

# Efficient Monte Carlo Simulations of Gas Molecules Inside Porous Materials

Jihan Kim<sup>\*,†</sup> and Berend Smit<sup>†,‡</sup>

<sup>†</sup>Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, United States

<sup>‡</sup>Department of Chemical and Biomolecular Engineering and Chemistry, University of California, Berkeley, California 94720, United States

**ABSTRACT:** Monte Carlo (MC) simulations are commonly used to obtain adsorption properties of gas molecules inside porous materials. In this work, we discuss various optimization strategies that lead to faster MC simulations with CO<sub>2</sub> gas molecules inside host zeolite structures used as a test system. The reciprocal space contribution of the gas–gas Ewald summation and both the direct and the reciprocal gas–host potential energy interactions are stored inside energy grids to reduce the wall time in the MC simulations. Additional speedup can be obtained by selectively calling the routine that computes the gas–gas Ewald summation, which does not impact the accuracy of the zeolite’s adsorption characteristics. We utilize two-level density-biased sampling technique in the grand canonical Monte Carlo (GCMC) algorithm to restrict CO<sub>2</sub> insertion moves into low-energy regions within the zeolite materials to accelerate convergence. Finally, we make use of the graphics processing units (GPUs) hardware to conduct multiple MC simulations in parallel via judiciously mapping the GPU threads to available workload. As a result, we can obtain a CO<sub>2</sub> adsorption isotherm curve with 14 pressure values (up to 10 atm) for a zeolite structure within a minute of total compute wall time.

## ■ INTRODUCTION

Porous materials such as zeolites, metal–organic frameworks, and ZIFs have been identified as important classes of materials suitable for a variety of new industrial applications, which include CO<sub>2</sub> capture, gas separations, and gas storage.<sup>1–6</sup> The pore topologies and chemical composition of the structures are key factors that determine the utility of a given material for a particular application. Given that small changes in the material can potentially lead to significantly different adsorption properties, the number of distinct materials that can in principle be synthesized is too large of a number for each of the structures to be studied in great detail. For example, although the number of known zeolite frameworks is around 201,<sup>7</sup> this number constitutes a very small fraction of more than 5.4 million structures within a hypothetical zeolite database that are feasible based on theoretical grounds.<sup>8,9</sup> Also, accordingly, it remains a daunting task for scientists to characterize millions of structures within a reasonable time and to successfully screen optimal structures for an application of interest.<sup>10–12</sup> In our previous work, we have demonstrated that our graphics processing units (GPU) code can characterize low pressure regime of a zeolite structure in just a few seconds via efficiently computing the Henry coefficient and the heats of adsorption values.<sup>13</sup> Subsequently, we managed to screen over 130 000 zeolite structures<sup>14</sup> (reduced from the original set of over a million based on the free sphere size via Zeo++<sup>15</sup>) to find the optimal structures within both the IZA and the hypothetical zeolite database.<sup>16</sup> However, for certain applications, there remains a need to fully characterize the adsorption properties of a material at a much wider pressure range, and, accordingly, the Henry coefficient and the heats of adsorption values are not sufficient enough to convey all of the important information. Alternatively, grand canonical Monte Carlo (GCMC) simu-

lations are commonly used to characterize adsorption properties of a material across a large range of pressure values, and, as a result, this work focuses on various techniques used to optimize these simulations. An efficient GCMC simulation allows full characterization of a very large database of materials possible and provides insights into material design for various applications. We again utilize the GPU as our hardware of choice, but many of the optimization techniques introduced here are hardware-independent as they can be applied to the CPU as well.

This Article is organized as follows. In section II, we discuss in detail our efficient MC algorithm and provide details regarding the optimization techniques to speed up the simulations. In section III, we provide the performance results of our implementation, and finally, in section IV, we summarize the important findings in our work.

## ■ GRAND CANONICAL MONTE CARLO METHOD FOR POROUS MATERIALS

Although this Article focuses on systems that contain CO<sub>2</sub> gas molecules inside zeolite materials, the optimization techniques that can speed up the MC simulations can readily be applied to other systems as well. Moreover, although we focus on grand canonical Monte Carlo (GCMC) simulations in our analysis, the results can be extended to other types of MC simulations (e.g., NVT, NPT) as well.

The allowed moves in the GCMC simulation consist of translation, insertion, and deletion with the following acceptance probabilities:

Received: May 9, 2012

Published: June 15, 2012

$$\text{translation:acc}(o \rightarrow n) = \min(1, e^{-\beta\Delta E})$$

$$\text{insertion:acc}(N \rightarrow N + 1) = \min\left(1, \frac{V\Lambda^{-3}}{N + 1} e^{\beta\mu} e^{-\beta\Delta E}\right)$$

$$\text{deletion:acc}(N \rightarrow N - 1) = \min\left(1, \frac{N}{V\Lambda^{-3}} e^{-\beta\mu} e^{-\beta\Delta E}\right)$$

where  $N$  represents the total number of gas molecules in the simulation box,  $o$  and  $n$  represent the old and the new states in the MC moves,  $\beta = 1/(k_B T)$  with  $T =$  temperature, set fixed at 300 K for all of the simulations,  $\mu = k_B T \ln \Lambda^3 \rho$  being the chemical potential of the system,  $\Lambda$  is the thermal de Broglie wavelength, which is a function of the partial pressure, and  $\rho$  is the particle density for the ideal gas system.  $\Delta E$  represents the change in the total energy of the system from the old to the new state after a single MC cycle. Throughout this work, 75% of the total MC moves are set to be insertion/deletion moves, while the remaining 25% are translation moves. The ratio of the MC moves can be varied and optimized to obtain faster convergence, but this part is ignored in our analysis and will not affect the performance analysis in this Article. For simplicity, we also ignore other commonly used MC moves such as swap and rotation.

Interactions between the gas–gas and the gas–host particles are modeled by the shifted Lennard-Jones and the periodic Ewald potentials. For systems with much larger number of particles as compared to what is typically observed in zeolites near full loading (e.g., less than 200 per unit cell), other methods such as multigrid methods and particle mesh Ewald are preferred.<sup>17</sup> The positions of the framework atoms are fixed during the MC simulations as the rigidity of zeolite framework has been shown to be a reasonable approximation in obtaining accurate adsorption properties.<sup>18</sup>

All of the force fields in this work come from Garcia-Perez et al. and Dubbeldam et al.<sup>19,20</sup> The gas–host interactions are precomputed and stored inside an energy grid in GPU global memory before the MC simulation, and utilized as a lookup table during the MC cycles.<sup>21</sup> The creation of the energy grid is especially important to block out (i.e., set to high energy) regions within porous materials that are diffusively inaccessible within the experimental time scale. The blocking algorithm is described elsewhere as we utilize an algorithm similar to one used in flood fill.<sup>13</sup>

## OPTIMIZATION TECHNIQUES

**Gas–Gas Tabulation Grid.** As mentioned in the previous section, the gas–host grid is necessary to correctly identify regions limited by diffusion inside the zeolite materials. In addition, the grid can be used to accelerate the MC simulations given that explicit gas–host interaction calculations will be avoided. With the gas–host grid in place, the most time-consuming part of the MC simulation takes place inside of a routine that computes the reciprocal components of the Ewald summation:<sup>22</sup>


$$U_{ij}^k(r) = \frac{1}{2L^3} \sum_{k \in K^3, k \neq 0} \frac{4\pi}{k^2} \exp(-k^2/4\alpha^2) q_i q_j \exp(-i\vec{k} \cdot \vec{r}_{ij}) \quad (1)$$

where  $U_{ij}^k(r)$  is the reciprocal  $k$ -space pairwise interactions between particles  $i$  and  $j$ . For each of these pairwise potential calculations, the total number of  $k$  vectors can range between

500 and 1000 for a typical zeolite structure, making their calculations significantly more expensive than the real-space Coulomb calculations. Similar to the gas–host grid, a gas–gas grid that contains the reciprocal components of the pairwise interaction terms as a function of pairwise distance can be created to provide significant speedup to circumvent computing explicit pairwise calculations.<sup>23</sup> The gas–gas grid can be generalized such that the charges,  $q_i$  and  $q_j$ , can be specified at the point of computation, which allows a single energy grid to account for many different types of gas molecule mixtures and thereby minimize GPU DRAM allocation. It turns out that within the GPU code, the energy grid array read operations from memory comprised of a large proportional overall wall time as all of the memory transactions are global, and unlike in the case of gas–host reads, the operation scales  $O(N)$  (as opposed to  $O(1)$ ) as for each proposed MC move, so there needs to be  $N - 1$  read operations for each unique pairwise interactions. However, the overall performance gained from reduced number of floating operations far outweighs the performance loss from additional memory read operations. As we will demonstrate later, the reciprocal energy components of the gas–gas interaction in zeolite are relatively very small as compared to the total pairwise interaction energy and accordingly can be omitted without much loss in accuracy. Thus, an additional speedup can be gained by neglecting to evaluate the reciprocal terms until at the end of the equilibration cycles, while accurately accounting for all of the interactions during the MC production cycles.

**Density-Biased Sampling.** In many of the zeolite structures (and porous materials in general), a large portion of the unit cell volume is considered to be inaccessible to the guest molecules due to dense packing of silicon and oxygen framework atoms in the crystalline structure. Moreover, even some of the low energy regions can be considered inaccessible due to diffusion limitations. To illustrate this point, we analyze zeolites MFI, LTA, RHO, and FAU whose coordinates are taken from the IZA database<sup>24</sup> and plot the proportion of gas–host energy grid points that are smaller than the specified energy values (horizontal axis) in Figure 1. Here,  $0k_B T$  represents the  $\text{CO}_2$  minima energy grid point values for each of the structures (MFI =  $-12.3k_B T$ , LTA =  $-11.1k_B T$ , RHO =  $-17.1k_B T$ , and FAU =  $-7.64k_B T$ ), where each of the grid points contains the Boltzmann-weighted average  $\text{CO}_2$  potential energies computed from collecting hundreds of randomized rotations of the molecule about the point. As can be seen from the curves, there is a large increase in  $\text{CO}_2$  occupation probability up until  $5\text{--}10k_B T$  above the minimum energy where the occupancy slowly saturates. Respectively, 23.9%, 36.8%, 37.0%, and 41.8% of the MFI, LTA, RHO, and FAU zeolite structures have regions where the average  $\text{CO}_2$  energies are below  $100k_B T$ . Even with the inclusion of gas–gas interactions that can bring down the potential energy of a given  $\text{CO}_2$  molecule,  $100k_B T$  is too high of an energy such that all of the regions that contain grid points above this value can be safely assumed to be inaccessible.

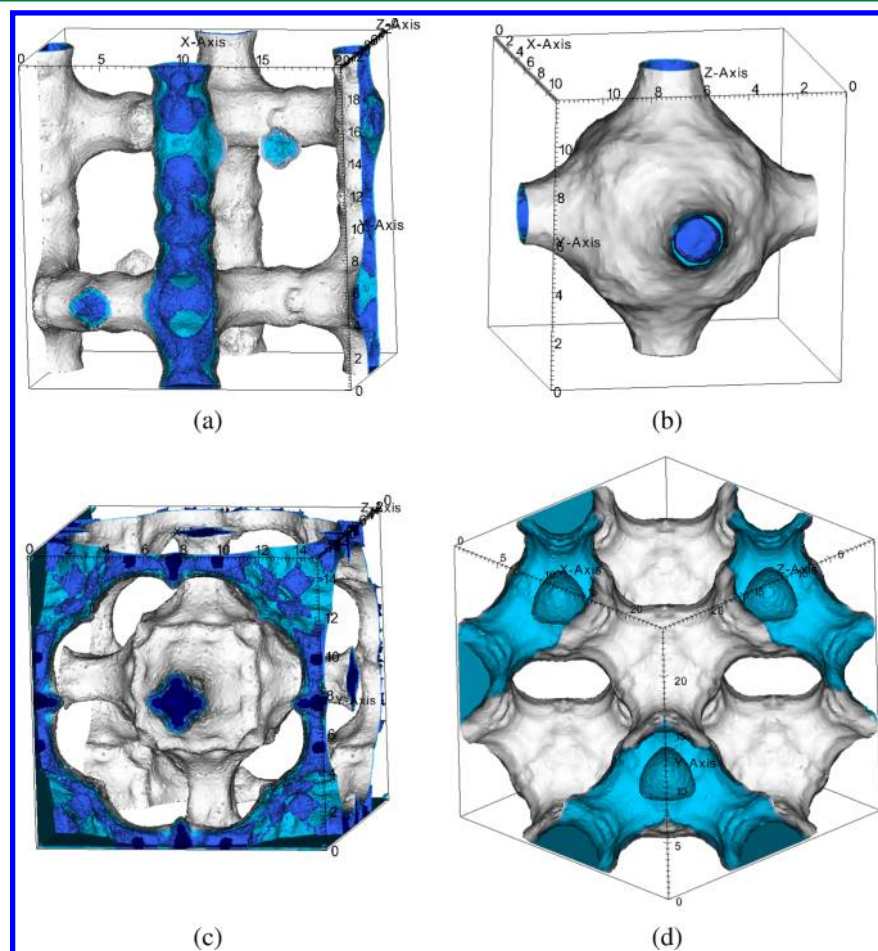
To illustrate these accessible regions within the zeolites, the  $\text{CO}_2$  potential energy contour profiles for MFI, LTA, RHO, and FAU are shown in Figure 2 with blue indicating low energy regions. These color-coded plots are generated from the GPU energy grid construction routine that is called before the start of the MC simulation. Accordingly, this information can be used to bias the GCMC insertion moves into low-energy regions such that the system can be equilibrated much more quickly



**Figure 1.** CO<sub>2</sub> potential energy distribution for zeolites MFI (red line), LTA (green line), RHO (blue line), and FAU (brown line). The minimum energy point for each of the four structures (i.e., MFI =  $-12.3k_B T$ , LTA =  $-11.1k_B T$ , RHO =  $-17.1k_B T$ , and FAU =  $-7.64k_B T$ ) is shifted to be  $0k_B T$  for illustration purposes. In the figure's vertical axis, probability indicates the proportion of energy grid points below a given  $k_B T$  from the horizontal axis. MFI, LTA, RHO, and FAU contain, respectively, 23.9%, 36.8%, 37.0%, and 41.8% of total energy grid points below  $100k_B T$ . Accordingly, the majority of the simulation volume is considered to be inaccessible due to very high potential energy.

from the density-biased sampling technique with quicker ascent toward the equilibrium uptake value. A similar technique has

been utilized in the lipid bilayer to improve the convergence of MC simulations.<sup>25</sup> The Metropolis criteria were changed to account for the modified insertion/deletion algorithm that no longer samples uniformly from the entire simulation volume to preserve detailed balance where the details are similar to what is shown elsewhere.<sup>25</sup> To eliminate possible biases, the reduced volume is uniformly sampled during the insertion moves and not just at the specific low-energy grid points. In our algorithm, we choose two CO<sub>2</sub> energy cutoff values,  $E_{cut,1}$  and  $E_{cut,2}$  (where  $E_{cut,1} = 5k_B T + E_{min}$  and  $E_{cut,2} = 300k_B T + E_{min}$ , with  $E_{min}$  indicating the global minimum energy point), to accelerate convergence. Above  $E_{cut,1}$ , the CO<sub>2</sub> occupation probability begins to decrease as seen from Figure 1, and the value was chosen to approximately capture most of the regions that are likely to be occupied by the CO<sub>2</sub>.  $E_{cut,2}$  was chosen such that any energy above this value is considered inaccessible ( $3\times$  larger than the  $100k_B T$  cutoff in Figure 1). For all  $i$  gas–host energy grid points where  $E_i < E_{cut,1}$ , we lumped the points that were distance-wise close to one another (i.e., within 2 Å) together into one single site. Accordingly, we identified between 5 and 50 local minimum sites within the zeolite materials where the likelihood of a single CO<sub>2</sub> molecule occupancy was very high. In our MC routine, the initial 1–2% of the equilibration cycles were spent with proposed insertion moves within the regions identified to be below the  $E_{cut,1}$  cutoff. During the rest of the equilibration cycles as well as the production cycles, insertion moves were sampled from all

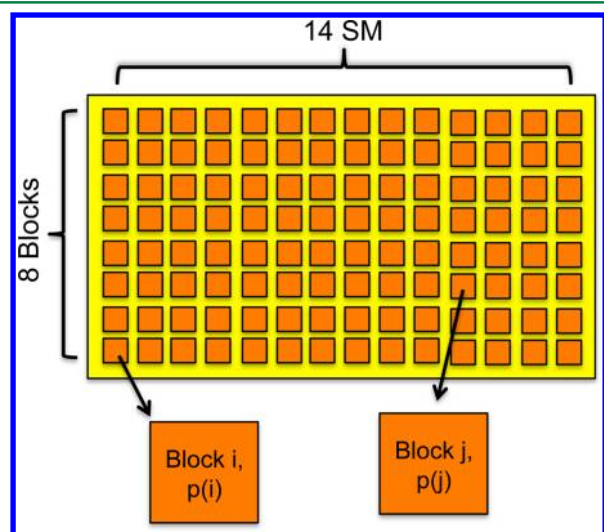


**Figure 2.** The CO<sub>2</sub> energy grids for zeolites (a) MFI, (b) LTA, (c) RHO, and (d) FAU. The blue regions indicate low energy regions where the likelihood of finding CO<sub>2</sub> is high.

energy grid points (roughly 10–50% of the total energy grid points depending on the structure) from values with  $E_i < E_{\text{cut},2}$ . Later in this Article, we demonstrate that the implementation of two-level density-biased sampling leads to faster convergence to the average loading values inside the zeolite structures.

**GPU Architecture Utilization.** The GPU architecture has at its disposal thousands of threads available for parallel processing. Unfortunately, there is not enough workload in our problem to efficiently allocate a large number of threads given that the ratio between the number of threads and number of particles remains large. In search for efficient utilization of the GPU threads, we decided to allow for multiple MC simulations to be conducted in parallel on the same material with independent simulations processing different pressure values to obtain a full isotherm calculation in a single simulation. By allowing the GPU threads to work on the same isotherm during a given simulation, the computational cost it takes to generate the energy grid can be incurred only once as these data can be shared by all of the parallel simulations.

The mapping between the GPU architecture and the GCMC system is illustrated in Figure 3. In the default setting, the GPU



**Figure 3.** The mapping between CUDA blocks and GCMC simulations. A given CUDA block,  $i$ , conducts its own independent GCMC simulation at pressure  $p(i)$ . In this particular configuration setting, eight blocks conduct GCMC simulation at the same pressure value to obtain more data for effective faster convergence in the GCMC simulations. The threads within the CUDA block parallelize the energy calculation routine, which is identified to be the bottleneck routine in MC calculations.

code launches  $14 \times 8$  CUDA blocks to concurrently process 112 independent MC systems. The number 14 (8) was chosen to be equal to the number of streaming multiprocessors (the maximum number of resident blocks) in the Tesla C2050 cards. Anywhere between 10 and 30 different pressure values are sufficient to obtain a general adsorption isotherm profile that can provide all of the essential adsorption information of a given material. In the default setting of the code, 14, user-defined pressure values can be processed by 8 different blocks with each conducting its own MC simulation with different randomized numbers. This implementation can effectively reduce the number of production cycles per block by 1/8, which can lead to a significant speedup in the overall wall time especially when the number of production cycles is far greater

than the equilibration cycles. The CUDA block size is set to be 32 threads as these threads work together to parallelize the energy calculations, which is the bottleneck routine in a given MC cycle.

All of our simulations were conducted on the Dirac cluster at the National Energy Research Scientific Computing (NERSC) Center utilizing the Nvidia Tesla C2050 Fermi cards. We utilized CUDA 4.0 and gcc 4.4.2 compiler, and the CURAND library is used to generate pseudorandom numbers during the Monte Carlo simulations.

## RESULTS

In all of our results, zeolites MFI, LTA, RHO, and FAU are chosen as test structures in the performance analysis as insights gained from them can extend to other zeolite structures well. The pressure values were chosen to vary from 0.001 to 100.0 atm, covering a wide range of values that are of interest to most applications related to gas adsorption.

First, we focus on the performance impact of utilizing the gas–gas grid as the computational cost associated with/without using the grid is summarized in Table 1 for a given  $\text{CO}_2$ – $\text{CO}_2$


**Table 1**

method	limiting factor	number of operations per pair potential
no grid	compute bound	$9k_x k_y k_z$ summations
gas–gas grid	memory bound	72 array reads

pair wise interaction. Without the gas–gas grid, the code computes  $9k_x k_y k_z$  (where typical numbers of  $k_x$ ,  $k_y$ , and  $k_z$  are around 8 or 9) summation terms for each of the atom–atom interactions within the two  $\text{CO}_2$  molecules. All of the terms associated with the explicit computation of the reciprocal components of the Ewald summation can be stored inside any of the fast GPU memory as the only variable, the position of the  $\text{CO}_2$  molecules, can fit inside the shared memory due to the small number of particles in our simulations. However, upon utilizing the gas–gas grid, the reciprocal component of the gas–gas interactions needs to be read from the precomputed grid inside the global memory (the number of gas–gas grid points typically are over a million and cannot be stored inside any of the fast GPU memory), and its evaluation can get very expensive. On the other hand, the evaluation of the gas–gas grid requires far less compute flops as the term does not scale with the total number of  $k$  vectors. For a given  $\text{CO}_2$ – $\text{CO}_2$  interaction, there are  $3(x, y, \text{ and } z) \times 3$  (three atoms within a  $\text{CO}_2$  molecule)  $\times 8$  (number of nearest energy grid points for interpolation) = 72 words that need to be read from memory to evaluate a given pair potential via interpolation. Some of these global memory transactions can be coalesced because some of the data are located within a contiguous block of the global memory. However, there does not seem to be a data structure in which all eight nearest neighboring energy grid points are located within a contiguous memory block, and, accordingly, multiple memory transactions from the gas–gas grid are required. Overall, utilizing the gas–gas grid makes the overall MC routine memory-bound as the number of memory transactions becomes very large as compared to the number of floating point operations.

The speedup (i.e., the ratio between compute wall time without utilizing the gas–gas grid and the time with the grid) varies depending on the uptake (i.e., loading) level of the gas molecule. To illustrate this point, the GCMC simulations for

the LTA zeolite structure were conducted at different pressure values to observe how speedup changes as a function of loading (Figure 4). Speedup numbers were obtained by comparing the




**Figure 4.** Speedup (no grid and with grid) as a function of loading for the LTA zeolite. The maximum speedup of 44.24 $\times$  occurs at the loading value of 0.0528 mol/kg. Similar behaviors (not included) are seen for other zeolite structures as well.

wall times (with/without grid) for the code to process 10 000 production MC cycles. Postequilibration wall times were used to ensure that loading values were similar during all 10 000 MC cycles to clarify the relationship between loading and performance.

The results from Figure 4 indicate that speedup does not increase monotonically with respect to CO<sub>2</sub> loading. Near zero loading, the system is empty during most of the MC cycles, and, subsequently, the time spent in evaluating the gas–gas grid is relatively short as compared to in other routines. Accordingly, in the infinite dilution case at low pressure regime, comparable wall times are observed for grid/no grid simulations. For larger loading values (i.e., 0.001–0.1 mol/kg), the speedup number increases significantly as the explicit evaluation of reciprocal Ewald summation terms become the bottleneck routine. On the other hand, in this regime, the number of CO<sub>2</sub> molecules is yet small enough such that the gas–gas grid evaluation was still comprised of a small portion of the overall MC (without the grid) time. These two effects combine to provide an overall increasing trend in the speedup curve, and a maximum speedup of 44.24 $\times$  occurs near the loading value of 0.0528 mol/kg. At even higher loading values, the wall time spent in reading from the global memory for the gas–gas grid becomes the bottleneck routine and the speedup starts to decrease. We observe that the speedup value drops to around 14 at high loading value (10 mol/kg), which corresponding to a pressure of 50 atm.

Next, we explore the impact that the various components of the Coulomb energy have on the isotherms. At low pressures, the number of gas–host interactions is much larger than gas–gas interactions as the total number of CO<sub>2</sub> molecules remain small in this pressure range. Accordingly, in practice, the coarse modeling of gas–gas interactions can still lead to accurate isotherm results while possibly speeding up the code as

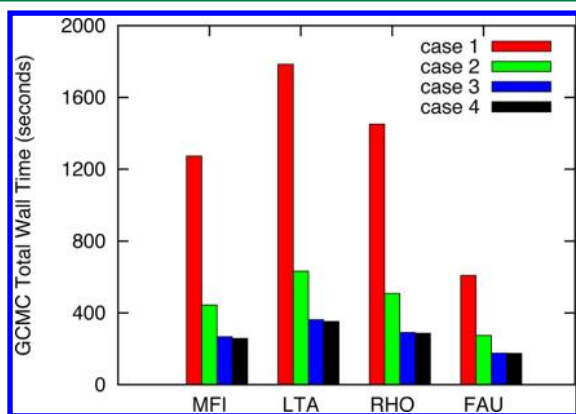


**Figure 5.** Adsorption isotherm curves for (a) MFI, (b) LTA, (c) RHO, and (d) FAU for case 1 (full calculations of reciprocal and direct Coulomb terms), case 2 (selective calculations of reciprocal Coulomb terms), case 3 (omission of reciprocal Coulomb terms), and case 4 (omission of reciprocal and direct Coulomb terms).

compared to cases where more accurate descriptions of the pairwise interactions are used. To determine both the accuracy and the performance for different ways of treating the gas–gas interactions, we take a look at four different cases (from the most accurate treatment to the least accurate): (1) full incorporation of both reciprocal and direct Coulomb terms during the whole equilibration and production MC cycles, (2) incorporation of the reciprocal Coulomb terms during the last 10% of the equilibration cycles and the entire production MC cycles, (3) omission of the reciprocal Coulomb terms, and (4) omission of both the reciprocal and the direct Coulomb terms. In this context, the direct Coulomb term refers to any real-space gas–gas interactions within a cutoff value of 12 Å. The number of equilibration (production) GCMC cycles were set to be 2 million (1.6 million) with the pressure values ranging from 0.075 to 25.0 atm.

The resulting adsorption isotherms for MFI, LTA, RHO, and FAU are shown in Figure 5. For case 4, the exclusion of all Coulombic gas–gas interactions results in underestimation of real loading at high pressure values for all of the structures. This omission plays less of an important role in MFI (Figure 5a), whose adsorption characteristics are predominantly dominated by the gas–gas Lennard-Jones van der Waals interaction and/or the gas–host interactions. Overall, as long as the direct gas–gas Coulomb terms are properly taken into account, there do not seem to be significant differences in the loading values as seen from cases 1, 2, and 3. This seems to indicate that the Lennard-Jones and the direct Coulomb pairwise interactions for the gas–gas interactions are sufficient to capture the important details of the material’s adsorption characteristics. For systems in which omission of the Fourier components causes larger differences in the loading values, we might need to adjust the equilibration with full potential to be more than 10%.

Next, we take a look at the performance analysis for MFI, LTA, RHO, and FAU for the four cases mentioned previously. Because the read operations from the gas–gas grid can become expensive, neglecting the gas–gas evaluation routine altogether can provide additional speedup on top of the speedup gained from using the grid in the first place. The GCMC simulation compute wall times in Figure 6 indicate that there exists about 4.5× average (that is, average for the four zeolites) speedup upon going from case 1 to case 4. The wall times for case 3 and 4 are similar to one another as excluding the direct Coulomb

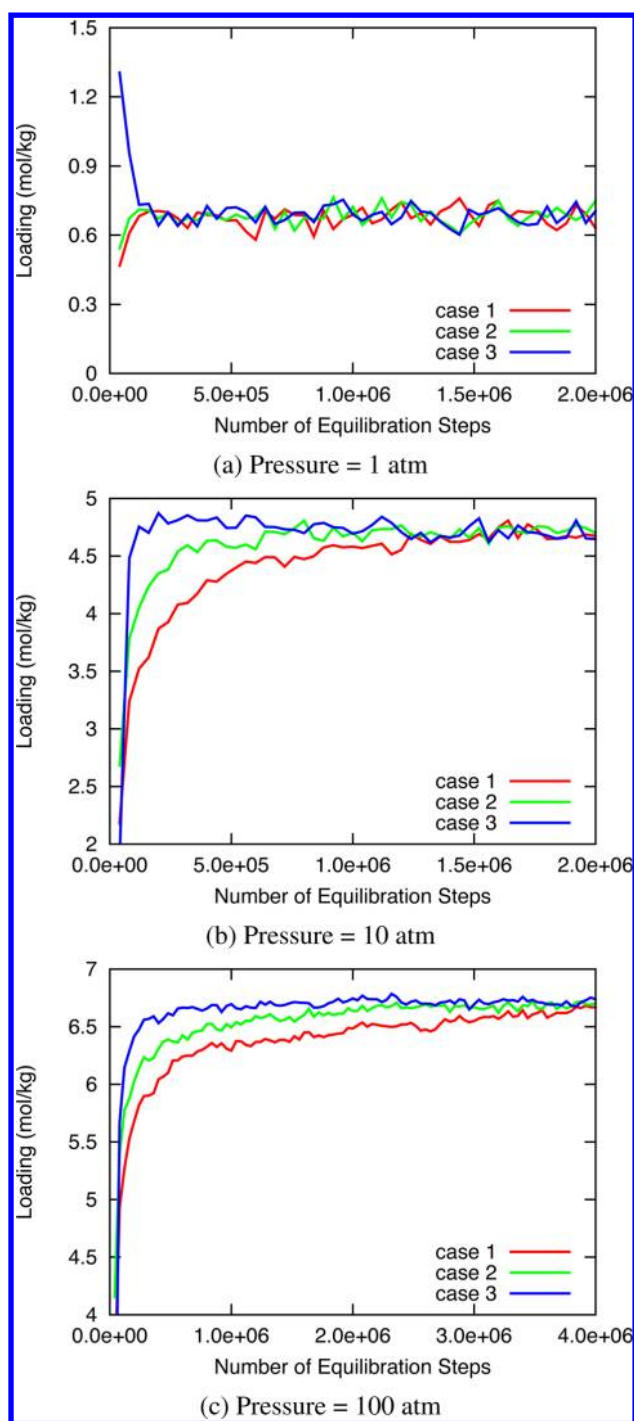


**Figure 6.** The total GCMC wall time in MFI, LTA, RHO, and FAU for four different cases with regards to treating the gas–gas interactions. For all simulations, the total number of MC cycles was set to be 2 million for equilibration and 1.6 million for production.

term in the calculations does not impact performance greatly. This can be explained from the fact the most expensive floating point operations (i.e.,  $1/r$ ) associated with evaluating the direct Coulomb terms are already computed for the Lennard-Jones interactions, and accordingly the Coulomb term does not incur much additional cost. Overall, to efficiently and to accurately compute the adsorption isotherms, we choose case 2 as an optimal strategy. In the limiting case where the number of production cycle steps is relatively small as compared to the equilibration steps, the performance loss from choosing case 2 over case 1 comes only from the last 10% of the equilibration cycles. Moreover, we believe that it remains important to correctly account for the reciprocal terms in the gas–gas evaluation even though their impact on the adsorption characteristics of the examined zeolite materials is negligible. Particularly because there might be some exceptional circumstances in which the adsorption characteristics would be significantly different especially at very high loading. Overall, there is an average improvement of about 2.7× in going from case 1 to case 2 as seen from Figure 6.

Finally, we discuss the impact of the density-biased sampling technique on the convergence behaviors of the GCMC simulations. In simulating porous materials, the amount of time it takes to equilibrate the system at high pressure values can become lengthy as most of the void space becomes occupied with the gas molecules. Accordingly, the main source of error for equilibration is underestimation, rather than overestimation, of uptake values at high pressure values. As mentioned in the previous section, we impose two cutoff values,  $E_{\text{cut},1} = 5k_{\text{B}}T + E_{\text{min}}$  and  $E_{\text{cut},2} = 300k_{\text{B}}T + E_{\text{min}}$  that can potentially aid in accelerating convergence. In this analysis, we take a look at three different cases: (1) uniform sampling of the entire simulation volume (i.e., no cutoff), (2) sampling energy smaller than  $E_{\text{cut},2}$  and (3) sampling energy smaller than  $E_{\text{cut},1}$  for 2% of the initial equilibration cycles and sampling energy smaller than  $E_{\text{cut},2}$  for the rest of the cycles. We specifically focused on the LTA structure at three different pressure values: 1, 10, and 100 atm. At 1 atm (Figure 7a), the number of steps needed to equilibrate the system to its convergent value of 0.7 mol/kg remained roughly the same in the three cases. However, at higher pressure values of 10 and 100 atm (Figure 7b and c), utilizing the two-level cutoff energies in density-biased sampling (case 3) results in noticeably faster convergence to the equilibrium loading values as compared to the other two cases. Overall, the utility of density-biased sampling becomes greater at higher pressure values due to the larger number of failed insertion attempts in these dense systems. Moreover, we can also conclude that gas molecules mostly reside in adsorption sites that are within  $E_{\text{cut},1}$  of the minimum energy value for the zeolite materials and that the adsorption sites are mostly determined by the gas–host interactions. We can further reduce the number of MC cycles needed for convergence by changing both the cutoff energy values and the proportion of MC cycles spent in each of the two-level sampling stages. However, we have not fully explored finding the optimal setting as there exists a complicated interplay between the aforementioned parameters as well as heavy dependence on the imposed pressure value and the structure itself.

Overall, with all of our optimization techniques in place, we can compute the CO<sub>2</sub> adsorption isotherms of LTA, MFI, RHO, and FAU utilizing 500 000 equilibration cycles and 800 000 production cycles in 116.58, 92.23, 100.78, and 47.6 s, respectively, at 14 pressure values from 0.005 to 10 atm. There



**Figure 7.** The CO<sub>2</sub> loading as a function of number of equilibration steps for LTA at three different pressures: 1, 10, and 100 atm. For case 1, the sampling for MC insertion moves is taken from the entire simulation volume. For case 2, the sampling for MC insertion moves is taken from regions that have CO<sub>2</sub> energies lower than  $E_{\text{cut},2} = 300k_{\text{B}}T + E_{\text{min}}$ . For case 3, insertion moves are taken from energies below  $E_{\text{cut},1} = 5k_{\text{B}}T + E_{\text{min}}$  for the initial 2% of the total equilibration cycles.

can be an additional decrease in general wall times if we are interested in nonpolar gas molecules such as methane or ethane (no Coulomb interactions) or are interested in adsorption isotherm curves at a smaller pressure range (e.g., up to 1 atm). The performance numbers reported here indicate that one can compute over 100 000 CO<sub>2</sub> adsorption isotherms within a

reasonable compute wall time of a week utilizing less than 8 GPUs.

## CONCLUSIONS

In this work, we have made important modifications to the Monte Carlo algorithms to efficiently compute the isotherms for various zeolite structures. The gas–gas grid was created to accelerate the routine for computing the reciprocal terms in the Ewald summation. Given that gas–host interactions are predominantly responsible for the adsorption characteristics of a zeolite material, we can obtain accurate adsorption isotherms and further speed up the algorithm by selectively calling the routine that computes these reciprocal terms. For porous materials, the majority of the simulation volume is considered to be inaccessible, and we can presample the low energy regions and accelerate convergence with biased insertion moves in our Monte Carlo algorithm. Finally, judicious parallelization algorithm allows for efficient utilization of threads inside the GPU. Overall, we have achieved 2 orders of magnitude speed up in our efficient GPU code as compared to the CPU simulation code without the optimizations. The limitation of the current GPU code comes mostly from the small amount of available DRAM (Tesla C2050 card has 3GB DRAM) as compared to CPU memory. Accordingly, for structures with relatively large unit cell size (e.g., each length greater than 80 Å), the mesh will become coarse and adversely affect the accuracy of the simulations. In cases such as above, we would need to reduce the amount of DRAM usage by storing only the energy grid points that are smaller than some cutoff energy value or utilizing multiple GPUs. Fortunately, we have not yet encountered a system where this has become an issue.

Although the analysis has been constrained to CO<sub>2</sub> molecules inside zeolite structures, our work can extend to other systems and other types of MC simulations, which allow scientists to fully characterize a large database of materials in reasonable compute time, paving the way to better material design and synthesis.

## AUTHOR INFORMATION

### Corresponding Author

\*E-mail: jihankim@lbl.gov.

### Notes

The authors declare no competing financial interest.

## ACKNOWLEDGMENTS

B.S. was supported as part of the Center for Gas Separations Relevant to Clean Energy Technologies, an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number DE-SC0001015.

During the initial stage of this work, J.K. was supported by the Director, Office of Science, Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and during the final stage of this work by the U.S. Department of Energy under contract DE-AC02-05CH11231 through the Carbon Capture Simulation Initiative (CCSI).

## REFERENCES

(1) Auerbach, S. M.; Carrado, K. A.; Dutta, P. K. *Handbook of Zeolite Science and Technology*; Marcel Dekker: New York, 2004.

- (2) Smit, B.; Maesen, T. L. M. *Nature* **2008**, *457*, 671–677.
- (3) Krishna, R.; van Baten, J. M. *Chem. Eng. J.* **2007**, *133*, 121–131.
- (4) Millward, A. R.; Yaghi, O. M. *J. Am. Chem. Soc.* **2005**, *127*, 17998–17999.
- (5) Li, J.-R.; Sculley, J.; Zhou, H.-C. *Chem. Rev.* **2012**, *2*, 869–932.
- (6) Sumida, K.; Rogow, D. L.; Mason, J. A.; McDonald, T. M.; Bloch, E. D.; Herm, Z. R.; Bae, T.-H.; Long, J. R. *Chem. Rev.* **2012**, *2*, 724–781.
- (7) (a) Baerlocher, C.; Meier, W. M.; Olson, D. H. *Atlas of Zeolite Framework Types*, 7th ed.; Elsevier: Amsterdam, NL, 2007. (b) <http://www.iza-online.org/> (accessed Jan 1, 2010).
- (8) Foster, M. D.; Treacy, M. M. J. <http://www.hypotheticalzeolites.net> (accessed Nov 13, 2009).
- (9) Earl, D. J.; Deem, M. W. *Ind. Eng. Chem. Res.* **2006**, *45*, 5449–5454.
- (10) Wilmer, C. E.; Leaf, M.; Lee, C. Y.; Farha, O. K.; Hauser, B. G.; Hupp, J. T.; Snurr, R. Q. *Nat. Chem.* **2012**, *4*, 83–89.
- (11) Keskin, S.; Sholl, D. S. *J. Phys. Chem. C* **2007**, *111*, 14055–14059.
- (12) Krishna, R.; van Baten, J. M. *J. Membr. Sci.* **2010**, *360*, 323–333.
- (13) Kim, J.; Martin, R. L.; Ruebel, O.; Haranczyk, M.; Smit, B. *J. Chem. Theory Comput.* **2012**, *8*, 1684–1693.
- (14) Pophale, R.; Deem, M. W.; Cheeseman, P. A. *Phys. Chem. Chem. Phys.* **2011**, *13*, 12407–12412.
- (15) Willems, T. F.; Rycroft, C. H.; Kazi, M.; Meza, J. C.; Haranczyk, M. *Microporous Mesoporous Mater.* **2012**, *149*, 134–141.
- (16) Lin, L.-C.; Berger, A.; Martin, R. L.; Kim, J.; Swisher, J. A.; Jariwala, K.; Rycroft, C. H.; Bhowan, A.; Deem, M. W.; Haranczyk, M.; Smit, B. *Nat. Mater.* **2012**, *11*, 633–641.
- (17) Dserno, M.; Holm, C. *J. Chem. Phys.* **1998**, *109*, 7678–7693.
- (18) Smit, B.; Maesen, T. L. M. *Chem. Rev.* **2008**, *108*, 4125–4184.
- (19) Garcia-Perez, E.; Parra, J. B.; Ania, C. O.; Garcia-Sanchez, A.; van Baten, J. M.; Krishna, R.; Dubbeldam, D.; Calero, S. *Adsorption* **2007**, *13*, 469–476.
- (20) Dubbeldam, D.; Calero, S.; Vlugt, T.; Krishna, R.; Maesen, T.; Beerdsen, E.; Smit, B. *Phys. Rev. Lett.* **2004**, *93*, 088302.
- (21) Kim, J.; Rodgers, J.; Athènes, M.; Smit, B. *J. Chem. Theory Comput.* **2011**, *7*, 3208–3222.
- (22) Frenkel, D.; Smit, B. *Understanding Molecular simulation: From Algorithms to Applications*, 2nd ed.; Academic Press: San Diego, CA, 2002.
- (23) Toukmaji, A. Y.; Board, J. A., Jr. *Comput. Phys. Commun.* **1996**, *95*, 73–92.
- (24) <http://www.iza-structure.org/databases> (accessed January, 2011).
- (25) Rodgers, J.; Webb, M.; Smit, B. *J. Chem. Phys.* **2010**, *132*, 064107.