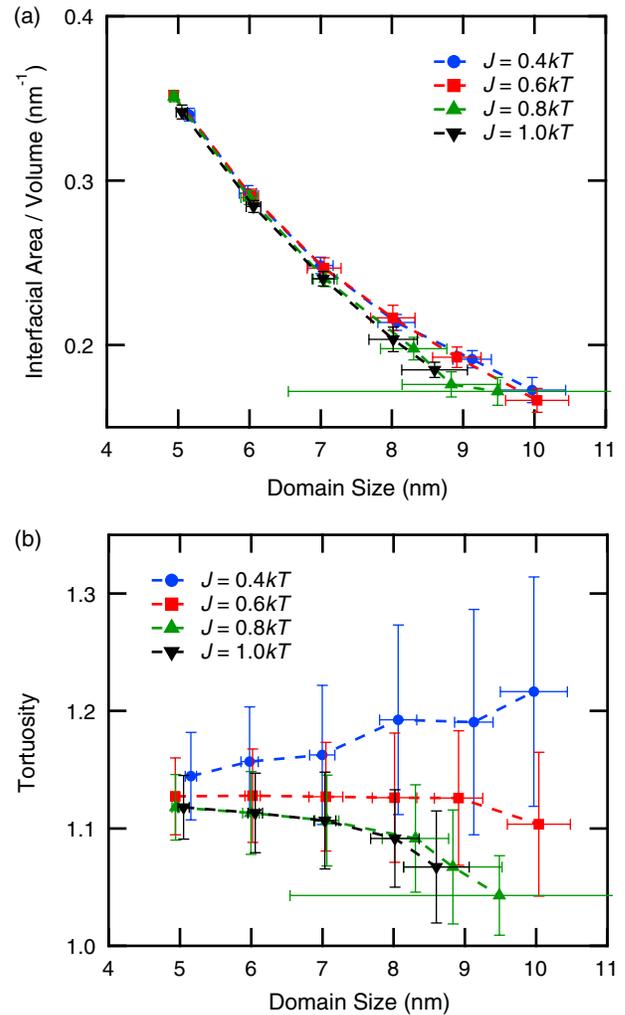
(a)

(b)

FIG. 3. The effect of the interaction energy J on final smoothed morphologies. (a) Interfacial-area-to-volume ratio and (b) tortuosity as a function of domain size after smoothing for $J = 0.4kT$ (blue circles), $J = 0.6kT$ (red squares), $J = 0.8kT$ (green triangles), and $J = 1.0kT$ (black inverted triangles).

domains reach about 8 nm or larger. In this regime, interaction energies of 0.8 and 1.0 kT produce domains with a significantly smaller interfacial-area-to-volume ratio than both 0.6 and 0.4 kT.

Figure 3(b) shows the tortuosity for the smoothed morphologies. Similar to the trend observed with the interfacial-area-to-volume ratio, the tortuosity obtained with each interaction energy is initially very similar but deviates as the domains grow in size. In particular, different trends are observed for each interaction energy. For 0.8 and 1.0 kT, the tortuosity decreases as the domains grow in size. For $0.6kT$, the tortuosity remains almost constant as the domains grow in size, and for $0.4kT$, the tortuosity increases as the domains grow in size. As a result, each interaction energy produces morphologies that have distinct differences and will be expected to produce different KMC simulation results. These potential effects on KMC simulations are tested and discussed in Sec. III C.

From these tests, it appears that using an interaction energy of 0.6kT is best for generating controlled model BHJ morphologies. The domains grow much faster than with higher interaction energies, which reduces computational time, and the domains are also more uniform in size. In addition, the tortuosity of the morphology is fairly constant as the domains grow in size, which allows one to look at the impact of the domain size independent from the tortuosity. We will also show in the next subsection that this constant tortuosity works especially well with the lattice rescaling method used to efficiently create morphologies with larger domain sizes. As a result, the subsequent section focuses only on morphologies generated using an interaction energy of 0.6kT.

B. Simplifications for computational efficiency

Regardless of the magnitude of the interaction energy, there are several methods for significantly reducing the

computation time required to generate a particular model morphology. The first method is to reduce the lateral dimensions of the lattice. Because periodic boundary conditions are used in the plane of the film, the choice of lateral dimensions is somewhat arbitrary. However, the number of sites in the lattice will affect the calculation time per MC step. As a result, it is common to use lateral dimensions that are smaller than the thickness dimension. Reducing two of the dimensions of the lattice can significantly reduce the total number of sites. For all previous tests, a 50 by 50 by 50 site lattice is used, but equivalent morphologies can be generated using a smaller lattice. However, at some point, it is also expected that too small of a lattice may introduce confinement effects that change the domain size and/or domain shape.

To probe this behavior, morphologies are created with domain sizes ranging from 5 to 9 nm on lattices where the length and width (lateral dimensions) are varied but always equal using lattice heights of 50, 75, and 100 nm. Twenty independent morphologies are created for each combination of domain size and lattice size, and after executing the swapping and smoothing algorithms, the final domain size is calculated and recorded for each morphology. Figure 4(a) shows how the domain size is affected by decreasing the lateral lattice dimensions when using a lattice height of 50 nm. When the lateral lattice dimensions are very large, the domain size obtained does not depend on the lattice dimensions as expected. However, as the lateral dimensions become smaller, changes to the domain size eventually start to be observed due to lattice confinement effects. The onset of confinement effects occurs at larger lateral dimensions when creating larger domains, and this same trend was observed for lattice heights of 75 and 100 nm (not pictured).

To characterize this relationship in more detail, for each target domain size and lattice height tested, the lateral dimensions at which the domain size start to be noticeably affected is recorded. We define this transition point as the minimum lateral dimensions. Figure 4(b) shows how the minimum dimensions change as a function of the target domain size for each lattice height tested. We find that the onset of lattice confinement effects is not dependent on the lattice height, and that as a general rule, as long as the lateral dimensions are greater than or equal to 4.5 times the target domain size, lattice confinement does not significantly impact the final morphology. As a result, the calculation time can be reduced by using smaller lateral dimensions, but the lattice size can only be safely reduced down to 4.5 times the desired domain size. We observe a similar limit when using other interaction energies, but do not perform a detailed characterization of these additional cases.

The methods described so far work well for creating relatively small domains, but as the domains continue to grow in size, the rate of domain growth also decreases, as
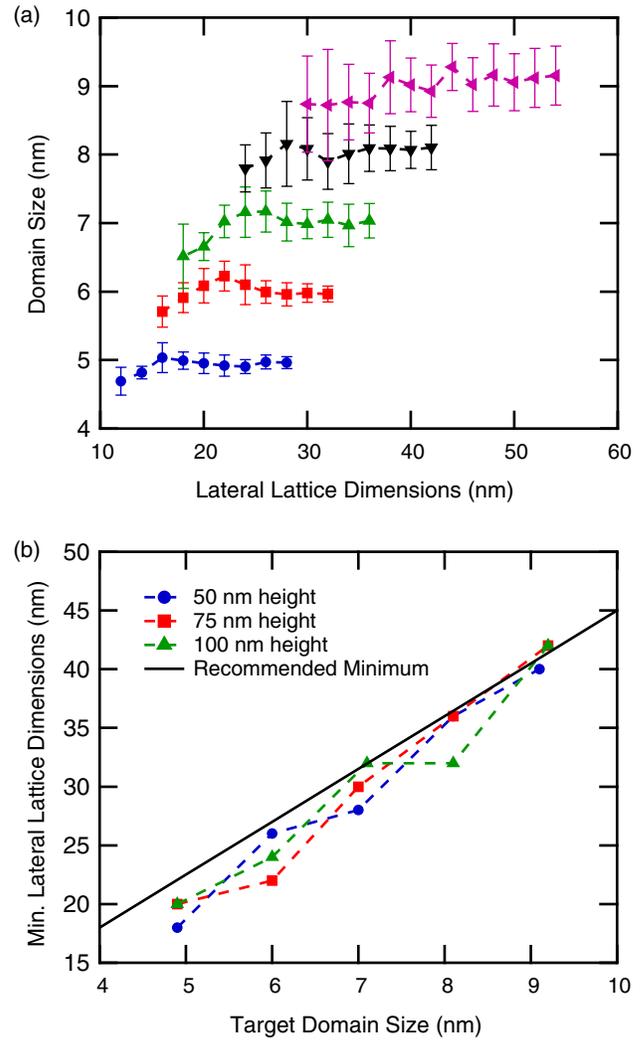

FIG. 4. The effect of the lateral lattice dimensions on morphologies generated with $J = 0.6kT$. (a) The effect of the lateral lattice dimensions on target domain sizes of 5, 6, 7, 8, and 9 nm using a lattice height of 50 nm and (b) the minimum lateral lattice dimensions as a function of target domain size for lattice heights of 50, 75, and 100 nm. The solid black line shows the recommended minimum lateral dimensions for the desired target domain size corresponding to 4.5 times the target domain size.

discussed previously. The final way to reduce the calculation time is to utilize a lattice rescaling method, as used by McNeill *et al.* [3]. This method essentially stretches the lattice equally in all three dimensions, making both the lattice and the domains larger without altering the shape and connectivity of the domains. For creating domains that are larger than 10 nm, which is more typical of many BHJ blend materials, a lattice rescaling method can dramatically reduce the calculation time. As an example, to create a morphology that is representative of a 100-nm film with 16-nm domains, an 80 by 80 by 100 lattice is needed. Even when using a lower interaction energy, without the

rescaling method, this calculation takes several days to create a single morphology on one processor. With the rescaling method, a morphology with 8-nm domains can be created on a much smaller 40 by 40 by 50 lattice and then rescaled by a factor of 2 to obtain the final desired morphology in about 2 h. Additionally, a rescaling factor of 3 can be used on the same initial morphology to create 24-nm domains with only slightly more calculation time.

However, it is also a concern if the rescaling method introduces major changes to the tortuosity. To characterize this, 24 independent morphologies are created with an interaction energy of $0.6kT$ on a lattice with a 100-nm height for a range of domain sizes using both the normal method and the rescaling method. Lateral lattice dimensions are set to 4.5 times the target domain size, and smoothing is applied before and after rescaling. Figure 5 shows that without rescaling, both the interfacial-area-to-volume ratio and the tortuosity begin to decrease for domain sizes over 10 nm. However, with the rescaling method, when initial domain sizes are created in the range of 5 to 9 nm, this does not occur and the tortuosity remains constant. If the rescaling method is used with interaction

energies or domain size ranges where the tortuosity is not constant, the tortuosity will fluctuate when switching between the normal method and each rescaling factor. However, with an interaction energy of $0.6kT$ and initial domain sizes in the range of 5 to 9 nm, the rescaling method can be safely used to efficiently create morphologies with a wide range of domain sizes.

C. Impacts on KMC device simulations

In Sec. III A, we show how the interaction energy can impact the simulated phase separation process, yielding morphologies with different domain shapes characterized by differences in the interfacial-area-to-volume ratio and tortuosity. To determine how these differences will impact device simulations, we generate morphology sets similar to typical OPV devices and perform two benchmark KMC simulations. Interaction energies of 0.4, 0.6, and 0.8 kT are used to create morphology sets with domain sizes of approximately 15, 18, and 21 nm on lattices representing a film thickness of 102 nm. A total of nine morphology sets with 24 morphologies each are generated. Domain smoothing, minimum lateral dimensions, and lattice rescaling are used, as described previously, to quickly generate the final morphologies. Instead of taking several days to generate each morphology, each morphology is created on one processor in only 1–2 h. Additional characterization of these morphology sets is shown in Sec. V of the Supplemental Material [34].

First, the simulated exciton-quenching efficiency is shown in Fig. 6(a), and we find that while the exciton-quenching efficiency is dependent on the domain size as expected, the interaction energy used to generate the morphology has almost no impact. As a result, for studies focused on modeling exciton diffusion and dissociation at the donor-acceptor interface in BHJ devices, the choice of interaction energy is not very significant as long as domain smoothing is applied to remove the island sites that will act as exciton-quenching sites.

However, the situation is significantly different for charge-transport simulations. The transit time distributions resulting from time-of-flight charge-transport simulations are shown in Fig. 6(b). It is clear that for all three domain sizes tested, a smaller interaction energy results in longer transit times and a more dispersive distribution. Looking back at Fig. 3(b), we observe noticeable differences in tortuosity for each of these interaction energies. In general, the tortuosity increases when decreasing the interaction energy. This same trend persists for the larger domain sizes generated here and is shown in Sec. V of the Supplemental Material [34]. This increase in tortuosity appears to significantly slow down the charge transport and increase the dispersion. We expect this effect to be enhanced when simulating thicker films or when the electric field is weaker. As a result, for simulation and modeling studies in which charge transport is an important factor, the choice of
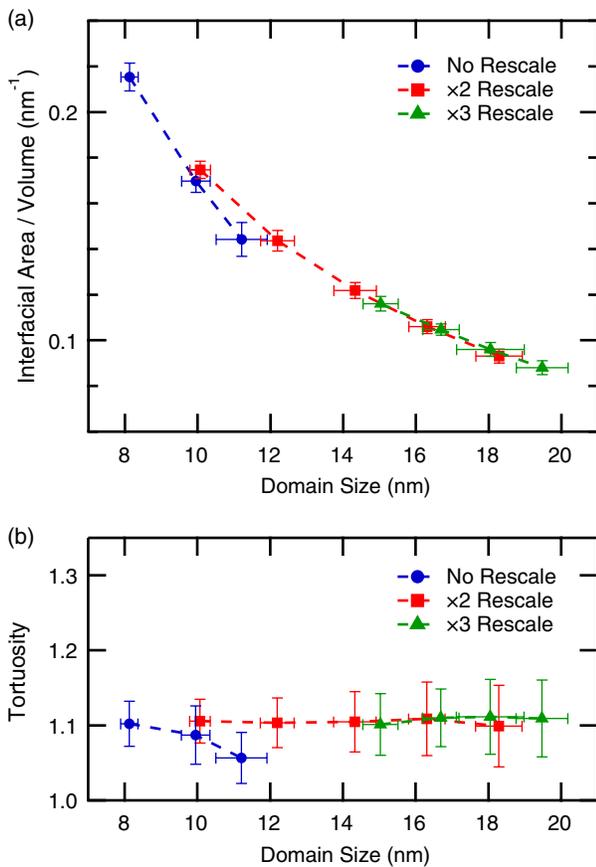

FIG. 5. Effect of lattice rescaling on morphologies. (a) Interfacial-area-to-volume ratio and (b) tortuosity as a function of domain size for $J = 0.6kT$ without rescaling (blue circles), ×2 rescaling (red squares), and ×3 rescaling (green triangles).
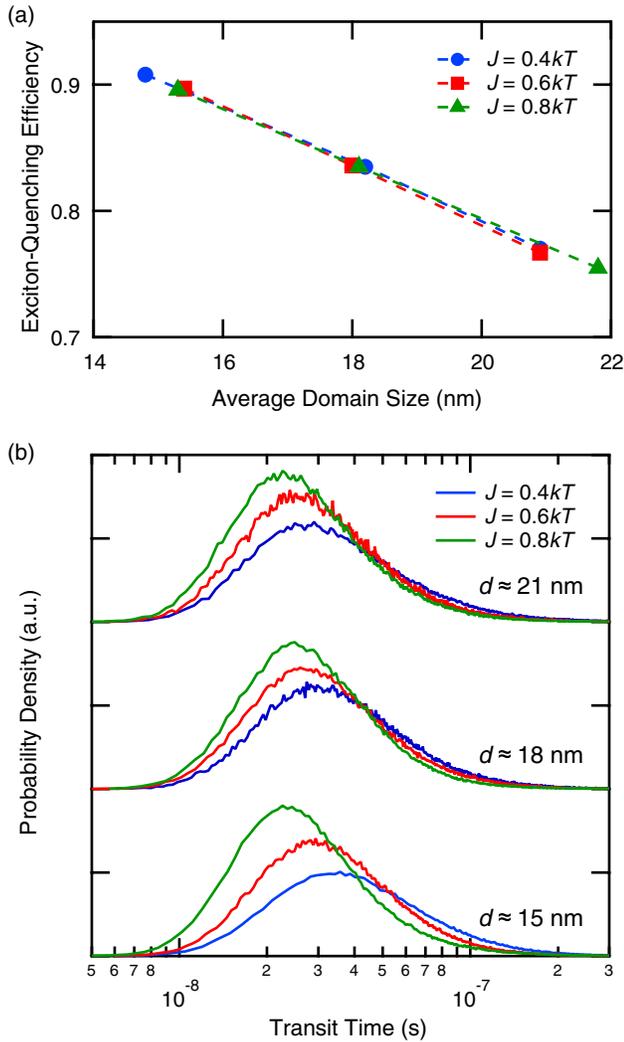
FIG. 6. The effect of the interaction energy J on KMC device simulation performance using model BHJ morphologies. (a) Exciton-quenching efficiency and (b) transit time distributions for $J = 0.4kT$ (blue circles), $J = 0.6kT$ (red squares), and $J = 0.8kT$ (green triangles) and domain sizes of 15, 18, and 21 nm.

interaction energy will have a significant impact. For example, in a full device simulation, slower transport should lead to reduced charge collection efficiency and increased charge recombination.

IV. CONCLUSIONS

Overall, we provide a detailed characterization of the Ising method for generating model BHJ morphologies. We investigate the effect of the interaction energy and demonstrate several methods for reducing the computation time required to generate model morphologies. We first introduce an algorithm called the bond formation algorithm for calculating the site-swapping probabilities, which gives rise to a major increase in calculation speed.

We then demonstrate how a smaller interaction energy of $0.6kT$, when used with a smoothing algorithm, produces pure domains with faster domain growth than previous work and a tortuosity that is almost independent of the domain size. Next, we probe the limits of using small lateral lattice dimensions with periodic boundary conditions to reduce the calculation time. Finally, we characterize the performance of a lattice rescaling method to be used when creating large domains (> 10 nm) and identify the conditions that allow for the creation of a wide range of domain sizes. In total, these developments reduce the morphology generation time by several orders of magnitude.

Combining all methods discussed here, morphologies with domain sizes and thicknesses typical of optimized BHJ OPVs are able to be efficiently generated for KMC simulations. We show how changes in domain size and tortuosity can significantly impact charge transport, which can have a broad impact on charge recombination and ultimately the power-conversion efficiency. With this in mind, it is imperative that future modeling studies are precise and forthcoming regarding the methods used for morphological modeling. In particular, studies should pay close attention to how the domain size is determined, how the tortuosity changes with increasing domain size, and how the lattice size is chosen.

Including detailed morphological features into device models continues to be an important step towards the ability to accurately analyze, simulate, and ultimately predict device performance. The advancements described here have been implemented and published in an open-source software code for supercomputer use [39] and in a user-friendly Web-based software tool [40]. With the methods and morphology generation tools freely available, other researchers can now easily generate model BHJ morphologies in a computationally efficient manner and apply them to novel systematic device modeling efforts. These developments will allow KMC simulations to be readily performed on large sets of morphologies created with a wide range of parameters, leading to increased understanding of the link between morphology and device performance.

[1] L. L. Kazmerski, Best research-cell efficiencies, http://www.nrel.gov/ncpv/images/efficiency_chart.jpg.

[2] H.-Y. Chen, J. Hou, S. Zhang, Y. Liang, G. Yang, Y. Yang, L. Yu, Y. Wu, and G. Li, Polymer solar cells with enhanced open-circuit voltage and efficiency, Nat. Photonics **3**, 649 (2009).

[3] C. R. McNeill, S. Westenhoff, C. Groves, R. H. Friend, and N. C. Greenham, Influence of nanoscale phase separation on the charge generation dynamics and photovoltaic performance of conjugated polymer blends: Balancing charge generation and separation, J. Phys. Chem. C **111**, 19153 (2007).

[4] Y. Zheng and J. Xue, Organic photovoltaic cells based on molecular donor-acceptor heterojunctions, Polym. Rev. **50**, 420 (2010).

[5] P. Peumans, S. Uchida, and S. R. Forrest, Efficient bulk heterojunction photovoltaic cells using small-molecular-weight organic thin films, Nature (London) **425**, 158 (2003).

[6] K. Kawasaki, Diffusion constants near the critical point for time-dependent Ising models. I, Phys. Rev. **145**, 224 (1966); Diffusion constants near the critical point for time-dependent Ising models. II, Phys. Rev. **148**, 375 (1966); Diffusion constants near the critical point for time-dependent Ising models. III. Self-diffusion constant, Phys. Rev. **150**, 285 (1966).

[7] K. Binder, Kinetic Ising model study of phase separation in binary alloys, Z. Phys. **267**, 313 (1974).

[8] A. B. Bortz, M. H. Kalos, J. L. Lebowitz, and M. A. Zendejas, Time evolution of a quenched binary alloy: Computer simulation of a two-dimensional model system, Phys. Rev. B **10**, 535 (1974).

[9] P. K. Watkins, A. B. Walker, and G. L. B. Verschoor, Dynamical Monte Carlo modelling of organic solar cells: The dependence of internal quantum efficiency on morphology, Nano Lett. **5**, 1814 (2005).

[10] K. Feron, X. Zhou, W. J. Belcher, and P. C. Dastoor, Exciton transport in organic semiconductors: Förster resonance energy transfer compared with a simple random walk, J. Appl. Phys. **111**, 044510 (2012).

[11] M. C. Heiber and A. Dhinojwala, Estimating the magnitude of exciton delocalization in regioregular P3HT, J. Phys. Chem. C **117**, 21627 (2013).

[12] C. Groves, R. A. Marsh, and N. C. Greenham, Monte Carlo modeling of geminate recombination in polymer-polymer photovoltaic devices, J. Chem. Phys. **129**, 114903 (2008).

[13] C. Groves, R. G. E. Kimber, and A. B. Walker, Simulation of loss mechanisms in organic solar cells: A description of the mesoscopic Monte Carlo technique and an evaluation of the first reaction method, J. Chem. Phys. **133**, 144110 (2010).

[14] C. Groves, J. C. Blakesley, and N. C. Greenham, Effect of charge trapping on geminate recombination and polymer solar cell performance, Nano Lett. **10**, 1063 (2010).

[15] H. Yan, S. Swaraj, C. Wang, I. Hwang, N. C. Greenham, C. Groves, H. Ade, and C. R. McNeill, Influence of annealing and interfacial roughness on the performance of bilayer donor/acceptor polymer photovoltaic devices, Adv. Funct. Mater. **20**, 4329 (2010).

[16] C. Groves and N. C. Greenham, Bimolecular recombination in polymer electronic devices, Phys. Rev. B **78**, 155205 (2008).

[17] K. Feron, C. J. Fell, L. J. Rozanski, B. B. Gong, N. Nicolaidis, W. J. Belcher, X. Zhou, E. Sesa, B. V. King, and P. C. Dastoor, Towards the development of a virtual organic solar cell: An experimental and dynamic Monte Carlo study of the role of charge blocking layers and active layer thickness, Appl. Phys. Lett. **101**, 193306 (2012).

[18] C. Groves, L. J. A. Koster, and N. C. Greenham, The effect of morphology upon mobility: Implications for bulk heterojunction solar cells with nonuniform blend morphology, J. Appl. Phys. **105**, 094510 (2009).

[19] D. Kipp and V. Ganesan, A kinetic Monte Carlo model with improved charge injection model for the photocurrent characteristics of organic solar cells, J. Appl. Phys. **113**, 234502 (2013).

[20] R. A. Marsh, C. Groves, and N. C. Greenham, A microscopic model for the behavior of nanostructured organic photovoltaic devices, J. Appl. Phys. **101**, 083509 (2007).

[21] B. Lei, Y. Yao, A. Kumar, Y. Yang, and V. Ozolins, Quantifying the relation between the morphology and performance of polymer solar cells using Monte Carlo simulations, J. Appl. Phys. **104**, 024504 (2008).

[22] F. Yang and S. R. Forrest, Photocurrent generation in nanostructured organic solar cells, ACS Nano **2**, 1022 (2008).

[23] L. Meng, Y. Shang, Q. Li, Y. Li, X. Zhan, Z. Shuai, R. G. E. Kimber, and A. B. Walker, Dynamic Monte Carlo simulation for highly efficient polymer blend photovoltaics, J. Phys. Chem. B **114**, 36 (2010).

[24] R. G. E. Kimber, A. B. Walker, G. E. Schröder-Turk, and D. J. Cleaver, Bicontinuous minimal surface nanostructures for polymer blend solar cells, Phys. Chem. Chem. Phys. **12**, 844 (2010).

[25] L. Meng, D. Wang, Q. Li, Y. Yi, and J.-L. Brédas, An improved dynamic Monte Carlo model coupled with Poisson equation to simulate the performance of organic photovoltaic devices, J. Chem. Phys. **134**, 124102 (2011).

[26] R. G. E. Kimber, E. N. Wright, S. E. J. O'Kane, A. B. Walker, and J. C. Blakesley, Mesoscopic kinetic Monte Carlo modeling of organic photovoltaic device characteristics, Phys. Rev. B **86**, 235206 (2012).

[27] L. J. A. Koster, Charge carrier mobility in disordered organic blends for photovoltaics, Phys. Rev. B **81**, 205318 (2010).

[28] J. M. Frost, F. Cheynis, S. M. Tuladhar, and J. Nelson, Influence of polymer-blend morphology on charge transport and photocurrent generation in donor-acceptor polymer blends, Nano Lett. **6**, 1674 (2006).

[29] C. Deibel, T. Strobel, and V. Dyakonov, Origin of the efficient polaron-pair dissociation in polymer-fullerene blends, Phys. Rev. Lett. **103**, 036402 (2009).

[30] I. C. Henderson and N. Clarke, On modelling surface directed spinodal decomposition, Macromol. Theory Simul. **14**, 435 (2005).

[31] B. P. Lyons, N. Clarke, and C. Groves, The quantitative effect of surface wetting layers on the performance of

organic bulk heterojunction photovoltaic devices, J. Phys. Chem. C **115**, 22572 (2011).

[32] B. P. Lyons, N. Clarke, and C. Groves, The relative importance of domain size, domain purity and domain interfaces to the performance of bulk-heterojunction organic photovoltaics, Energy Environ. Sci. **5**, 7657 (2012).

[33] D. P. Olds, P. M. Duxbury, J. W. Kiel, and M. E. Mackay, Percolating bulk heterostructures from neutron reflectometry and small-angle scattering data, Phys. Rev. E **86**, 061803 (2012).

[34] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevApplied.2.014008 for additional details about the calculation methods and supplementary results and analysis.

[35] I. Maqsood, L. D. Cundy, M. Biesecker, J. -H. Kimn, D. Johnson, R. Williams, and V. Bommisetty, Monte Carlo simulation of Förster resonance energy transfer in 3D nanoscale organic bulk heterojunction morphologies, J. Phys. Chem. C **117**, 21086 (2013).

[36] O. Wodo, S. Tirthapura, S. Chaudhary, and B. Ganapathysubramanian, Computational characterization of bulk heterojunction nanomorphology, J. Appl. Phys. **112**, 064316 (2012).

[37] M. C. Heiber and A. Dhinojwala, Dynamic Monte Carlo modeling of exciton dissociation in organic donor-acceptor solar cells, J. Chem. Phys. **137**, 014903 (2012).

[38] O. V. Mikhnenko, M. Kuik, J. Lin, N. van der Kaap, T.-Q. Nguyen, and P. W. M. Blom, Trap-limited exciton diffusion in organic semiconductors, Adv. Mater. **26**, 1912 (2014).

[39] M. C. Heiber, Ising_OPV, https://github.com/MikeHeiber/Ising_OPV.

[40] M. C. Heiber, Bulk heterojunction morphology generator, https://nanohub.org/resources/bhjmorphology.

[41] K. Binder and D. Stauffer, Theory for the slowing down of the relaxation and spinodal decomposition of binary mixtures, Phys. Rev. Lett. **33**, 1006 (1974).