Appendix A: Dimensionless Units

The energy of a bond in our model is not established from a physical system. Therefore, we carefully define it with regard to the thermal energy. For this work, the temperature of the simulation is the variable T which is scaled by

$$T_r \equiv \varepsilon / k \equiv 1 \tag{1}$$

where ε is the energy of a bond and k the Boltzmann constant. In the stochastic simulations we perform the internal energy u per lattice site is compared to the thermal energy β^{-1} using the definition:

$$\beta u = -\left(\frac{\varepsilon}{kT}\right)\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right) = -\left(\frac{T_r}{T}\right)\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right) \tag{2}$$

Reported values of the internal energy are multiplied by k^{-1} as seen by:

$$u = -(\varepsilon)\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right) = -(kT_r)\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right) \tag{3}$$

The denominator 2N reflects that each bond is shared by two sites, where there are N sites on the lattice.

We calculate the residual entropy by extracting estimates heat capacity in two ways. One is to estimate the slope of our data and employ the thermodynamic definition:

$$C_V \equiv \left.\frac{\partial U}{\partial T}\right|_V \approx \frac{\Delta u}{\Delta T / T_r} = \left(\frac{kT_r}{2N}\right)\frac{\Delta\langle \#\text{ of bonds}\rangle}{\Delta T} \tag{4}$$

In the penultimate equality U and T comes from the formal definition; the estimation retains the factor of k^{-1} in our reported quantities as it refers to the definition provided in equation

(3). The alternate definition comes from harnessing the measured susceptibility within our simulation runs:

$$C_V \equiv \frac{\left\langle \left(U - \langle U \rangle\right)^2 \right\rangle}{kT^2} = \frac{\left\langle U^2 \right\rangle - \langle U \rangle^2}{kT^2} \tag{5}$$

We derive this from our simulations by calculating:

$$\begin{aligned} C_V &= \frac{1}{k\left(T/T_r\right)^2}\left(\left\langle\left(-\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right)\right)^2\right\rangle - \left\langle -\left(\frac{\langle \#\text{ of bonds}\rangle}{2\text{N}}\right)\right\rangle^2\right) \\ C_V &= \frac{kT_r^2}{4N^2}\left(\frac{\left\langle \langle \#\text{ of bonds}\rangle^2\right\rangle - \langle\langle \#\text{ of bonds}\rangle\rangle^2}{T^2}\right) = \frac{kT_r^2}{4N^2}\left(\frac{\text{var}\left(\langle \#\text{ of bonds}\rangle\right)}{T^2}\right) \end{aligned} \tag{6}$$

Therefore to return it to the units of N^{-1} so that it scales properly on a per site basis, we multiply by an extra factor of N with the final result:

$$C_V = \frac{kT_r^4}{4N}\left(\frac{\text{var}\left(\langle \#\text{ of bonds}\rangle\right)}{T^2}\right) \tag{7}$$

In Appendix F, we discuss the importance of averaging over only samples which are uncorrelated in time. We take care insure that the observations of the samples used for our analysis are taken at time points much further apart than the relaxation time.

Appendix B: Comparison with Bond Percolation

As we consider the connectivity of our simulation results, it is helpful to relate this to the bond percolation threshold. It has been noted that percolation is a necessary, but not sufficient condition for structural arrest in a reversible gel. Bond percolation is defined as a connected network of bonded sites which spans the lattice. The case in which there is no interaction energy between bonds (the bonds are randomly placed on edges) is equivalent to the high temperature limit in our model. The threshold for bond percolation is defined as the minimum probability that an edge is occupied required for network. This probability p_c is 1/2 on a square lattice and approximately 1/4 on a cubic lattice[145-147]. Converting to the number of bonds per site, the percolation threshold density would be 1 on a square and 3/4 on a cubic lattice. We notice that the dynamic phenomena we observe occur well below the high temperature limits of our model and that the bond densities were larger than the percolation threshold. It is therefore reasonable to assume (although not rigorously proven), that all our simulated samples are percolating networks. This is consistent with the observation that the percolation threshold is above the gelation threshold for reversible bonding systems[40]. The dynamic transitions which we measure are not a direct consequence of percolation.

Appendix C: Landau Theory Calculations in Two Dimension

During the interpretation of our simulation results, it is useful to orient ourselves to the relevant values of the scaled temperature $T/T_r = \beta/\varepsilon$ by considering the mean-field results. It is important to acknowledge the simple Landau theory, which expands the energy in terms of an order parameter has severe and well documented limitations. However, the results provide insight into the order parameters in our model and a starting place for our continued analysis. We note that this is a form of the mean-field approximation, as is clearly seen in the values of the results. We note that in our results, there is a locally preferred structure at lower temperature. This is consistent with an order/disorder transition, and reinforces the less obvious lack a macroscopic symmetry breaking with two of the states dominating at all vertices (preponderance in either the (1,3) ┤, ├ or (2,4)┴, ┬). Indeed, in all cases there is little skewing, even when structure on the order of the length scale of the simulation sample size develops. However it is not clear the number of order parameters needed to describe these results.

In this implementation, we define two order parameters, α and δ, to allow for the possibility of two forms of order/disorder symmetry breaking in the system behavior. Both indicate the same limit of stability limit $T = 0.25$. Thus, in our mean-field results presented in Chapter 1, we assert that the orientations not involved in a locally ordered structure will have degenerate probabilities. Employing this assumption we recover the spinodal temperature, $T = 0.25$, further validating that symmetry breaking occurs only in one direction (either ┤, ├ or ┴, ┬).

We use classical Landau theory[148] to identify when the random state becomes unstable relative to the ordered lamella state for the fully-dense lattice. This describes a liquid to crystal phase transition. It is important to note that at this value of T (or conversely temperature), the system will spontaneously order into a lamella. Other satisfied structures might spontaneously form at temperatures near this transition, but they are neglected in this estimate. Thus, the predictions of this analysis would not be inconsistent with observation of a lamella state in a simulation sample at a higher temperature, or a transition to a different energy minima altogether.

At the heart of Landau theory is a functional form of the free energy as an expansion in powers of an order parameter. We can then use the functional form to identify when the free energy minimum transitions from the zero to a finite value of that parameter. Defining an order parameter to distinguish between a glassy and a polycrystalline state remains a fundamental challenge in this field[149]. We recognize that we cannot identify order parameters to discuss generically this transition because of the degeneracy of the potential energy minimum. However, we can identify order parameters for the specific transition from the disordered state to the lamella state, where all the rows (or columns) are identical in one direction and in the opposite direction they alternate. If we assume periodic boundary conditions, we can reduce the lattice to two sites without loss of generality. In this state two molecule orientations (e.g. ┤, ├) will dominate the lattice and the probability of the other two (e.g. ┴, ┬) will be equal everywhere, but smaller. This introduces the first-order parameter, α. A second order parameter, δ, is introduced when we choose either the left-hand or right-hand site to be in state 1 (alternately the top or bottom site to be in state 2).

We use these order parameters to write the free energy expansion. Both α and δ are even functions so no linear term appears in the free energy expansion. The transition occurs when the second order terms change from positive to negative (or the global free energy minimum changes from zero to nonzero). The value of ε required for this transition is 4 (or the temperature of this transition is T =0.25). .

Using the Hamiltonian, we find the internal energy for our model is:

$$
\begin{aligned}
\beta H = -\left(\frac{\beta}{2}\right)\sum_{i=1}^{N}\sum_{j=1}^{N}\{\ & p_1(i,j)\big[\ p_1(i,j+1)+p_2(i,j+1)+p_3(i,j+1)+p_2(i+1,j)\\
&+p_3(i+1,j)+p_4(i+1,j)+p_1(i,j-1)+p_3(i,j-1)+p_4(i,j-1)\ \big]+\\
& p_2(i,j)\big[\ p_2(i+1,j)+p_3(i+1,j)+p_4(i+1,j)+p_1(i,j-1)+p_3(i,j-1)\\
&+p_4(i,j-1)+p_1(i-1,j)+p_2(i-1,j)+p_4(i-1,j)\ \big]+\\
& p_3(i,j)\big[\ p_1(i,j+1)+p_2(i,j+1)+p_3(i,j+1)+p_1(i,j-1)+p_3(i,j-1)\\
&+p_4(i,j-1)+p_1(i-1,j)+p_2(i-1,j)+p_4(i-1,j)\ \big]+\\
& p_4(i,j)\big[\ p_1(i,j+1)+p_2(i,j+1)+p_3(i,j+1)+p_2(i+1,j)+p_3(i+1,j)\\
&+p_4(i+1,j)+p_1(i-1,j)+p_2(i-1,j)+p_4(i-1,j)\ \big]\ \}\\
-\left(\frac{\beta\eta}{2}\right)\sum_{i=1}^{N}\sum_{j=1}^{N}\ & p_5(i,j)\big[\ p_5(i,j+1)+p_5(i+1,j)+p_5(i,j-1)+p_5(i-1,j)\ \big]
\end{aligned}
\qquad (1)
$$

With the value of the reduced temperature given in chapter 2 . The free energy is therefore given by:

$$
\beta f = \beta u + \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{4}\ p_{k,i,j}\ln p_{k,i,j} \qquad (2)
$$

We substitute the following relations to describe the lamella state in terms of one site:

$$
\begin{aligned}
p_{2,i,j} &= p_{4,i,j}\\
p_{2,i,j} &= p_{2,i,j+1} = p_{2,i+1,j}\\
p_{1,i,j} &= p_{1,i+2,j} = p_{3,i+1,j} = p_{3,i+3,j}\\
p_{1,i,j} &= p_{2,i,j+1}
\end{aligned}
\qquad (3)
$$

and write the remaining terms of the ordered parameters α and δ:

$$\begin{aligned} p_1 &= ¼\ (1\ +\ \alpha)\ +\delta \\ p_2 &= ¼\ (1\ -\ \alpha) \\ p_3 &= ¼\ (1\ +\ \alpha)\ -\delta \\ p_4 &= ¼\ (1\ -\ \alpha) \end{aligned} \tag{4}$$

We then expand the free energy in terms of the ordered parameters. This gives:

$$\beta f = \frac{9\beta J}{8} + \ln 4 + \left(\frac{\beta J}{8} + \frac{1}{2}\right)\alpha^2 + (\beta J + 4)\delta^2 + \ldots$$

Thus the onset of instability with respect to the lamella phase is at $T^{-1} = -4$ and both ordered parameters become unstable at the same time. This result foreshadows the general mean-field solution which does not specify which final state of the lattice is preferred (not restricted to the ordered parameter of a single energy minima- that of the lamella). It also supports the results of the DMF simulations which find that even once a minima is found and one orientation is strongly preferred at vertex $\langle \vec{i} \rangle$.

Appendix D: Mean-Field Calculation in Three Dimensions

A natural extension of our work is to consider the model in other lattices forms, either in two dimensions (e.g. hexagonal) or in three dimensions (e.g. cubic). We also can consider other conformations of the bonding sites or additional bonding sites. For example, a molecule with 4 or 5 bonding sites would still be valence-limiting on a cubic lattice (6 edges at each vertex). The complexity this adds to the mean-field calculations is dramatic if we explicitly consider the full Hamiltonian as we did in Chapter 1. During our work, we found the mean-field diagram was a very useful guide for our understanding of the simulations. Others have recently shown that despite the simplicity of this approximation, it provides a more accurate estimate of simulation results than other more detailed calculations. A more detailed discussion of this can be found in Chapter 2.

A further provocative thought is that the glass transition occurs as lack of percolation of mobile regions around glassy states on the potential energy landscape (PEL). The large dimensionality of the PEL exceeds the upper critical dimension limiting a mean field analysis such the mean-field conditions are met and the approximation when used in this way provides the exact results[142]. We have not formulated our approximation to address a percolation transition in accessible portions of the PEL explicitly; however the connection between the ideas is appealing.

As the Hamiltonian becomes larger, we find we can simplify our perspective by considering there are two types of occupied orientations. The first is one in which is part of a locally ordered domain and thus is energetically stabilized. The second is a non-ordered liquid like phase.

In this Section we will consider the same T-shaped molecule, but now on a cubic lattice. Exploration beyond the current T-shape model in 2-dimensions into 3-dimensions or on alternate lattice structures is motivated by both the expectation of more physically relevant systems but also by the increase in accessibility of the equilibrium gel phase. This phase occupies a larger range of densities on the phase diagram as the valency is lowered[7]. Indeed, access to this phase can be shown to require anisotropic interactions[7]. In colloidal gels, represented by potentials that have directional attractive interactions on the smaller than the length scale of the particle, kinetic arrest can be achieved through non-equilibrium routes during phase separation at lower densities (the high density phase becomes an attractive glass) and also via reversible paths at higher density[94, 95, 133].

On the cubic model there are 12 possible orientations for each molecule. We will refer to the probability that the molecule at the vertex is in a locally ordered domain as p_α. Given the density of molecules ρ, and recalling that the other orientations are degenerate as they are still in a liquid-like phase, the value of these probability of the other occupied states are:

$$p_\beta = \frac{(\rho - p_\alpha)}{11} \tag{1}$$

When in the fully-dense region ρ=1. Else, we have used the random dilution approximation by assuming that the interaction with the solvent is neutral. In addition to calculating the phase diagram, we also demonstrate that the model retains the same degeneracy of ground states which we believe to be fundamental to the nature of this class of materials. First we will consider the fully-dense state. At high temperature each orientation is equally probable so:

$$p_\alpha = (1/12) \tag{2}$$

We can calculate the random potential energy by recalling that there are 6 near-neighbors on a cubic lattice, with 3 possible bonds that the molecule at the vertex we are considering could form. The multiplicative factor or 12 arises from the 12 possible orientations and we divide by 1/2 because each bond is shared by two molecules.

$$\beta u_{random} = -\frac{\beta}{2}\left(p_\alpha \times 6 \times 3 \times 12 \times p_\alpha\right) = \frac{-3}{4}\beta \tag{3}$$

Adding to this the entropy we find

$$\begin{aligned} \beta f_{random} &= \beta u_{random} + s_{random} \\ \beta f_{random} &= -\frac{3}{4}\beta + 12 p_\alpha \log(p_\alpha) \\ \beta f_{random} &= -\frac{3}{4}\beta + \log(12) \end{aligned} \tag{4}$$

Recall that we are using the same reduced units as described in the other portions of this work. At lower temperatures, there will be a fraction of molecules that may adopt a locally oriented structure. This results in

$$p_\alpha > (1/12)\,;\, p_\beta = \frac{\left(1 - p_\alpha\right)}{11} \tag{5}$$

$$\beta u_{oriented} = -\frac{\beta}{2}\left(3p_\alpha^2 + 30 p_\alpha \frac{(1-p_\alpha)}{11} + 183\left(\frac{(1-p_\alpha)}{11}\right)^2\right) \tag{6}$$

It would be quite difficult to tally by hand the 216 possible combinations, but the use of locally ordered states helps considerably. First, we consider the interactions resulting from the molecule of interest being in a locally preferred orientation, p_α. We notice that there must be three bonds between the central molecule and its near-neighbors which are also in the locally preferred orientation by definition resulting in a factor of $3p_\alpha^2$. Next, we consider the central molecule and the possible bonds it can form with near neighbor molecules which

are not in locally preferred orientations. There are 12 orientations at each near neighbor, however, only 6 are oriented toward the center molecule, and one of these configuration must already represent the locally oriented structure. There are three bonds which might form between the central locally oriented structure and these liquid like neighbors. Thus there are 15 possible bonds adding $15 p_\alpha \frac{(\rho - p_\alpha)}{11}$.

Now we consider the liquid like states in the from the center molecule. There are 11 states remaining. As before, along the three vertices from which locally oriented bonds are oriented, there are 5 remaining liquid states that can occupy one of the 3 bonds originally between the locally preferred direction and 5 liquid orientations at the near neighbor sites that meet the geometrical requirements. Thus, a first term of $75\left(\frac{(\rho - p_\alpha)}{11}\right)^2$. There are 6 orientations at the central molecule that could present a bond forming direction along the vertice that did not have a bond in the original locally oriented structure. There are the 6 near-neighbors that could take advantage of this. This provides a second term of $108\left(\frac{(\rho - p_\alpha)}{11}\right)^2$ Collecting these terms we find that

$$\begin{aligned} \beta f_{oriented} &= \beta u_{oriented} + s_{oriented} \\ \beta f_{oriented} &= -\frac{1}{2}\left(3p_\alpha^2 + 30 p_\alpha \frac{(1-p_\alpha)}{11} + 183\left(\frac{(1-p_\alpha)}{11}\right)^2\right) \\ &\quad + p_\alpha \log(p_\alpha) + (1-p_\alpha)\log\left(\frac{(1-p_\alpha)}{11}\right) \end{aligned} \tag{7}$$

Extension into the solvated region involves merely replacing $1 - p_\alpha$ with ρ. We then calculate the chemical potential from the free energy as

$$\beta\mu_{liquid} = \frac{\partial f_{liquid}}{\partial \rho}; \beta\mu_{oriented} = \frac{\partial f_{oriented}}{\partial \rho} \quad (8)$$

This allows us to write the grand potential as $w = f - u$. The grand free potential energy describes the relationship between temperature and density. We consider several different coexistence cases. First, we calculate the low density high density coexistence. Then, as temperature lowers, we consider a low density solution in coexistence with a high density locally ordered liquid. These relationships are shown in Figure D-1.

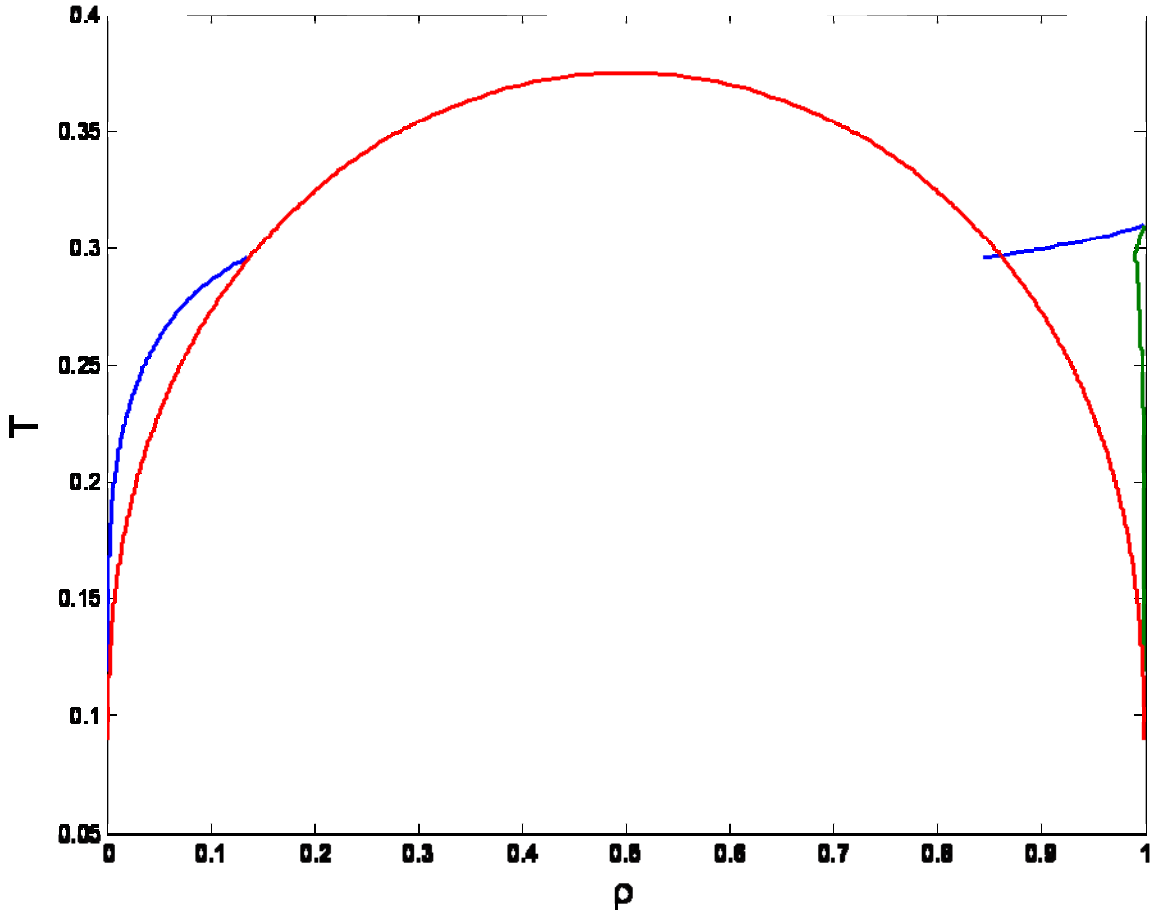

Figure D-1: The three dimensional phase diagram

The red line (——) represents the low density/high density solution coexistence curve. It is superseded by a coexistence between a low density solution represented as a blue line (——) and a high density ordered phase as shown as a green line(——). The fully-dense bimodal is T=0.31 and the spinodal is at T=0.14.

Just as before in the two dimensional case, the quench method whether a continuous quench from a random state, or in a step-wise pattern between previously cooled samples, does not change the average energy of the simulation sample as seen in Figure D-2. This indicates a path independence that further suggests the samples have achieved equilibrium.

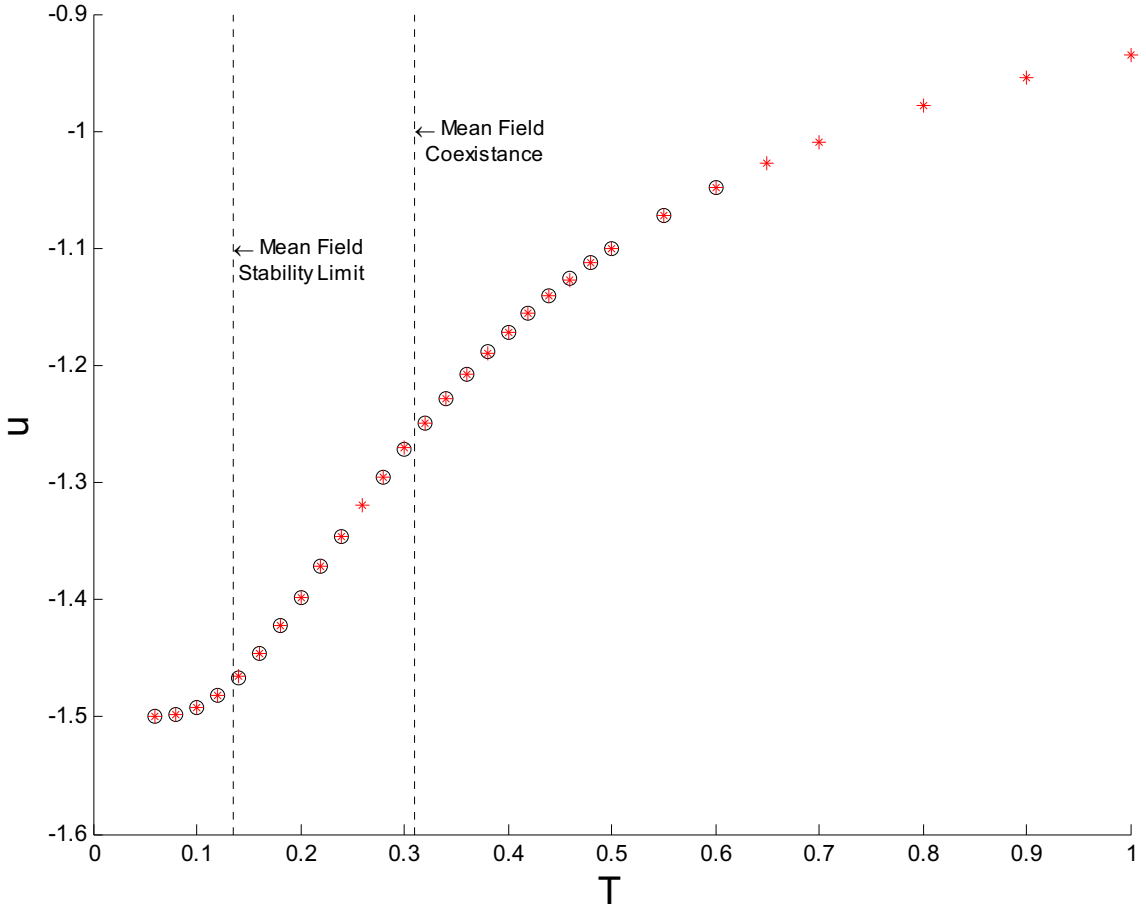

Figure D-2: Simulation results for different quench methods
The Metropolis MC method was used to quench the simulation samples to low temperature based on the choice of path. In the first method, simulation samples were initially in a random configuration and then quenched against the final temperature (o). In the second method, a simulation sample was equilibrated at high temperature and then cooled in a successive series of steps (*). The mean-field bimodal and spinodal are marked as reference temperatures.

We can take an preliminary look at the value of the heat capacity C_v. In our two-dimensional work, we considered three methods of evaluating C_v. In the end, the one we felt was the most representative in the end was fitting the internal energy curve to an analytic fitting function and taking the derivative. We have not shown that in Figure D-3. Instead, we have represented the calculation of C_v using the variance method taking the information directly from the instantaneously quenched simulation samples. As a reference, we also use the quick method of calculating the slope between two measurements with a simple difference method. These results suggest that the C_v maximum is higher than the mean-field spinodal. Whether this is due to the smaller sampling size, not using the preferred method of calculating C_v or some other reason, the impact of this observation on the fragile-to-strong crossover is not evaluated here.

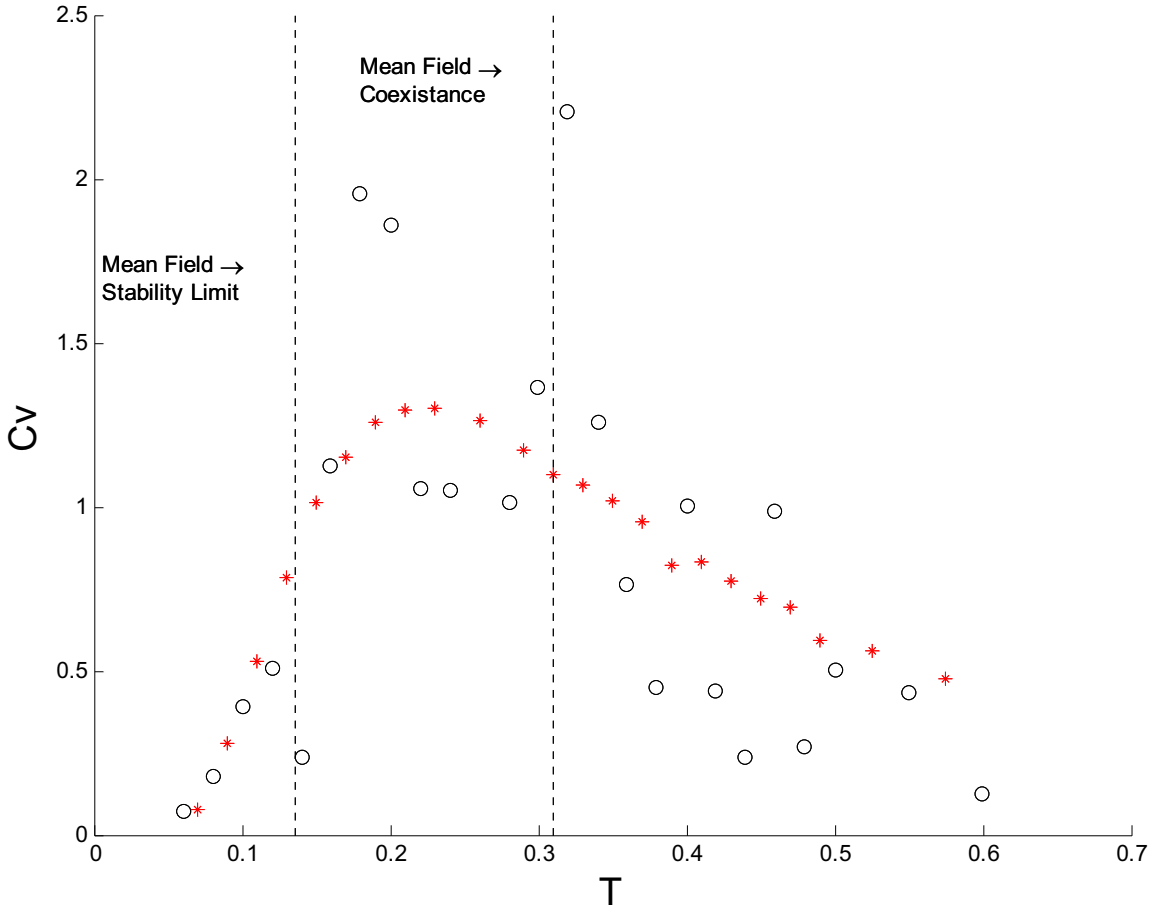

Figure D-3: Heat capacity measurements in the three dimensional model
The heat capacity is calculated by two methods. First, we use the variance from the simulation data itself to calculate the heat capacity (o). As a comparison, we employ a simple difference method to calculate the slope between data from the internal energy (*). The mean-field bimodal and spinodal are once again provided as reference.

One of the key aspects of this model is that it provides a greater than exponential number of ground states connected by very small barriers in energy. Each of the follow groupings is a locally oriented structure that can be tiled across a simulation sample and would result in every bond being satisfied. The boundaries of the tiles all are such that they are interchangeable so that the number of minima is greater than 4^N. There are also other tiling groupings that are not accounted for by combinations of these structures (not shown). Thus, increasing the dimensionality retains this important characteristic.

Figure D-4: Samples of possible locally oriented structures in three dimensions
These four structures could interchangeable be used to tile a large sample. Their boundary conditions match. Each structure is a 4 x 4 x 4 cube.

Appendix E: Stochastic Simulation Techniques

While increases in processor speed and new computational techniques have made simulations a powerful tool in analyzing many systems, they are still plagued by the diversity in size and time scales[110, 150, 151]. Thus, despite the intuitive appeal of molecular dynamics (MD) simulations or off-lattice calculations as incorporating 'more accurate physics', the increased efficiency of on-lattice kMC simulations allow for a much longer time-scale to be evaluated. In the following discussion, we shall parallel the description of MC simulation techniques outlined in Section 2.3.

Master Equation

We can visualize all possible states of the lattice as defining a hyper-dimensional energy surface (albeit discrete) in conformational space. For the purpose of describing this space, the state of the lattice will be identified as a lower case roman character, while the orientation of a specific molecule as a Greek character with a subscript describing its location on the simulation. A specific state m is located at the conformation coordinate $\vec{x}_m$, where $\vec{x}_m \equiv \left\{\sigma_{(1)}, \sigma_{(2)}, \cdots, \sigma_{(N)}\right\}$. The numerical subscripts define the location of the molecule σ on the square lattice. The energy at $\vec{x}_m$ is defined by the Hamiltonian previously described and motion along the surface occurs by the change in orientation of a single molecule. The system is coupled to a thermal bath allowing for transfer between states. This gives us at equilibrium results distributed according to the canonical ensemble; the probability $P_m(eq)$ that the system is located at $\vec{x}_m$, is given by

$$P_m(eq) = \frac{1}{Z}\exp\left(-\beta E_m\right)$$
$$Z = \sum_k e^{-\beta E_k} \tag{1}$$

where k includes all states.

The equilibrium probabilities are not known a priori, instead we use the master equation formalism such that any initial distribution explores conformation space as defined by[115]:

$$\frac{dP_m(t)}{dt} = \sum_k q_{k\to m} P_k(t) - \sum_k q_{m\to k} P_m(t) \tag{2}$$

The transition probabilities, $q_{k\to m}$, specify the choice of motion along the landscape from $\vec{x}_k$ to $\vec{x}_m$. At long-times, steady state should be reached, so the left hand side of (5) must go to zero and the probabilities $P_k(t)$ and $P_m(t)$ adopt their equilibrium value. This results in the condition of detail balance with:

$$q_{k\to m}\exp\left(-\beta E_k\right) = q_{m\to k}\exp\left(-\beta E_m\right) \tag{3}$$

We notice that the choice of transition probabilities is not unique. Our choice of Hamiltonian does not result in an analytically tractable master equation, so we solve for the dynamic and equilibrium properties using MC simulations[46, 116].

Metropolis Monte Carlo

Two different Monte Carlo simulation methods were employed for these studies. First, a simple Metropolis recipe[117] was employed for equilibration of the lattices. In this method, a specific molecule, $\sigma_\alpha(t;\vec{i})$, is chosen at based on a random choice of lattice vertex $\vec{i}$. Then a change in orientation of the molecule from $\alpha \to \gamma$ is proposed at random. We define the number of bonds that $\sigma_\alpha(t;\vec{i})$ forms with its near-neighbors as $b_\alpha(t;\vec{i})$. The transition state between state m and state n is defined as the difference in the energy between the two orientations is compared with a random number.

$$q_{m\to n} = \begin{cases} 1 \text{ if } e^{(1/T)(b_\gamma(t+1;\vec{i})-b_\alpha(t;\vec{i}))} > u \text{ (attempt accepted)} \\ \text{otherwise 0 (attempt rejected)} \end{cases} \tag{4}$$

If the attempt is accepted, the molecule adopts the new orientation, $\sigma_\gamma(t+1;\vec{i})$. Otherwise the original orientation is retained, $\sigma_\alpha(t+1;\vec{i})$. N attempts (where N is the number of vertices on the lattice) is one MCS. The acceptance rate for this method was sufficiently large to cool to very low temperatures relative to the predicted mean-field transition as shown in Appendix C.

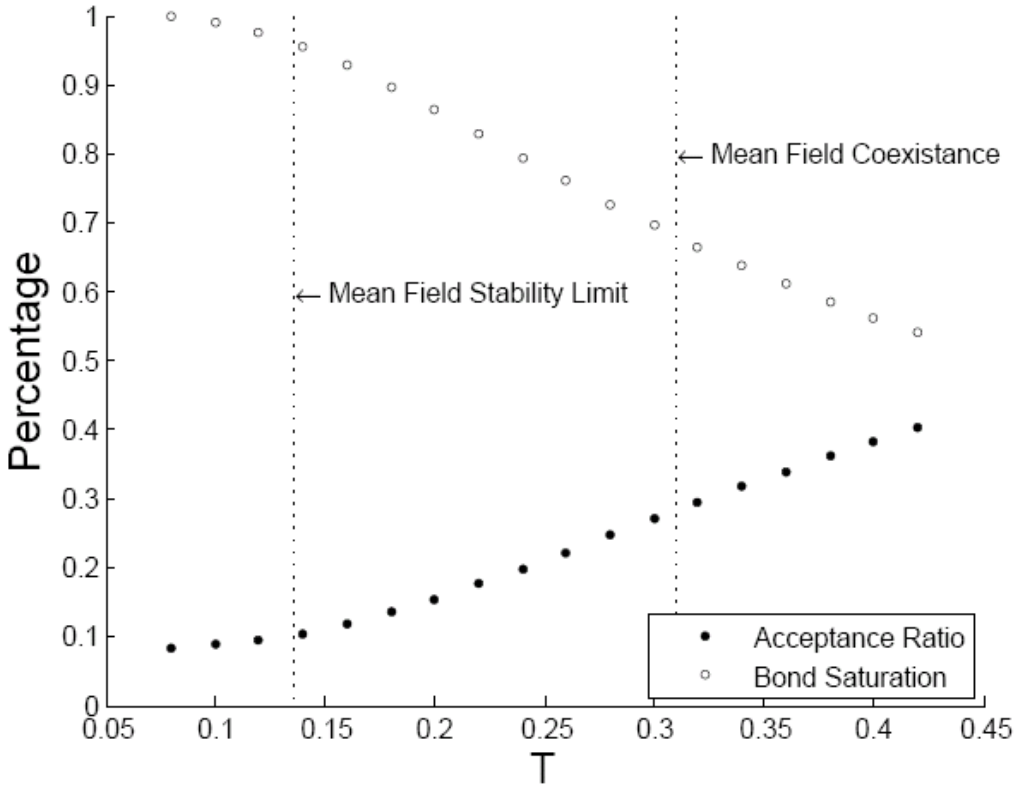

Figure E-1: Metropolis MC acceptance ratio and bond saturation temperature
Despite the 'bond saturation' or the percent of possible bonds which have been formed at low temperature, the 'acceptance ratio' or number of MC attempts which result in a change to the number of attempts made overall, remains relatively large.

Under the Metropolis algorithm, the probability of acceptance is a comparison of the initial and final states only. Although the rate of change of breaking a single bond increases exponentially with decreasing temperature, all transitions to lower or degenerate energy states are accepted. Thus, the method allows for a lattice to explore a large number of configurations quickly from a computational stance. While in our notation we indicate an attempt as an increment in time, $t+1$, the time has no explicit connection to physical time and therefore the dynamics do not have a clear definition. However, we note at equilibrium the distribution of states is expected to be canonical and any method which leads to the correct detail balance is acceptable[118, 119].

Kinetic Monte Carlo:

As we are investigating the properties of the vitrified state, which by definition are time dependent, we need to choose a simulation method which allows us to investigate the dynamics of our system. If the system resides for a considerable length of time within a stable or metastable state and then experiences a very quick (on the time scale of the residence time) transition between such states, the dynamics of the system may be highly dependent on rare transitions. An example of such an event is the difficulty in overcoming high potential energy barriers requiring large activation energies[110, 150]. In such cases, the runtime for a traditional MD simulation is prohibitive; the rare events are separated by an impractical number of MD steps and the sampling necessary to create an accurate thermodynamic ensemble exceeds conventional resources[150]. The wide variety of systems studied that include rare events, and the overall utility of the appropriate kMC simulation results has lead to the development of refined understanding of choice transition states and improved computational techniques [110, 119, 150, 152-161].

In addition to the two conditions described by (5) and (6), extracting dynamic details via kMC simulations requires us to define the of the conformational surface between $\vec{x}_m$ and $\vec{x}_k$ for all states m and k [118, 119]. We then form a dynamical hierarchy of transition rates $r_{m\to k}$ that would reflect the expected behavior of the model system both at equilibrium as well as in non-equilibrium case. The transition probabilities are defined as

$$q_{m\to k} = \frac{r_{m\to k}}{\xi_{\max}} \tag{5}$$

where $\xi_{\max}$ is greater than or equal to the largest transition rate[118]. A wide variety of systems have been studied using this technique[27, 118, 119, 151, 162, 163]. Many examples such of reaction

rates have been directly proposed using the general properties of transition state theory or meet its requirements. Transition-state theory proposes that the transition rate $r_{m\to k}$ is dependent on the details of reaching the saddle point $x^*_{m\to k}$ along the path in between $\vec{x}_m$ and $\vec{x}_n$ [115, 163]. It has been noted that the specific form of the transition probabilities influences the non-equilibrium simulation time[156], even leading to incorrect growth exponents[119]. Kang and Weinberg point out that in most cases the energetic barrier from the initial state, $\vec{x}_m$, to the transition state $x^*_{m\to k}$, describes the microscopic transition process[119].

In our model, we neglect any microscopic energetic details of a simulation sample in state m favoring instead a common vibration frequency, ω. The transition probabilities $q_{m\to n}$ reflect of the potential energy difference between $\vec{x}_m$ and $\vec{x}^*_{m\to n}$

$$q_{m\to n} = \omega \exp(-\beta(E^* - E_m)) \tag{6}$$

where ω^{-1} becomes the fundamental time scale. Conveniently, because each state is linked via the rotation of a single molecule σ_m, we define $E^* - E_m$ as the number of bonds that specific molecule had at state m, which are broken to move to the saddle point. With this choice of the transition probability, there is no impact of the final orientation of that molecule on the transition probability and meets the requirements of detail balance.

We have several large separations between time scales on our lattice. The fastest processes are not activated or cooperative events. The exploration of the local minima, the fastest process, has already been reduced to a constant. We accomplished this by incorporating all β-relaxations (exploration of the local minima) into a single thermal variable, ω, which we assume to be temperature independent over the range of our simulations and the same for all

wells. This establishes a reference time for our system. The transitions between basins do have a significant dependence on the local topology of the PEL. Even with our simple set of transition rates, at low temperatures the difference in breaking a single bond as opposed to three bonds is very significant. We can see this clearly from the plateaus in Figure 2.5.2. Additionally, in the case of a molecule has no bonds with its neighbors there is also fast event taking place on a time scale distinct from the kinetically activated rotations. This is discussed below and illustrated in Figure E-2 and Figure E-3.

By definition, any kinetic MC simulation uses the current transition rate dependent relationship:

$$\Delta t = \frac{-\log(u)}{\sum_m n_m q_m} \tag{7}$$

where u is a random number, q_i is the transition rate for each possible process, in our case given by equation (6). The number of states which can change due to that process is n_m, for our model this is $\sum_{\vec{i}} b_\alpha(t;\vec{i})$. It is important to remember that while at most only one transition occurs between steps, the 'clock' of elapsed time is running for all molecules in parallel, so we need to include all of them in the time increment.

Because we are investigating the low temperature properties of these simulations, using traditional importance methods of proposing a new state[46] and then accepting or rejecting that move does not improve the simulation time significantly because of the large number of rejected moves. Instead, a residence time or 'n-fold'[118, 120, 121] technique in which an event is selected with the appropriate frequency based on the transition rate and every event is accepted improves the time considerably. The transition probabilities and determination of elapsed real time is the same as any other kMC technique.

This technique is requires that the set of available configuration states and their transition properties are know a priori, [27, 151, 164]. The most computationally intensive portion of the method is creating the list of possible transitions with their relevant rates. We are favored in our case by the limited range of our potential, which makes updating this list reasonable.

To enact this process for our, we evaluate the bond energy at every vertex $b_\alpha\left(t;\vec{i}\right)$ and then grouped the same energies and weight them by the appropriate transition probability as given in (6). Therefore we have the intervals:

$$\left[0,\frac{p_0(t)}{p_{tot}(t)}\right),\left[\frac{p_0(t)}{p_{tot}(t)},\frac{\sum_{m=0}^{1}p_m(t)}{p_{tot}(t)}\right),\left[\frac{\sum_{m=0}^{1}p_m(t)}{p_{tot}(t)},\frac{\sum_{m=0}^{2}p_m(t)}{p_{tot}(t)}\right),\left[\frac{\sum_{m=0}^{2}p_m(t)}{p_{tot}(t)},\frac{\sum_{m=0}^{3}p_m(t)}{p_{tot}(t)}\right] \quad (8)$$

where

$$p_{tot}(t)=p_0(t)+p_1(t)+p_2(t)+p_3(t)$$

$$p_{tot}(t)=\omega n_o(t)+n_1(t)\omega e^{-1/T}+n_2(t)\omega e^{-2/T}+n_2\omega(t)e^{-3/T}$$

We then generate a random number u and find the corresponding interval. A second random number is chosen to decide which specific monomer will change. The rotation change is proposed by a third random number. Finally the physical time is updated using equation (7). This completes one step. The values of n_m are updated for the new time step and the cycle repeats.

An ‘instantaneous’ quench rate, in which we initialize a simulation sample in a random configuration and run the kMC simulation at the final temperature shows an extreme example of this behavior. This choice was used in both the results presented in Sections 2.5.2 and 0. The two figures below are representations of the initial portion of the these simulations to demonstrate the effect of the various rates on time progression. One advantage of the residence time kMC technique is that the time between steps changes with the configuration of the lattice. When there are many slow processes, the time step is small.

However if these processes are reduced, the time step becomes larger in a dynamic fashion. In Figure E-2, the 'physical time' and is shown after 1000 MCS at a variety of temperatures. Notice that there is a non-monotonic relationship between the progressed time and temperature.

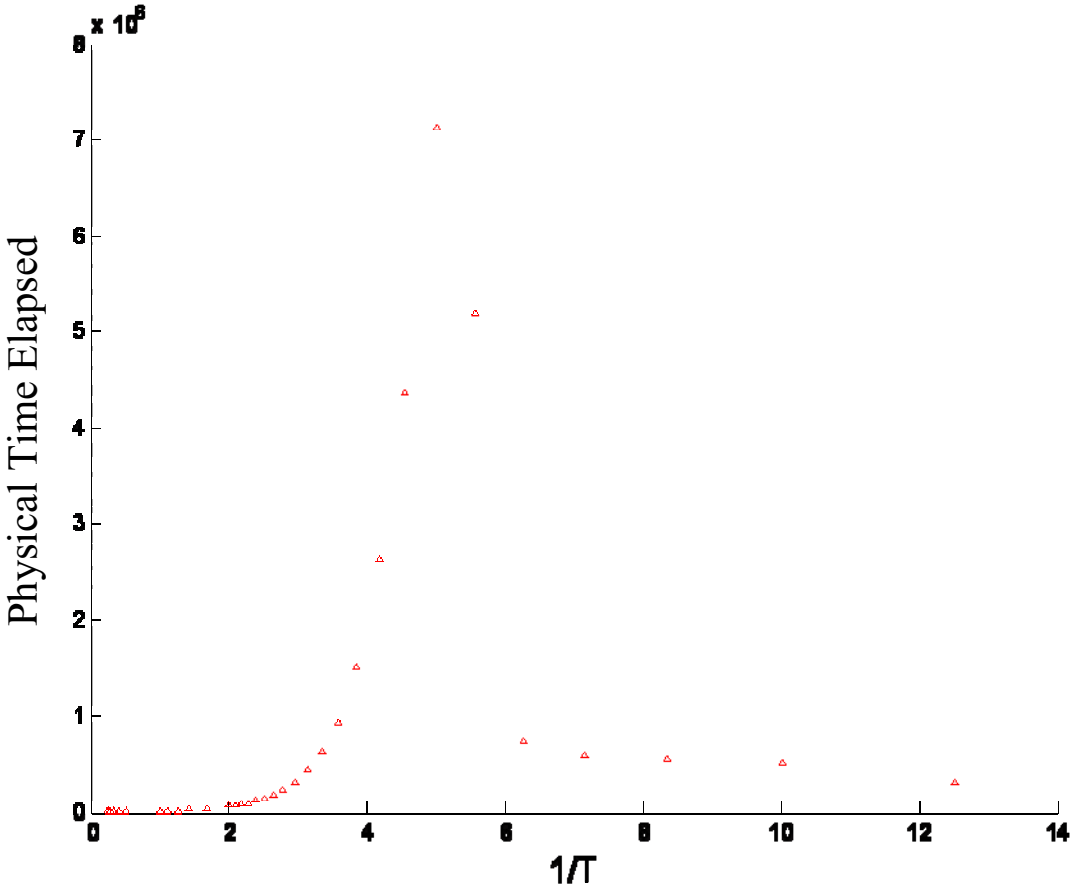

Figure E-2: Physical time at 1000MCS for a variety of temperatures
Despite the same number of successful attempts, the amount of physical time that has elapsed is dramatically different at different temperatures.

For the same 'snap shot' in time above, we report the number of monomers we report the fastest process, which would result in the smallest progression in physical time for each MCS.

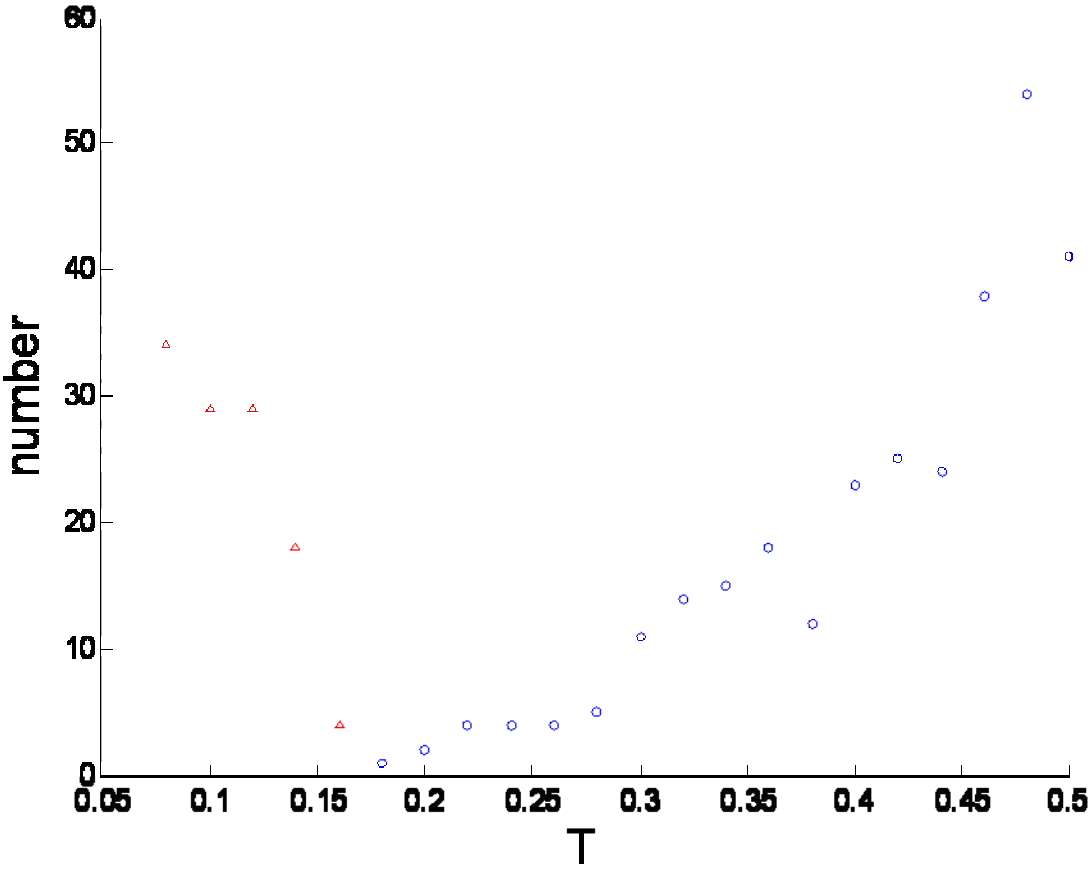

Figure E-3: The number of non-bonded molecules at 1000 MCS as a function of T
Notice that the smallest number of non-bonded molecules after the same number of MC attempts which were successful is the smallest where the elapsed time in Figure E-2 is the largest. This indicates that this fast process was selected with greater frequency. The change in color of symbols indicates whether the simulation was close to its final temperature (o) or still had a long way to progress before it would relax fully (Δ).

The elapsed physical time in each simulation is a dynamic quantity, responding to the distribution of molecules in different bonding states. Thus the accepted simulation step drives will be smaller when there are more fast processes, but then become larger if the number of these processes decreases and the accepted steps are related to the slower, activated events. This is one of the advantages of this technique. There also may be a true dynamic consequence to this pattern of non-bonded states, a relationship that should be explored further in subsequent work.

Appendix F: Calculation of the Heat Capacity from Monte Carlo Simulations

Employing a kinetic Monte Carlo recipe, we can define the relaxation time of a simulation sample by noting that there is a difference of the expectation value of the error when subsequent observations are correlated in time. When deriving the correlation function, we need to exercise great caution in specifying what assumptions are being made. We start our analysis of the correlation function by defining an observable A_t which may be a continuous value such as the energy, or a discrete value such as the state of the lattice sample at time *t*. Properly, an ensemble average would imply averaging over configurational space or time, which are equivalent in an equilibrium system. However we are going to be generous with our definition of ensemble average to include systems which are metastable, that is, are not strictly at equilibrium, but are arrested in a subset of configurational space.

As an MC simulation reflect movement over the PEL, subsequent observations are highly correlated. This provides a difficulty in using the variance, σ, to measure the heat capacity of a simulation because the calculation of expectation value of the variance, $\overline{\delta A^2}$, deviates from σ non-trivially. As we note in our work, the C_v calculation from the derivative of an arbitrarily fit to the energy follows the same trend as the C_v calculated with the measured variance, a relationship shown in Figure F-1. However we also note that as more observations are made, the two values become closer.

How closely does the expectation value of the variance match the actual value as a function of the number of observations? We shall answer that question in two parts. Initially we will ignore the correlated values and see how the accuracy of expectation value of the variance to

the relates to the number of measurements during simple sampling. Then we shall address the more complex case of how the variance is related to the correlation (or relaxation) time.

Formally, we start with the assumption that we are measuring independent observations of a quantity A which obeys a Gaussian distribution with mean $\langle A \rangle$ and variance, $\sigma = \langle A^2 \rangle - \langle A \rangle^2$. If we make a set of n observations of the distribution, $\{A_i\}$, the unbiased estimator of the mean is:

$$\overline{A} = \frac{1}{n}\sum_{i=1}^{n} A_i \tag{1}$$

as we have no other information about the sampling or the distribution. Similarly, the deviations from the observations quantity, $\delta A_i = A_i - \overline{A}$, allow us to find the expectation value of the variance.

$$\begin{aligned} \overline{(\delta A)^2} &= \frac{1}{n}\sum_{i=1}^{n}\left(A_i - \overline{A}\right)^2 \\ &= \frac{1}{n}\sum_{i=1}^{n}\left(A_i^2 - 2A_i\overline{A} + \overline{A}^2\right) \\ &= \overline{A^2} - 2\overline{A}\left(\overline{A}\right) + \overline{A}^2 \\ &= \overline{A^2} - \overline{A}^2 \end{aligned} \tag{2}$$

If we assume that there is no correlation between measurements, the order of summation of the set may be changed (i.e. indices i and j may be interchanged if this violates no other rule of summation). However we do have the restriction that the observations are not equally weighted, although we do not have further information about the value of the weights. Thus this correlation, that the term $A_i A_i$ has more information than $A_i A_j$, which provides the first correction to the expectation value.

$$\begin{aligned}\left\langle \overline{\delta A^2} \right\rangle &= \left\langle \frac{1}{n}\sum_{i=1}^{n} A_i^2 - \left(\frac{1}{n}\sum_{i=1}^{n} A_i\right)\left(\frac{1}{n}\sum_{j=1}^{n} A_j\right) \right\rangle \\ &= \left\langle \frac{1}{n}\sum_{i=1}^{n} A_i^2 - \frac{1}{n^2}\sum_{i=1}^{n} A_i A_i - \left(\frac{1}{n}\sum_{i=1}^{n} A_i\right)\left(\frac{1}{n}\sum_{i\neq j=1}^{n-1} A_j\right) \right\rangle \end{aligned} \tag{3}$$

The ensemble average can be moved inside the summations. The individual values $A_i A_j$ and A_i^2 are each replaced by the values $\langle A \rangle^2$ and $\langle A^2 \rangle$.

$$\begin{aligned}\left\langle \overline{\delta A^2} \right\rangle &= \left\langle \langle A^2 \rangle \frac{1}{n}\sum_{i=1}^{n} 1 - \langle A^2 \rangle \frac{1}{n^2}\sum_{i=1}^{n} 1 - \langle A \rangle^2 \left(\frac{1}{n^2}\sum_{i=1}^{n}\sum_{i\neq j=1}^{n-1} 1\right) \right\rangle \\ &= \left\langle \langle A^2 \rangle - \frac{1}{n}\langle A^2 \rangle - \left(\frac{n(n-1)}{n^2}\right)\langle A \rangle^2 \right\rangle \\ &= \left(\langle A^2 \rangle - \langle A \rangle^2\right)\left(1 - \frac{1}{n}\right) \end{aligned} \tag{4}$$

From the expression (4) we can conclude that improving our estimate of the susceptibility goes as $1/n$ with uncorrelated observations. This shows one computational limit of these methods.

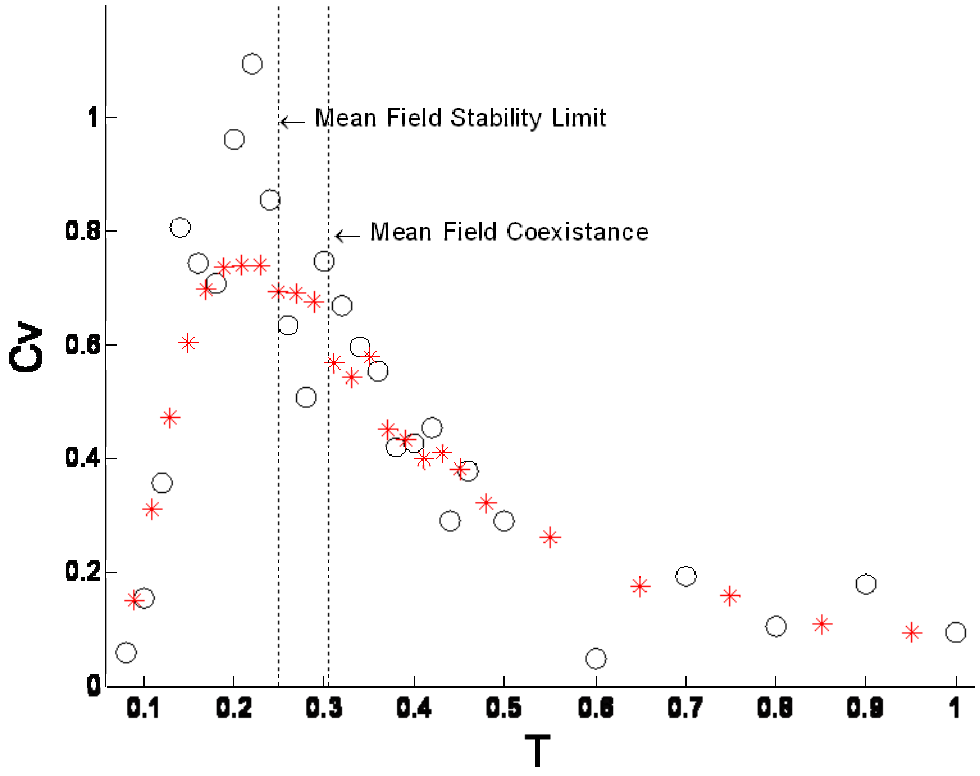

Figure F-1: Specific heat per site calculated by two different methods
The first method was using the slope of u found by simple difference method (*) and using the susceptibility calculated from the variance in the measurements (O)

We can now tighten again our assessment by removing the assumption that subsequent observation are uncorrelated. Returning to equation our original definitions, we remember that the set order is now important, and we cannot interchange indices in sums because they carry information about the order of observation.

$$\begin{aligned}
\overline{(\delta A)^2} &= \left\langle \left[\frac{1}{n}\sum_{i=1}^{n}\left(A_i - \langle A\rangle\right)\right]^2 \right\rangle \\
&= \left\langle \frac{1}{n^2}\sum_{i=1}^{n}\left(A_i - \langle A\rangle\right)\sum_{k=1}^{n}\left(A_k - \langle A\rangle\right) \right\rangle \\
&= \left\langle \frac{1}{n^2}\sum_{i=1}^{n}\sum_{k=1}^{n}\left(A_iA_k - A_i\langle A\rangle - A_k\langle A\rangle + \langle A\rangle^2\right) \right\rangle \\
&= \left\langle \frac{1}{n^2}\sum_{i=k}^{n}\left(A_i^2 - 2A_i\langle A\rangle + \langle A\rangle^2\right) + \frac{2}{n^2}\sum_{i=1}^{n}\sum_{k=i+1}^{n}\left(A_iA_k - \left(\frac{A_i + A_k}{2}\right)\langle A\rangle\right) \right\rangle \\
&= \frac{1}{n}\left[\langle A^2\rangle - \langle A\rangle^2\right] + \frac{2}{n}\sum_{i=1}^{n}\sum_{k=i+1}^{n}\left(\langle A_iA_k\rangle - \langle A\rangle^2\right)
\end{aligned} \tag{5}$$

Next we notice that there is no specific event that breaks the time symmetry between two observations. Thus only the magnitude of $\delta t = |t_k - t_i|$ is important. Therefore, $\langle A_iA_k\rangle = \langle A_0A_{k-i}\rangle$. Therefore, if we define $t' \equiv k - i$ we can remove one of the summations.

$$\begin{aligned}
\left\langle \overline{(\delta A)^2} \right\rangle &= \frac{1}{n}\left[\langle A^2\rangle - \langle A\rangle^2\right] + \frac{2}{n^2}\sum_{i=1}^{n}\left[\sum_{k=i+t}^{n}\left(\langle A_iA_k\rangle - \langle A\rangle^2\right) - \sum_{k=i+1}^{t}\left(\langle A_iA_k\rangle - \langle A\rangle^2\right)\right] \\
&= \frac{1}{n}\left[\langle A^2\rangle - \langle A\rangle^2\right] + \frac{2}{n}\sum_{t=1}^{n}\left(1 - \frac{t}{n}\right)\left(\langle A_0A_t\rangle - \langle A\rangle^2\right)
\end{aligned} \tag{6}$$

This already has the form of a response function embedded in it. If we proceed carefully through the integration we find that the integrated response function can be directly related to the expectation value of the response as[46]:

$$\left\langle \overline{(\delta A)^2} \right\rangle = \frac{\sigma^2}{n}\left(1 + 2\frac{\tau_A}{\delta t}\right) \tag{7}$$

Appendix G: Inherent Structures on the Energy Landscape

While there is compelling evidence for the existence of spatial heterogeneity, it is not known if it is a universal characteristic of glassy systems or how it originates. Sastry et al. [165] ran molecular dynamics (MD) simulations of a binary Lennard-Jones mixture. They found that the exploration of the potential energy landscape by the system was distinctly different in the high temperature (above T_g) and low temperature (below T_g) region. In the high temperature region, the average energy was insensitive to temperature; however, the distribution in energy of the inherent structures of the explored basins was high. At low temperature the average energy again changed little with temperature, however the system explored only basins with low energy inherent structures. The transition between these two plateaus in the average energy of the inherent structures was dramatic.[165]

We develop a Metropolis based scheme for exploring the configuration space of the fully-dense lattice. In these simulations we first equilibrate the lattice at a particular temperature. Then we can explore the energy landscape using zero temperature quench[166] and evaluate relaxation time with different correlations. Results indicate that the simulations have glassy dynamics, even using the Metropolis recipe.

We perform a zero temperature quench to identify the inherent structures. During the zero temperature quench, no annealing is possible; the system can never gain potential energy[75]. In our model there are numerous (and often broad) saddles in our potential energy surface. Thus not all basins are uniquely defined. We modified an algorithm used by Glotzer et al. for these non-unique minimizations[166].

After we allow the lattice to equilibrate, we have the lattice continue to explore configuration space using the algorithm above. Every 1000 MCS we identify the inherent

structure. First we perform one MCS accepting all states with lower potential energy. If a suggested change in molecule orientation has the same energy it is not accepted. Then we perform a second MCS using the heat bath method[46] suggested by Glotzer et al. The heat bath method accepts all moves with lower potential energy, ½ of the moves with equal potential energy and none of the moves with higher potential energy. These two steps are repeated until the energy no longer decreases between steps or a minimum of 100 times.

Preliminary data shows that the difference in the potential energy of the inherent structures with temperature is very small, on the order of a tenth of a percent, see

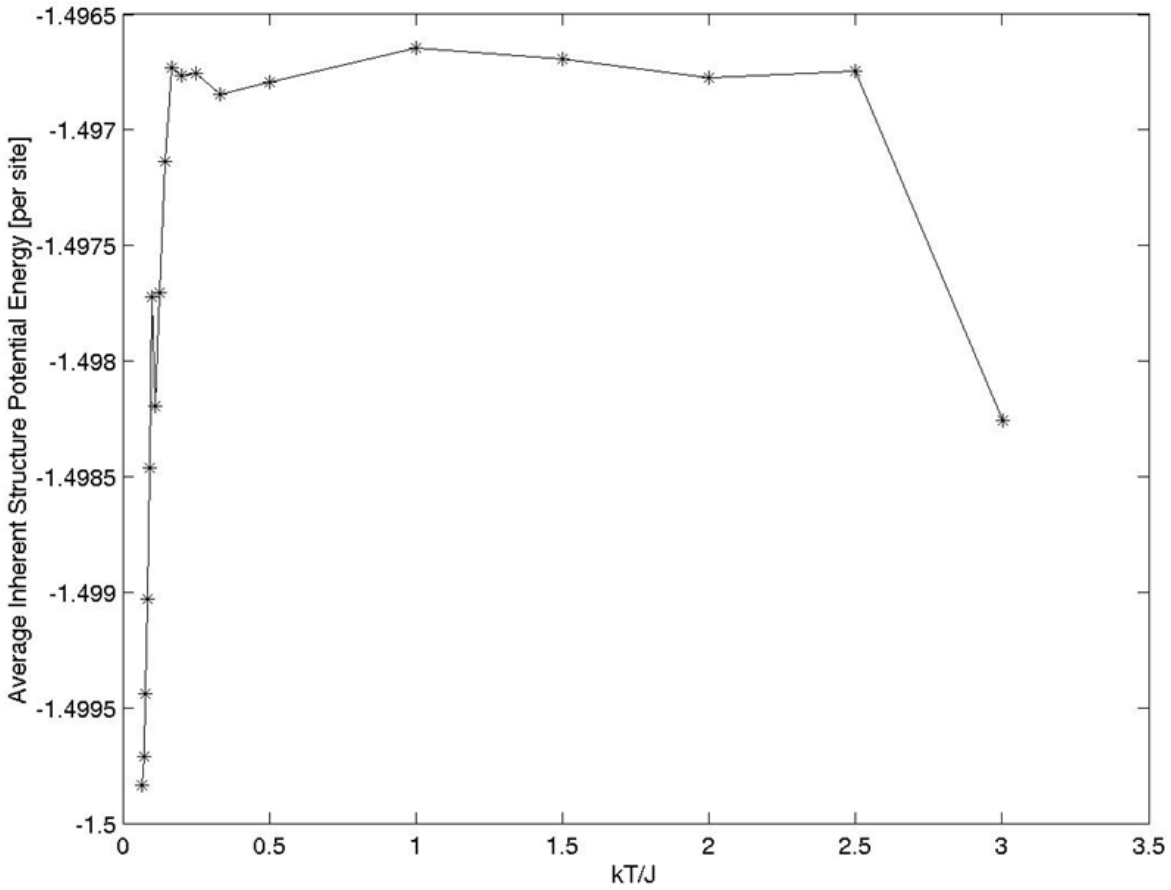

Figure **G-1**. This is smaller than those reported elsewhere. We do not yet have sufficient data to resolve a trend above the error bars. But it is very unlikely that with more data we could see the types of trends shown with both a high and low temperature plateau[165, 166]. This

suggests that our data is more consistent with the change in behavior appearing at the FSC as has been documented in silica[72].

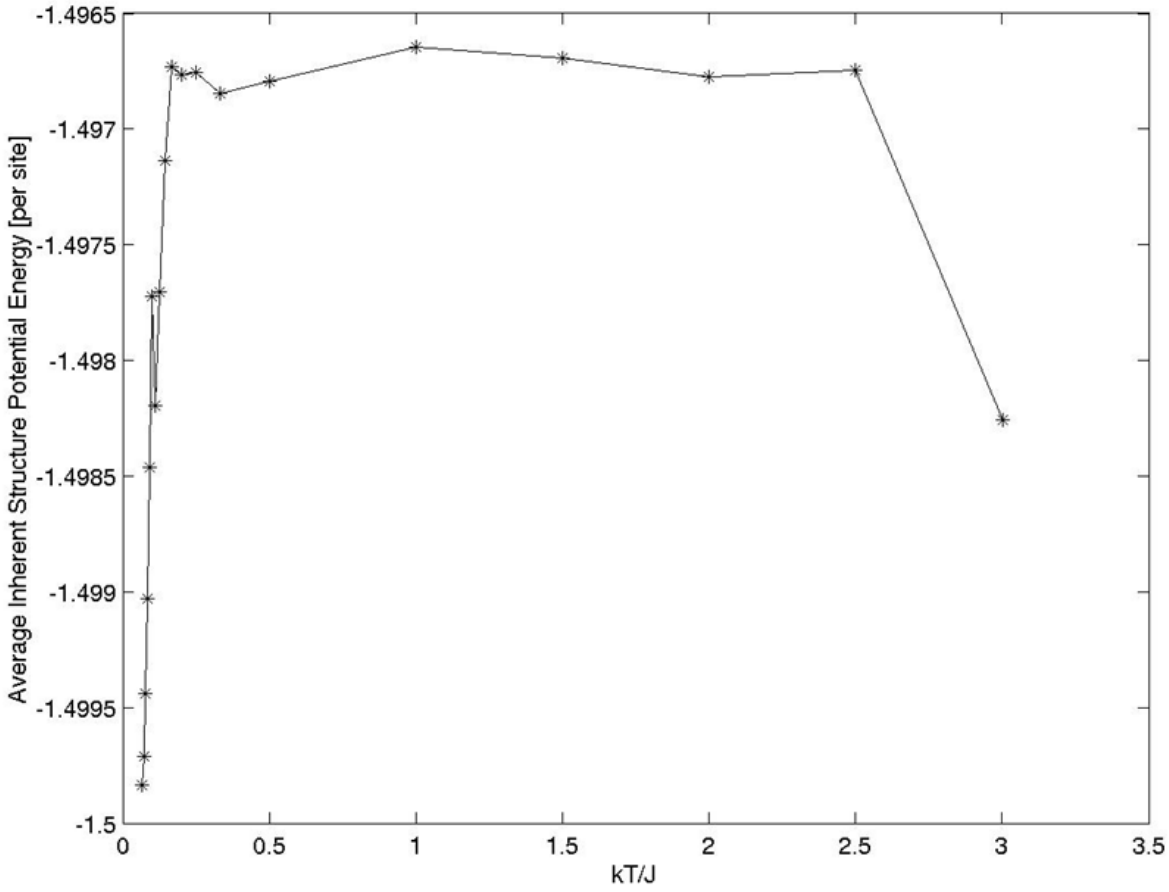

Figure G-1: Average inherent structure potential energy as a function of T
Beginning with an equilibrated Metropolis MC simulation at the T reported. The lattice size is 40 by 40 with a TAM at every site. Average potential energy of the 100 inherent structures is shown. Recall that the random (infinite temperature) potential energy of this model is –1.125 and the potential energy of a fully satisfied structure is –1.5.