Speeding up simulation of diffusion in zeolites by a parallel synchronous kinetic Monte Carlo algorithm

Andrea Gabrieli, Pierfranco Demontis,* Federico G. Pazzona, and Giuseppe B. Suffritti

Dipartimento di Chimica, Università degli Studi di Sassari and Consorzio Interuniversitario Nazionale per la Scienza e Tecnologia dei Materiali (INSTM), Unità di Ricerca di Sassari, via Vienna, 2, I-07100 Sassari, Italy

(Received 10 February 2011; published 11 May 2011)

Understanding the behaviors of molecules in tight confinement is a challenging task. Standard simulation tools like kinetic Monte Carlo have proven to be very effective in the study of adsorption and diffusion phenomena in microporous materials, but they turn out to be very inefficient when simulation time and length scales are extended. In this paper we have explored the possibility of application of a discrete version of the synchronous parallel kinetic Monte Carlo algorithm introduced by Martínez *et al*. [J. Comput. Phys. **227**, 3804 (2008)] to the study of aromatic hydrocarbons diffusion in zeolites. The efficiency of this algorithm is investigated as a function of the number of processors and domain size. We show that with an accurate choice of domains size it is possible to achieve very good efficiencies thus permitting us to effectively extend space and time scales of the simulated system.

DOI: 10.1103/PhysRevE.83.056705 PACS number(s): 02.70.Tt, 47.56.+r, 66.30.−h, 82.75.Jn

I. INTRODUCTION

Zeolites are crystalline microporous aluminosilicates that have found a large number of uses in the chemical industry [1,2]. Their crystal structure consists of a definite channel and cage network extending in one, two, or three dimensions. The presence of regular micropores provides an environment where the adsorbed molecules no longer move freely, but are restricted to reduced spatial dimensions where peculiar many-body effects make zeolites behave as solid solvents [3]. The diverse physical phenomena occurring in these systems embrace heterogeneous catalysis, percolation, and even a dramatic change of the phase diagram [4]. All of them are ruled by the dimensionality resulting from the specific network of channels and cages that largely determines the nature of the local interactions and of the long-range order. Moreover, the molecular mobility is strongly influenced by the topology of the surrounding medium, which provides the energy landscape through the multifarious interplay between adsorbent-adsorbate and adsorbate-adsorbate interactions. Ranging from electronic transitions to slow molecular migration, a hierarchy of time scales and distances are involved in the many processes happening in the interior of the crystal, whose consequences are at the same time essential and difficult to quantify. These phenomena are still far from being understood, and despite a great deal of effort in theory and computation [5], a fundamental description of the confinement effect in zeolites is not yet available. In recent years a growing research in multiscale modeling and simulation schemes simple enough to be analyzed and able to capture the essential features of the real physical systems has been reported [6,7]. Effective and efficient multiscale modeling that bridge the gap between molecular-level interactions and macroscopic properties are essential to advance zeolite science and technology [8]. In this spirit this study investigates the efficacy and potential of using a parallel kinetic Monte Carlo (kMC) algorithm for multiscale zeolites modeling, and addresses some of the challenges involved in designing *competent* algorithms that solve hard problems quickly, reliably, and accurately.

This paper is organized as follows. Section II shortly summarizes the standard kMC method and outlines the most significant challenges in improving its performance. In Secs. III and IV we introduce the basis of the architecture and the design limits of a parallel version of the algorithm on a discrete system, and in Sec. V we discuss an application to a selected system.

II. THE MODEL

In a kMC simulation [9,10] a state of the system is represented by a configuration of molecules in a discrete network of sites, and a random walk is performed from state to state [11]. The most widely adopted kMC algorithm is *rejection-free*, meaning that at every step the system makes a transition from one state to another and the time t is advanced by extracting an interevent time from an exponential distribution, that is, $t = -\ln(u)/R_{\text{tot}}$, where u is a uniform (pseudo) random number in $(0,1)$ and $R_{\text{tot}} = \sum_i^n r_i$ is the sum of the rates r_i of all possible events n. The standard kMC method scales badly with the size of the system (i.e., with the number of events) because of two factors, namely (i) the time spent in generating, searching, and updating the list of events, and (ii) the proportionality of the interevent time to $1/R_{\text{tot}}$ which implies that, given the same number of iterations, the trajectory length decreases with increasing system size.

In a large system it is a fair hypothesis to assume that distant regions do not interact significantly with each other. This is the ground for a parallel kinetic Monte Carlo algorithm able to improve the standard method by overcoming its limitations. The underlying idea in parallelizing kMC is the partitioning of the system in domains, where it is possible to execute a sequential algorithm. The domains are independent of each other, consequently by assigning each domain to a different processor the number of events will be reduced, along with the value of R_{tot}. This in turn will raise the efficiency and lengthen the trajectory.

*demontis@uniss.it

The major problem in parallelizing kMC is the complete asynchronicity of the algorithm. In the rejection-free kMC, at every time step an event is selected and realized. The corresponding time depends on the rates of *all* the possible events. This implies that a parallel approach consisting only of executing serial kMC algorithms independently of each other is correct only if the noninteraction condition between domains is rigorously respected. In real systems this is unachievable since interactions or transfers of matter at the boundaries between domains cannot be avoided. Moreover, each domain has its own timeline and in order to avoid causality errors it is necessary to synchronize and correct them. Despite that, many methods were developed to rigorously treat these problems (see, for example, [12,13]) thus permitting us to have a parallel kMC procedure able to solve the same master equation of a sequential one. The major drawback of these methods is that they can be highly expensive and complicated to be implemented. The work of Martínez *et al*. [14] shows that ignoring the interaction between domains introduces an error that can be controlled through a careful choice of the domain size. This leads to a great simplification of the algorithm and improves the efficiency. Moreover, this method avoids causality errors by synchronizing the time across domains through the introduction of a *null event*.

III. PARALLEL ALGORITHM

Our parallel algorithm is a manipulation of the continuous synchronous kMC introduced by Martínez *et al*. [14] adapted to a discrete lattice. The first step is the spatial decomposition of the lattice in K domains (where K is equal to the number of processors) named Ω_k, $k = 1, \ldots, K$. The domain shape is arbitrary and the optimal choice, aimed to minimize the communication between domains, is strictly problem dependent. In principle, domains do not necessarily have to be equivalent. They can be assigned heterogeneous sizes and shapes to attain the best optimization possible. In the present case the domains are chosen to be (all equivalent) parallelepiped shaped (see Sec. V B). The simulation proceeds as follows.

(1) In each domain, say Ω_k, a list of the possible events n_k and relative rates r_{ik} ($i = 1, \ldots, n_k$) is generated. Rates can be summed to give a total rate R_k for each of the K domains:

$$R_k = \sum_{i}^{n_k} r_{ik}. \tag{1}$$

It is worth noting that if the system was not subdivided into domains, the value of R_k would be simply equal to the sum of the rates of all the events as in the sequential case. This implies that the subdivision does not alter the set of states the system can reach.

(2) The synchronicity of time horizon between domains is ensured by selecting the greatest among all values of R_k:

$$R_{\max} = \max_{k=1,\ldots,K} \{R_k\}, \tag{2}$$

and introducing in each domain for which $R_k < R_{\max}$ the possibility of a *null event*, that is an event in which no particle moves. The rate of the null event in the kth domain is defined as

$$r_{k0} = R_{\max} - R_k. \tag{3}$$

As a consequence, the domain having the greatest relative total rate equal to $R_{\max}$ will have no null event. Introducing null events is necessary to align the interevent time for the entire system on the time of the *fastest* evolving domain: this way the same interevent time can be chosen for all the domains as a function of only the maximum rate $R_{\max}$ and a random number. Despite the presence of null events in every of the $K - 1$ domains having $R_k < R_{\max}$, *globally* the algorithm is still rejection-free since inside the domain with $R_k = R_{\max}$ there is no null event, so that at each time step at least one molecule movement is realized.

(3) In each domain an event is selected out of the list of n_k events available to the kth domain, and realized with probability $p_{ik} = r_{ik}/R_{\max}$.

(4) If the outcoming configurations of two or more domains conflict with each other at their shared boundary, they are subjected to a correction procedure. Section IV is devoted to this topic.

(5) Interevent time is extracted from an exponential distribution:

$$\tau = \frac{-\ln(u)}{R_{\max}} \quad \text{with } u \text{ random number} \in (0,1). \tag{4}$$

(6) The entire procedure is iterated until final time is reached.

IV. CONFLICTING SITUATIONS AT THE DOMAIN BOUNDARIES IN DISCRETE SYSTEMS

In a continuous system a boundary conflict can arise if at a given time step the global outcoming configuration contains at least a pair of particles extremely close to each other. In discrete systems where a strict exclusion principle holds this translates to the much more likely situation where two or more particles are attempting to occupy the same lattice site.

There are basically two possible strategies for solving such a conflict: (i) the synchronous sublattice method [15] and (ii) a *rollback* procedure (see, for example, [13]). In the former every domain is further divided into sublattices having a size larger than the range of interactions. Conflicts are avoided by executing moves only in a randomly selected sublattice. In the latter instead conflicts are treated only when they occur. Indeed, rollbacks have a high computational cost. To speed up the simulation, Martínez *et al*. [14] avoided rollbacks by simply ignoring the conflicts. We have instead chosen to implement them anyway to avoid loss of synchronicity, and then to minimize their number by properly choosing the domains shape, thus compensating for the consequent slowing down of the simulation. The full time-horizon synchronicity of the domains allows the use of that procedure only when a violation of the exclusion principle occurs. In that case one proceeds as follows.

(1) Check for conflicting events across boundaries.

(2) For each conflicting pair of domains, the move to be undone is chosen through random selection of one of the two domains.

(3) Undo the chosen move, that is, restore the previous state of the list of rates.

(4) By using the same random number we have used for the realization of the conflicting event, a new move is performed [16] and the simulations goes on.

In other synchronous methods [13,17] each domain has its own history and time. At a fixed time interval one has to check the boundary events in order to verify if the generated timeline is consistent, then correct possible problems and eventually assign the proper time to obtain synchronicity. This procedure can lead to a certain number of moves to be undone, and its implementation is rather complicated. On the contrary, the method presented in this work looks very simple since the time horizon has been set up to be flat. This permits boundary events to be communicated immediately, and the maximum number of moves to be undone at each time step to be just one. Moreover, no causality error can arise. The main drawback is the increased communication cost, but this can be minimized by properly choosing the shape and the dimension of the domains. Even though the best domain choice is problem dependent, in general the ideal shape is the one that minimizes the number of communicating domains, and the ideal size is the largest possible in order to reduce the probability of a boundary event while still benefiting from the use of multiple processors in parallel, as we will show in Sec. V B.

The method is not rigorous for interacting particles where when a move happens to change the configuration at the boundaries then performing a rollback may not lead back to the starting configuration. This changes the value of $R_{\max}$ due to the addition of a particle in the selected domain, thus introducing an error. Nevertheless, the range of values the change in $R_{\max}$ might fall in is limited and independent of the size of the domains. Therefore a domain size can be found that minimizes such a range (e.g., enlarging the domain reduces the overall effect of the change). Moreover, such situations will happen with a relatively low frequency during the simulation if the domain size is chosen large enough so that the number of nonboundary sites is much greater than the number of sites at the boundaries, thus making the effect of conflicts negligible.

As its major strength and main advantage with respect to more complicated procedures, the nonrigorous approach presented here enables the error to be easily controlled thus allowing the same results of a rigorous method to be obtained [14], but with a simpler implementation and a faster execution.

V. APPLICATION TO A SELECTED SYSTEM: BENZENE IN NaX

Aromatic hydrocarbons are among the crucial ingredients of many plastic and allied materials. With the increase in the prices of crude oil, there is an urgent need to reduce the processing costs of aromatics while increasing the efficiency. This makes it necessary to bring about new technologies. The proposed [18] high activity for alkylation reactions of benzene with ethylene of Faujasite (FAU)-type zeolites to make styrene, one of the most relevant industrial monomers, offers the advantage of a high selectivity toward the desired product due to the shape-selective properties of their microcrystalline pore structures. In these applications the diffusive molecular transport through zeolites needs to be described accurately for a predictive design of the processes. However, the number of fundamental studies that investigate aromatics diffusion and adsorption in porous solids is limited due to the complexity of the system, the sluggish motion of aromatics in zeolites caused by the strong interactions between π electrons and extra framework cations and the large size of the aromatic species. Adsorption properties of aromatics in zeolites and other porous solids have been relatively less investigated as compared to alkanes in zeolites. This is particularly true if we consider only theoretical or computational studies. Demontis *et al*. were the first who investigated diffusion of benzene in NaY belonging to the FAU-type zeolites [Fig. 1(a)]. Their simulations suggest that benzene is frequently localized near the sodium cation and the 12-ring windows [19], in excellent agreement with the neutron diffraction study of Fitch *et al*. [20]. Auerbach *et al*. [21–25] studied the jump motion of the guest benzene molecules in a lattice site model of NaY, proving that cost-effective modeling techniques to simulate diffusive phenomena across multiple space and time scales lead to a significant gain, even if the price is loosing information at the intermediate scales. As a consequence, a multiscale modeling approach seems to be the proper choice to deal with this problem. It is our purpose in this paper to test the parallel synchronous kMC method with the aim of extending the modeling to the micro-millisecond time (and corresponding length) scales.

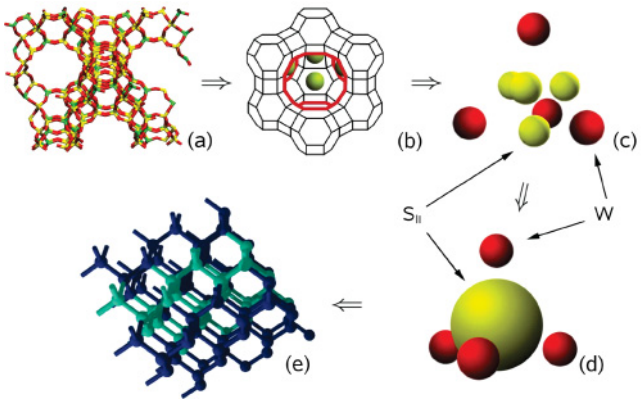

FIG. 1. (Color online) (a) Molecular structure of FAU-type zeolites. The three-dimensional framework of zeolites is constituted by a network of cages (b) connected by windows. The cages can accommodate a number of guest molecules adsorbed in well-defined binding sites. In NaX zeolite there are two types of these sites: four S_{II} (yellow/light gray spheres) inside the cage and four $S_{\mathrm{III}'}$ located in the window W connecting two cages (red/dark gray rings). In kMC simulations these sites are mapped on a detailed lattice (c) but it is possible to coarse grain the inner sites by stacking it on the center of the cage (d). Each particle can move from there to one of the four W sites (red/dark gray spheres). W to W moves are also possible. A move from S_{II} to S_{II} is possible but it produces no position change. (e) Schematic representation of zeolite FAU framework. Spheres represents coarse-grained S_{II} sites, while sticks represents W sites (for details refer to Sec. V). Distances are proportional to the real distances among cages. Different colors correspond to different domains.

TABLE I. Activation energies and preexponential factor at infinite dilution for benzene in NaX [26].

Jump	Activation Energy (eV)	Preexp. Factor (s^{-1})
$S_{II} \rightarrow S_{II}$	0.15	0.8×10^{13}
$S_{II} \rightarrow W$	0.25	0.8×10^{13}
$W \rightarrow S_{II}$	0.10	1.1×10^{12}
$W \rightarrow W$	0.10	2.4×10^{11}

A. Sequential Algorithm

We applied our method to the study of benzene diffusion in NaX, belonging to the FAU-type zeolites.

The diffusion of benzene in this type of system can be represented in the framework of the rare events dynamics, since residence times are much longer than travel times between adsorption sites. This implies that kMC is best suited to study it.

To test the parallel algorithm we first developed a model based on a previous work on this subject [26]. The zeolite framework is represented by a three-dimensional lattice of binding sites in bi-univocal correspondence with real ad sites. In the case of NaX and NaY there are two types of sites, S_{II} located over the Na^+ cation inside the cage [Fig. 1(b)] and W located on each window connecting two adjacent cages. NaY and NaX zeolites differ in the Na content, but the same lattice can be used for both zeolites [26]. Although it is difficult to determine the exact distribution of cations in NaX, experiments show that benzene is adsorbed at the S_{II} and $S_{III'}$ sites [27] [Fig. 1(b)]. The latter is very close to the 12-term oxygen ring, thus permitting this site to be viewed like the W site of zeolite Y. Assuming an Arrhenius behavior the dynamics of benzene is represented by jumps from site to site, with the rate constant calculated through considerations about the difference in energetic and geometric features of the two types of sites (see Table I).

The Hamiltonian for this lattice is [26]

$$H(\vec{s},\vec{\sigma}) = \sum_{i=1}^{M_W} s_i f_W + \frac{1}{2}\sum_{i,j}^{M_W} s_i J_{i,j}^{WW} s_j + \sum_{i=1}^{M_W}\sum_{j=1}^{M_{S_{II}}} s_i J_{i,j}^{WS_{II}} \sigma_j + \frac{1}{2}\sum_{i,j=1}^{M_{S_{II}}} \sigma_i J_{i,j}^{S_{II}S_{II}} \sigma_j + \sum_{i=1}^{M_{S_{II}}} \sigma_i f_{S_{II}}. \quad (5)$$

In this equation $\vec{s}$ and $\vec{\sigma}$ are the number of particles adsorbed in W and S_{II} sites, respectively (occupation numbers), $f_i = \epsilon_i - T\tilde{s}_i$ (Table II) is the free energy associated with the site i (ϵ_i is the adsorption energy and $\tilde{s}_i$ is the entropy), J is the interaction energy between nearest neighbor particles, and $M_W = 2M_{S_{II}}$ are the number of W and S_{II} adsorption sites, respectively. It is a common choice to ignore attractive interactions between particles leading to a simple site blocking model, but in the present case this cannot be done due to the critical temperature of benzene being 560 K [28]. To account for these interactions a parabolic jump model is adopted [26,29] where the change in the activation energy caused by the interactions is calculated as a function of the configuration in the neighboring sites. It assumes the transition state for a jump being located at the intersection of two parabolas, which is chosen to represent the minimum energy path among each pair of sites. The new value for the activation energy is obtained by [26]

$$E_a(i,j) = E_a^{(0)}(i,j) + \Delta E_{i,j}\left(\frac{1}{2} + \frac{\delta E_{ij}^{(0)}}{k_{ij}a_{ij}^2}\right) + \Delta E_{ij}^2\left(\frac{1}{2k_{ij}a_{ij}^2}\right). \quad (6)$$

$E_a^{(0)}(i,j)$ is the activation energy at the limit of infinite dilution. ΔE_{ij} represents the variation in adsorption energy between sites i and j due to interactions $\Delta E_{ij} = \delta E_{ij} - \delta E_{ij}^{(0)} = (E_j - E_i) - (\epsilon_j - \epsilon_i)$. In this equation $E_k = \epsilon_k + \sum_{l=1}^{M} J_{kl} n_l$ for a given configuration $\vec{n}$. Finally a_{ij} is the distance between two sites and k_{ij} is the harmonic force constant [26]

$$k_{ij} = \left(\frac{2}{a_{ij}}\right)^2 \left\{\frac{1}{2}\left[E_a^{(0)}(i,j) + E_a^{(0)}(j,i)\right] + \sqrt{E_a^{(0)}(i,j)E_a^{(0)}(j,i)}\right\}. \quad (7)$$

Previous works of our group with cellular automata models applied to the study of zeolites [30–32] have shown that it is possible to *coarse-grain* space and time scales by treating adsorption sites inside a cage as one single site. This leads for large systems to improving the efficiency without losses of physical information. Application of this coarse-graining paradigm to the model presented here leads to a lattice where all the S_{II} sites competing to each cage are grouped into a multiple-occupancy site placed at the cage center. The correct time evolution is guaranteed through the use of kMC rates for all the possible jumps between the various sites making up the central multiple-occupancy site, which is the scenario for all the intracage motions, while every intercage move requires the passage through a W site (Fig. 1).

Our sequential kMC was validated first by running several simulations to obtain self-diffusion coefficients to be compared with the experimental results [33]. We stress that the purpose of this comparison is to verify the correctness of the sequential algorithm and not to get new insight on the physical behavior of the system. The accomplishment of this task is postponed to a future work through the application of the parallel kMC method presented here and validated.

The simulations were carried out in a system containing 256 S_{II}-type sites and 128 W-type sites, corresponding to eight unit cells of NaX. As one can see in Fig. 2 (where the self-diffusivity D_{self} is plotted vs the coverage θ which is the number of molecules divided by the number of sites) the diffusion isotherms are in good qualitative agreement with the experimental data. The difference in the shape between the model and the experiments are expected due to the coarse graining of the S_{II} sites.

After that, other simulations were carried out to check the correctness of the parallel algorithm implementation. All calculation were executed on a cluster with Intel Xeon E5420

TABLE II. Adsorption energies and entropies [26].

ϵ_W (eV)	$\epsilon_{S_{II}}$ (eV)	$\tilde{s}_W$ (eV/K)	$\tilde{s}_{S_{II}}$ (eV/K)
−0.63	−0.78	1.7×10^{-4}	0

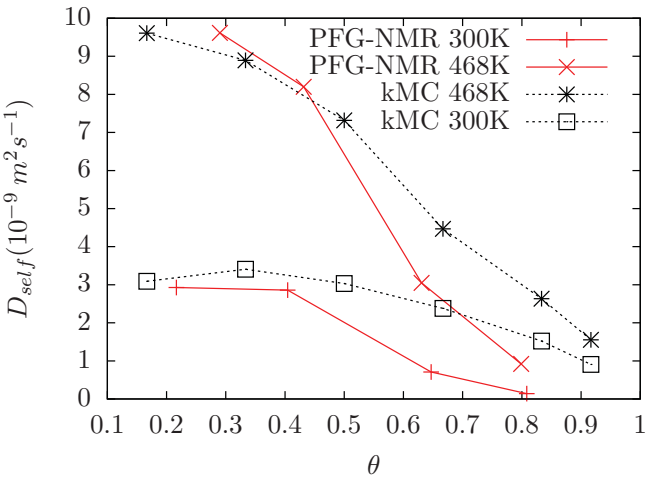

FIG. 2. (Color online) Diffusion coefficient as a function of coverage. Interaction parameter $J = -0.03$ eV. Experimental values (red solid lines, taken from Germanus *et al.* [33]) are multiplied by a factor of 10 and are shown only for a qualitative comparison with the simulation data (black dashed lines).

2.50 GHz processors and Infiniband communication link. For communications we made use of the MPI libraries. In Fig. 3 diffusion isotherms obtained by simulating a system of 4096 cages with an increasing number of processors are reported. As one can see from the plot, reducing the dimension of domains causes a slight shift in the value of the diffusion coefficient. The origin of this behavior is the error introduced with the parallel algorithm that can be easily controlled by choosing appropriate dimension for the domains. The choice is strictly problem dependent and must be assessed in each case.

B. Efficiency

To determine the efficiency of the method we made use of two definitions in order to better quantify the factors involved.

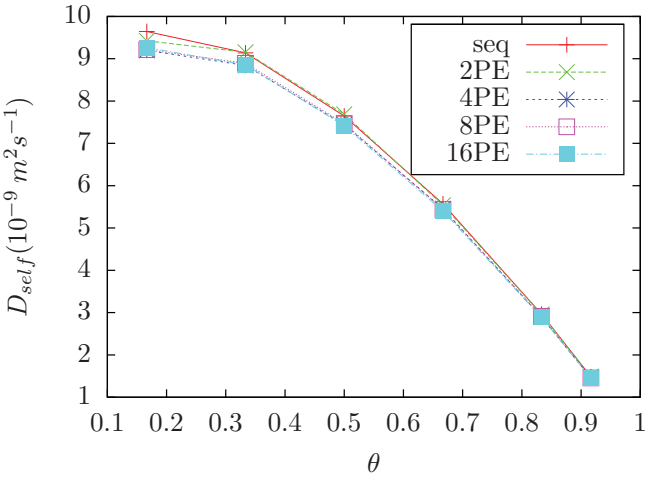

FIG. 3. (Color online) Diffusion isotherm for various number of processors on the same system size. Differences between sequential and parallel simulations are always small even for a relatively large number of processors (relative to the size of the system), and tend to converge for small numbers of processors.

In the first one $\tilde{\eta}$ each processing unit involved in the parallel runs simulates a portion of the system having the same size as the system simulated in the single-processor runs. This way the definition quantifies the efficiency on the basis of the cost of communications between processing elements, holding fixed the size factor [14,17]

$$\tilde{\eta} = \frac{t_{S,n}}{t_{K,nK}} \times 100\%, \tag{8}$$

where $t_{S,n}$ and $t_{K,nK}$ are, respectively, the time spent in executing the serial algorithm with n particles and the time spent in executing the parallel version with K processors on a system containing nK particles. Clearly the ideal efficiency of 100% is obtained when the time required to run the parallel code is the same as that required to execute the sequential one. This cannot be achieved in a real simulation because of the additional time required by the processors to communicate. In this particular implementation of the algorithm the major limitation is the need of global communications for updating the value of $R_{\max}$. It is important to note that despite this limitation, the impact of communication time over the global efficiency can be minimized by tuning the communication/calculation ratio. This is an easy task since almost every kMC algorithm scales with the size of the simulated system [34,35], so that it can be achieved by just finding the optimal value for the size of each domain.

The second definition is the speedup [17]:

$$S = \frac{t_S}{t_K}, \tag{9}$$

where t_S is the time required to execute the serial code and t_K the time required to execute the parallel code on K processors. Here the simulated system is assumed to be exactly the same (i.e., same size and same number of iterations) for both the serial and the parallel simulation. With this definition the speedup plot depends essentially on the scaling law the algorithm obeys [17] (in the present case the algorithm scales with the total number of particles) when the number of processors is low, and on the communication cost for higher numbers of processors.

Two sets of simulations were performed (see Table III) to study systematically the behavior of the parallel algorithm, starting with executing the sequential algorithm (used as a reference), and then increasing both the domain size and the number of processors (e.g., when using two processors the system consists of two identical replicas of the reference system and so on).

All the simulations have been carried out at a temperature of 468 K and a value of -0.02 eV for J (the nearest neighbor interaction energy). As one can expect the effect of increasing the size of the domains is an improved efficiency. This is because the communication/calculation ratio decreases. The possibility of modifying the system size to change this ratio is limited by the efficiency of the sequential algorithm adopted, that is the maximum system size that one can simulate with the parallel algorithm without reducing the length of the trajectory can be estimated as K times the maximum size attainable with a standard simulation (we recall that K is the number of processors).

TABLE III. The simulation sets considered in this work to estimate the importance of the communication/calculation ratio on the efficiency. Every subset spans several simulations at the same system size but different loadings starting from 1 ($\theta = 0.17$) up to 5.5 ($\theta = 0.92$) molecules per cage. The average number of sites per cage is six.

	Set A					Set B					
	Subset 1	Subset 2	Subset 3	Subset 4	Subset 5	Subset 1	Subset 2	Subset 3	Subset 4	Subset 5	Subset 6
No. cages	64	128	512	1024	2048	512	1024	4096	8192	163 84	327 68
No. processors	1	2	8	16	32	1	2	8	16	32	64

A determining factor of the efficiency is the domain shape. The domains must be chosen carefully on the basis of the system *topology* rather than the geometry. In the present case the node-to-node connections of the diamond-lattice topology of the FAU zeolite can be easily mapped onto a cubic grid. At this point it is straightforward to notice that such a grid can be better partitioned into *slices* rather than cubes [13], so that every domain (i.e., every slice) does communicate with two neighboring domains instead of six (see Fig. 4). Anyway, since the main reason of efficiency loss is the global communication caused by the need of synchronizing the domains, the choice of a cubic or a parallelepiped decomposition does not significantly affect the overall value of the efficiency $\tilde{\eta}$, but the cubic shape presents as one can expect a greater number of conflicting events.

The behavior of the efficiency is similar to that of the original method [14] and some others given in the literature [17] with a fit of the form

$$\tilde{\eta} = \frac{1}{1 + a(\ln K)^b}, \tag{10}$$

where a and b are two constants. For the first set of simulations a ranges from 0.020 to 0.268 and b ranges from 2.177 to 4.472, while for the second set the values range from 0.038 to 0.152 and from 1.889 to 2.743 for a and b, respectively. Differences in the value of b among different simulations can be related to the different values of the communication/calculation ratio. In the A set (see Table III) this ratio is greater than in the B set because domains are smaller, whereas the information exchanged between domains is the same (with a fixed number of processors), therefore the efficiency decays more steeply. Within each set the efficiency is influenced also by the communication/calculation ratio, which in this case, however, results from the combination of two opposite effects depending on the change of the total number of molecules adsorbed in the system. By increasing that number there can be more moves between domains that require more information exchanges and more rollbacks, thus increasing the ratio. On the other hand, increasing the number of molecules leads the number of events to increase as well, requiring then more computation. The balance between the different weights of these two effects causes the value of the parameters a and b to fluctuate. Moreover, these values differ from that obtained by Martínez and Merrick [14,17] mainly because of technical and algorithmic differences.

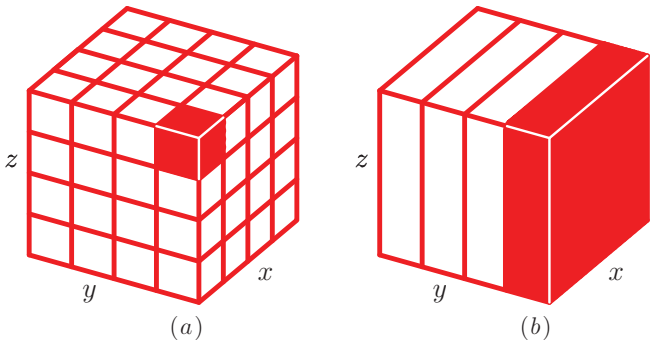

FIG. 4. (Color online) A comparison between (a) cubic domains and (b) slices. Each slice is a parallelepiped-shaped portion of the system spanning its whole extension in the y direction. This way, since with periodic boundary conditions each slice has no domain boundaries in the y direction, it does communicate with two domains only against the six of the cubic domain case.

In Fig. 5(a) the parallel efficiency $\tilde{\eta}$ [Eq. (10)] is reported as a function of the number of processors (K) for the simulation set A. A comparison with Fig. 5(b), where the same data are reported for the set B, makes clear the importance of the communication/calculation ratio which favors the B-set simulations (where the computation is much more expensive than in the A set). This leads to a greater efficiency of the algorithm when applied to set B in all the cases studied here. As expected, differences in the efficiency are more pronounced for numbers of processors greater than eight, since the increased cost of communication is not compensated by an equal increase in the computation cost. Finally, in Fig. 6 the speedup is reported. As stated before, its value is determined by a combination of two contributions, the cost of communications and the scaling law of the algorithm implemented. If we do not take into account the communication cost we would get the same computing time for both the serial and the parallel algorithm only if the algorithm were *not* scaling with the system size. Therefore the parallel implementation of our algorithm gives a substantial gain in the execution time, for example, a simulation of the set B3 (Table III) requires 59 h when using the serial algorithm and 8 h when using the parallel one on eight processors.

The only factor that may reduce the speedup is the number of rollbacks, since each roughly doubles the time spent for the current cycle. Anyway, this does not represent a problem in the present case where the number of rollbacks is kept relatively low (Fig. 7) by the particular topology of the system.

As a consequence, its influence over the efficiency of the method is limited and we obtain speedup values really close to the ideal efficiency of 100% (dashed line in Fig. 6, however we remark that the speedup is expected to decrease for a very large number of processors).

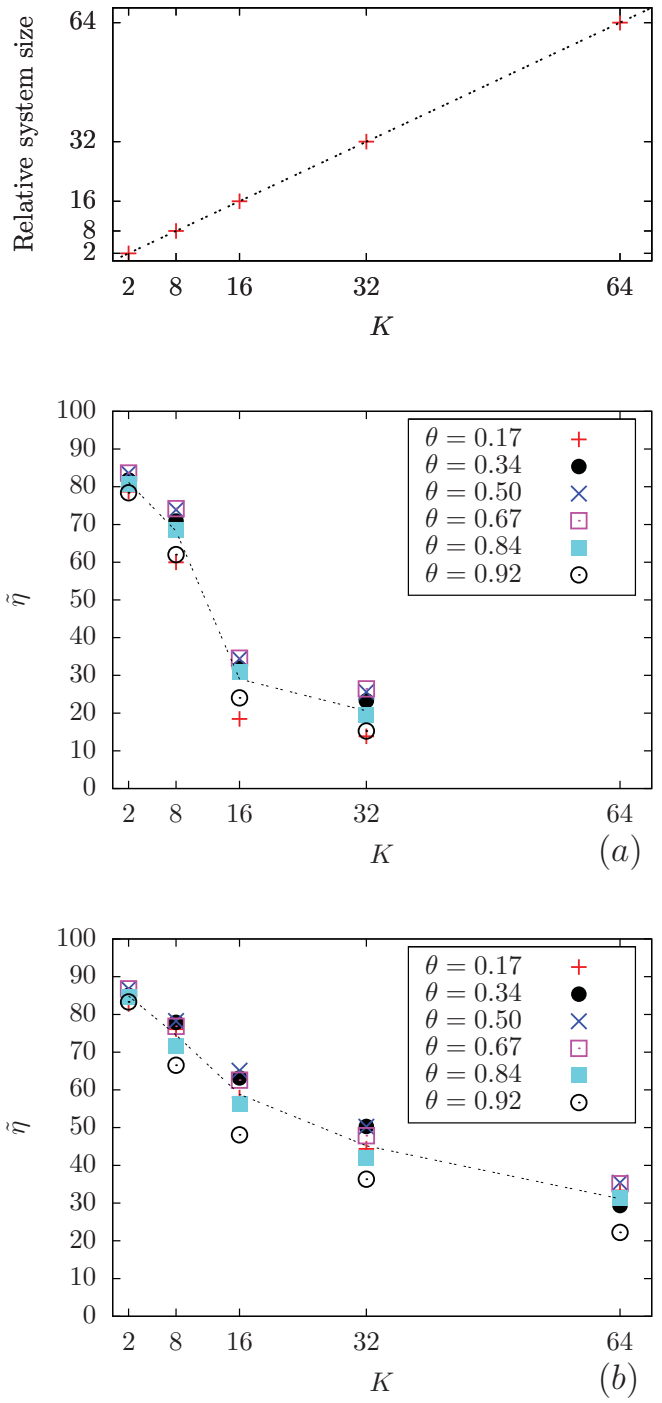

FIG. 5. (Color online) Parallel efficiency for simulation set (a) A and (b) B as a function of the number of processors at different loadings. For the parallel runs only, the system size increases linearly with the number of processors. On the top plot the ratio between the system size in the parallel runs and the system size in the serial, single processor run is shown. The ideal efficiency of 100% would be obtained only if the time required by the single-processor simulation of a system of a given size were the same as the time required by a parallel simulation on K processors of a system K times larger. Best results are obtained in set B because of the more favorable communication/calculation ratio. Dashed lines have been drawn to guide the eye.

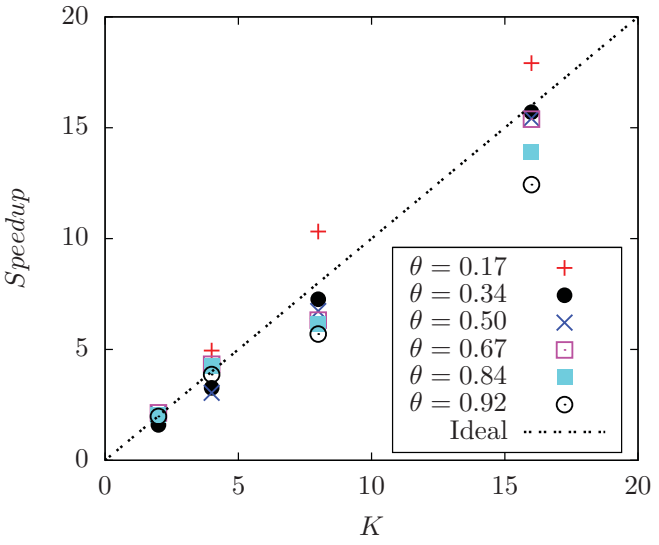

FIG. 6. (Color online) Speedup S [defined in Eq. (9)] of the parallel algorithm with respect to the sequential one shown as a function of the number of processors. This plot refers to the B3 system (Table III), other simulations show similar behavior and were not reported.

The number of null events is small as well (Fig. 7) and can be controlled through a proper choice of the system size. Consequently, during a simulation of set B approximately 90% of the possible moves are performed with no need of redefining the domain shape. The boost in the number of events per cycle is reported in the inset of Fig. 7.

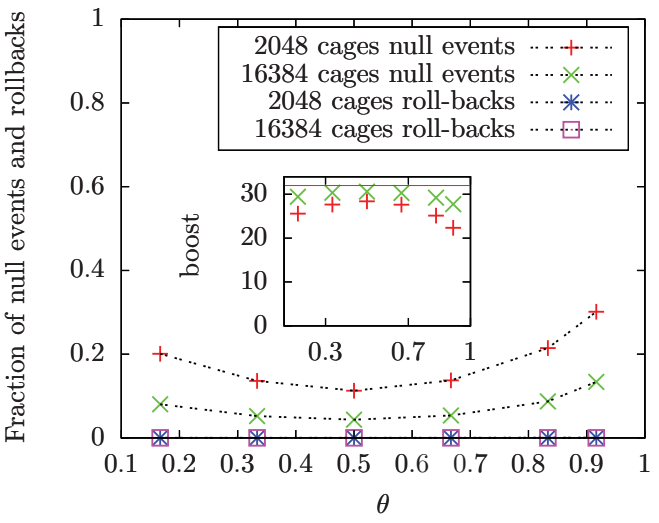

FIG. 7. (Color online) Fraction of null events and rollbacks for parallel simulations of systems A5 and B5 (Table III) on 32 processing units. The fraction of null events decreases with increasing the system size (leading to a 90% of effective moves), whereas the fraction of rollbacks is always negligible, even for a small system. Discrepancies among different coverages reflect the change in the relative number of possible events. Other simulations showing analogous behavior are not reported. The inset shows the average number of moves realized for every kMC step. The maximum of those values equals the number of processors, which is 32 in this case.

VI. CONCLUSIONS

A parallel kinetic Monte Carlo algorithm, originating from the synchronous algorithm of Martínez *et al*. [14], has been applied to the study of benzene diffusion in zeolite NaX. We have shown that, despite the presence of a rollback procedure in the algorithm, high efficiencies can be reached by exploiting the local nature of the molecule-molecule interactions inside the zeolite, allowing the need of rollbacks to be minimized through a proper spatial decomposition. In the present form the algorithm is still approximate, but the correct tuning of the domains size leads to obtaining results with the desired accuracy. We believe that the algorithm outlined here is applicable in general with little modification to other types of zeolites. Even better performances are expected to be found for other zeolites like the Linda Type A (LTA) family, ZSM5 [36], or for zeolitic imidazolate frameworks (ZIF) [37] because of the absence of shared sites between communicating cages. Adsorbate-adsorbate interactions does not extend significantly outside the cages, thus permitting an ideal domain decomposition. As for other similar methods [14,17], the efficiency of the algorithm is very sensitive to the value of the communication/calculation ratio that can be easily controlled by changing the size or the shape of the domains.

ACKNOWLEDGMENTS

This work has been carried out with financial support provided by the India-European Union collaborative project AMCOS (NMP3-SL-2009-233502), Italian Ministero dell'Istruzione, dell'Università e della Ricerca, by Università degli Studi di Sassari, and by Istituto Nazionale per la Scienza e Tecnologia dei Materiali (INSTM), which are gratefully acknowledged. F.G.P. is thankful to Regione Autonoma della Sardegna for the financial support within the "Master and Back" program.

[1] J. Kärger and D. M. Ruthven, *Diffusion in Zeolites and Other Microporous Materials*, 1st ed. (Wiley, New York, 1992).
[2] D. M. Ruthven, *Principles of Adsorption and Adsorption Processes*, 1st ed. (Wiley, New York, 1984).
[3] G. Sastre and A. Corma, J. Mol. Catal. A: Chem. **305**, 3 (2009).
[4] A. Huwe, F. Kremer, P. Behrens, and W. Schwieger, Phys. Rev. Lett. **82**, 2338 (1999).
[5] B. Smit and T. L. M. Maesen, Chem. Rev. **108**, 4125 (2008).
[6] A. C. Stuart, D. Collins, and D. G. Vlachos, J. Chem. Phys. **129**, 184101 (2008).
[7] F. G. Pazzona, *A Cellular Automata Model for Diffusion and Adsorption in Zeolites: Construction of a Mesoscopic Model* (LAP Lambert Academic, Saarbrücken, Germany, 2010).
[8] C. Colella, P. Aprea, B. de Gennaro, and B. Liguori, eds., *Proceedings of the 16th International Zeolite Conference* (A. De Frede, Naples, Italy, 2010).
[9] K. A. Fichthorn and W. H. Weinberg, J. Chem. Phys. **95**, 1090 (1991).
[10] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comput. Phys. **17**, 10 (1975).
[11] A. Voter, *Radiation Effects in Solids*, in *Proceedings of the NATO Advanced Study Institute on Radiation Effects in Solids*, edited by Kurt E. Sickafus, Eugene A. Kotomin, and Blas P. Uberuaga, NATO Science Series II: Mathematics, Physics and Chemistry, Vol. 235, pp. 1–23.
[12] B. D. Lubachevsky, J. Comput. Phys. **75**, 103 (1988).
[13] Y. Shim and J. G. Amar, Phys. Rev. B **71**, 115436 (2005).
[14] E. Martínez, J. Marian, M. H. Kalos, and J. M. Perlado, J. Comput. Phys. **227**, 3804 (2008).
[15] Y. Shim and J. G. Amar, Phys. Rev. B **71**, 125432 (2005).
[16] B. D. Lubachevsky and A. Weiss, CoRR **cs.DC/0405053** (2004).
[17] M. Merrick and K. A. Fichthorn, Phys. Rev. E **75**, 011606 (2007).
[18] S. Namuangruk, P. Pantu, and J. Limtrakul, J. Catal. **225**, 523 (2004).
[19] P. Demontis, S. Yashonath, and M. L. Klein, J. Phys. Chem. **93**, 5016 (1989).
[20] A. N. Fitch, H. Jobic, and A. J. Renouprez, J. Phys. Chem. **90**, 1311 (1986).
[21] S. M. Auerbach, N. J. Henson, A. K. Cheetam, and H. I. Metiu, J. Phys. Chem. **99**, 10600 (1995).
[22] S. M. Auerbach, L. M. Bull, N. J. Henson, H. I. Metiu, and A. K. Cheetam, J. Phys. Chem. **100**, 5923 (1996).
[23] S. M. Auerbach and H. I. Metiu, J. Chem. Phys. **105**, 3753 (1996).
[24] C. Saravanan and S. M. Auerbach, J. Chem. Phys. **107**, 8120 (1997).
[25] C. Saravanan and S. M. Auerbach, J. Chem. Phys. **107**, 8132 (1997).
[26] C. Saravanan and S. M. Auerbach, J. Chem. Phys. **110**, 11000 (1999).
[27] G. Vitale, C. F. Mellot, L. M. Bull, and A. K. Cheetam, J. Phys. Chem. B **101**, 4559 (1997).
[28] S. M. Auerbach, Int. Rev. Phys. Chem. **19**, 155 (2000).
[29] E. S. Hood, B. H. Toby, and W. H. Weinberg, Phys. Rev. Lett. **55**, 2437 (1985).
[30] P. Demontis, F. G. Pazzona, and G. B. Suffritti, J. Phys. Chem. B **110**, 13554 (2006).
[31] P. Demontis, F. G. Pazzona, and G. B. Suffritti, J. Chem. Phys. **126**, 194709 (2007).
[32] P. Demontis, F. G. Pazzona, and G. B. Suffritti, J. Chem. Phys. **126**, 194710 (2007).
[33] A. Germanus, J. Kärger, and H. Pfeifer, Zeolites **4**, 188 (1984).
[34] A. Chatterjee and D. Vlachos, J. Comput.-Aided Mater. Des. **14**, 253 (2007).
[35] T. P. Schulze, Phys. Rev. E **65**, 036704 (2002).
[36] C. Baerlocher, W. Meier, and D. Olson, *Atlas of Zeolite Framework Types* (Elsevier, Amsterdam, 2001).
[37] A. Phan, C. J. Doonan, F. J. Uribe-romo, C. B. Knobler, M. O'Keeffe, and O. M. Yaghi, Acc. Chem. Res. **43**, 58 (2010).