JCTC Journal of Chemical Theory and Computation

Article

pubs.acs.org/JCTC

Structure Formation of Toluene around C60: Implementation of the Adaptive Resolution Scheme (AdResS) into GROMACS

Sebastian Fritsch,* Christoph Junghans,‡ and Kurt Kremer

Max Planck Institute for Polymer Research, Ackermannweg 10, 55128 Mainz, Germany

ABSTRACT: For the example of C60 solutes in toluene, we present the implementation of the adaptive resolutions scheme (AdResS) for molecular simulations into GROMACS. In AdResS a local, typically all-atom cavity is coupled to a surrounding of coarse-grained, simplified molecules. This methodology can not only be used to reduce the CPU time demand of atomistic simulations but also to systematically investigate the relative influence of different interactions on structure formation. For this, we vary the thickness of the all atom layer of toluene around the C60 and analyze the first toluene layers in comparison to a full bulk simulation.

1. INTRODUCTION

Simulating large-scale systems on an all- or united-atom level still remains a very demanding task. Besides the problem of parametrizing appropriate and precise molecular interactions (force fields), often the time and length scales needed for a realistic description of the systems are inaccessible by 'brute-force' approaches. Because of this, methodologies have been developed over time to overcome these problems by devising simple and more idealized models, which can be treated much more efficiently in a computer simulation.[1−4] These include highly idealized models, which address generic or universal questions, such as the motion of a long polymer chain in a melt of identical polymers. More recently systematic coarse-graining procedures have been developed and very successfully applied to many (macro-) molecular problems. Through simplifying the descriptions, the properties of many materials can be appropriately studied on that level. Still, in all of these cases, the whole system is treated on one level of resolution, preventing the user from tackling problems where, in some regions of space, a more detailed description is needed. This issue can be addressed by the adaptive resolution scheme (AdResS), where within one simulation different regions coexist, which correspond to different levels of resolution.[5,6]Here, molecules can freely move from one region to the other without any barrier, which allows them to change their properties from fine grained to coarse grained and vice versa (Note that other approaches directly couple all-atom simulation to an implicit solvent representation).[7] By calculating the full detail only in some 'interesting' regions of the system and reducing the complexity for the rest, the computational cost can be reduced significantly. In the present contribution, we introduce the implementation of the AdResS concept into GROMACS and test this with the example of a C60 fullerene in a toluene solvent.

The example of C60 in toluene allows to discuss a second option to employ the AdResS concept. By studying the fullerene and the surrounding toluenes on an all atom and the toluenes further away on a coarse-grained level, we investigate static and dynamic properties of such a solution in a rather efficient way. AdResS however offers another different view on such a problem, which is difficult to address by other methods. By systematically varying the thickness of the all-atom layer around the C60 molecule, one can study the influence of structural support of the bulk solvent on the organization of the first toluene layers close to the solute and test what is needed to obtain experimental boundary conditions, where an infinite reservoir is supplying particles.

Understanding such complex fluids is of great interest for fundamental as well as industrial and biomedical research and goes far beyond the test of simulation concepts. Fullerenes[8] inspired many potential applications in organic electronics[9] as well as biopharmaceutics as carriers in drug delivery.[10] For this it is important to know their solubility, diffusion, and aggregation behavior. Computational models can be a guide to get an insight into these microscopic processes. C60 in toluene is especially interesting as X-ray scattering experiments[11,12] predict a local ordering of the liquid around the fullerene resulting in suppressed local density fluctuations, making it a perfect system to demonstrate AdResS in GROMACS.

We performed adaptive resolution simulations for different sizes of the 'all-atom' zone where the full atomistic degrees of freedom are present around the fullerene. By varying the thickness of the atomistic layer, the relevance of the surface interactions compared to the bulk structural influence can be investigated. The local ordering of the solvent is also correctly reproduced. We also show that the diffusion behavior can be closely reproduced even with the smallest all-atom zone. Thus in a large-scale simulation with many solute molecules, the major part of the solvent can be treated as coarse grained without affecting the local ordering at all and without affecting the diffusion behavior significantly even for the smallest all-atom zone. The computational efficiency and the speedup are discussed as well.

2. METHODS

2.1. AdResS. The AdResS method[5,6] allows for the smooth coupling of spatial domains of different degrees of resolution.

Received: October 4, 2011
Published: January 18, 2012

This (particle-based) method was successfully applied to couple coarse-grained[6] as well as path-integral[13] representations in molecular dynamics simulations. The former is the basis of the present study, in which a small spherical atomistic region around C60 is coupled to a coarse-grained region.

Figure 1 illustrates the simulation setup. The key aspect is a force interpolation scheme which describes how the pair forces

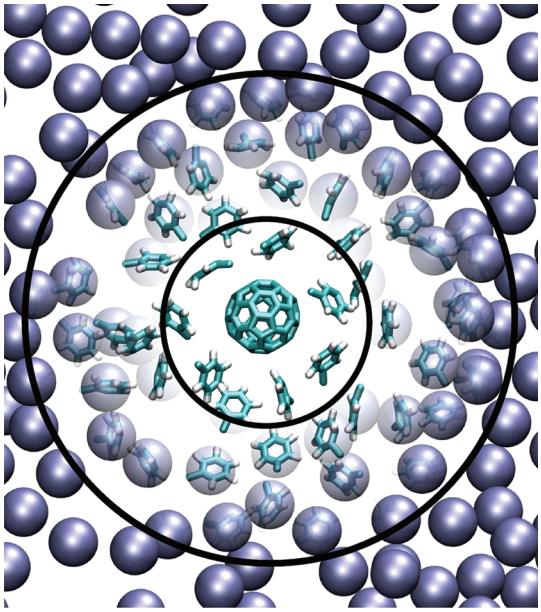

Figure 1. Illustration of the adaptive simulation setup with spherical atomistic zone. The solvent molecules close to the solute are treated with full detail. A force interpolation is employed which smoothly couples the full detail description to an effective center of mass interaction in the bulk. Molecules are constantly exchanged between the regions.

between molecules are treated when molecules diffuse from the all-atom (AA) to the coarse-grained (CG) representation and vice versa:

$$\mathbf{F}_{\alpha\beta}^{\text{total}} = w(r_\alpha)w(r_\beta)\sum_{i\in\alpha}\sum_{j\in\beta}\mathbf{F}_{ij}^{\text{AA}} + [1 - w(r_\alpha)w(r_\beta)]\mathbf{F}_{\alpha\beta}^{\text{CG}} \tag{1}$$

where $r_\alpha = |\mathbf{X}_\alpha - \mathbf{X}_{\text{C60}}|$, $r_\beta = |\mathbf{X}_\beta - \mathbf{X}_{\text{C60}}|$ are the distances of the center of mass of the molecules α and β relative to the C60 center of mass. The all-atom force field $\mathbf{F}_{ij}^{\text{AA}}$ is the force between two atoms. $\mathbf{F}_{\alpha\beta}^{\text{CG}}$ is an effective force field for the center of mass (see Section 2.2 for details), and $w(r)$ is a function of the center of mass of the molecule that goes smoothly from 1 in the all-atom to 0 in the coarse-grained region:

$$w(r) = \begin{cases} 0 & : r > d'_{\text{at}} + d_{\text{hy}} \\ \cos^2\left(\dfrac{\pi}{2d_{\text{hy}}}(r - d'_{\text{at}})\right) & : d'_{\text{at}} + d_{\text{hy}} > r > d'_{\text{at}} \\ 1 & : d'_{\text{at}} > r \end{cases} \tag{2}$$

where $d'_{\text{at}} = d_{\text{at}} + d_{\text{skin}}$ is the size of the all atom and d_{hy} the size of the force-interpolation region. Equation 1 thus smoothly couples the all-atom force field $\mathbf{F}^{\text{AA}}$ to the coarse-grained force-field $\mathbf{F}^{\text{CG}}$ (see refs 5 and 14 for more details on the method).

It is well-known from previous studies that this simple interpolation induces small density artifacts in the hybrid zone. If necessary, these can be corrected by applying an external thermodynamic force which acts in the hybrid region on the centers of mass of the molecules. This force is based on the local differences in chemical potential[15] and can in practice be obtained iteratively from the density profile of the adaptive simulation.[16] Starting from the uncorrected density profile a series of short simulations of pure toluene is performed, where in each step the correction to the thermodynamic force is

$$\mathbf{F}_{\text{th}}^{i+1}(\mathbf{r}) = \mathbf{F}_{\text{th}}^{i}(\mathbf{r}) - \frac{M_{\text{toluene}}}{\rho_0^2\kappa_T^{\text{at}}}\nabla_{\mathbf{r}}\rho_i(|\mathbf{r}|) \tag{3}$$

where ρ_i is the normalized radial density measured from the center of the atomistic zone and κ_T^{at} is the all-atom isothermal compressibility. The thermodynamic force is converged when the density profile is sufficiently flat. In some cases it is necessary to extend the thermodynamic force into the atomistic zone, because up to a cutoff distance away from the hybrid zone, $w(r_\alpha)w(r_\beta) \neq 1$ and so hybrid interactions are present (see eq 1). Thermodynamic force is then applied in the range $d_{\text{at}} < r < d_{\text{at}} + d_{\text{skin}} + d_{\text{hy}}$.

As the knowledge of the coarse-grained force field $\mathbf{F}^{\text{CG}}$ is required for AdResS, we will briefly introduce this approach in Section 2.2.

2.2. Coarse-Grained Model of Liquid Toluene. Using structure-based coarse-graining, a simplified description for molecules can be found.[1−3,17,18] For the present study the gain of coarse graining is two-fold: On the one hand, the simplified description is cheaper computationally, and on the other hand, it also can be seen as a tool to study influence of the lack of detailed chemical features.[19]

To coarse grain, the atomistic toluene interactions were mapped onto one single interaction site in the center of mass of the molecule. The interaction potential was obtained from the iterative Boltzmann inversion method which reproduces the radial distribution function (rdf) $g(r)$ for the centers of mass of the toluenes. The procedure starts from an initial guess for the potential from the distribution in the canonical ensemble:

$$U_0(r) = -\,\mathrm{k_B}T\ln g(r) \tag{4}$$

where $g(r)$ is the solvents center of mass rdf. Several iteration steps consisting of a simulation with the new potential are performed. After each step the potential is updated based on the difference of current and target distributions:

$$U_n(r) = U_{n-1}(r) + k_B T\ln g_{n-1}(r)/g_{\text{target}}(r) \tag{5}$$

In this way the coarse-grained two body potential implicitly includes the multibody effects occurring in the liquid state. The effective coarse-grained potential between the toluene's centers of mass is shown Figure 2. using the VOTCA package, 150 iteration cycles of the potential of 2 ns length were performed.[20] A pressure correction term with target pressure $p = 1$ bar, as in ref 21, was applied during the last 20 cycles. The effective potential was chosen to be 0 beyond 1.32 nm, the position of the second minimum in the rdf.

3. COMPUTATIONAL DETAILS

All simulations were carried out using the GROMACS[22] simulation package. The AdResS algorithm was implemented in GROMACS 4.5 and will be part of GROMACS 4.6.

3.1. Atomistic Simulation. As a reference system, 3987 toluene molecules were simulated in a cubic box of dimensions 8.9 × 8.9 × 8.9 nm. The interaction parameters for toluene

dx.doi.org/10.1021/ct200706f | J. Chem. TheoryComput. 2012, 8, 398−403

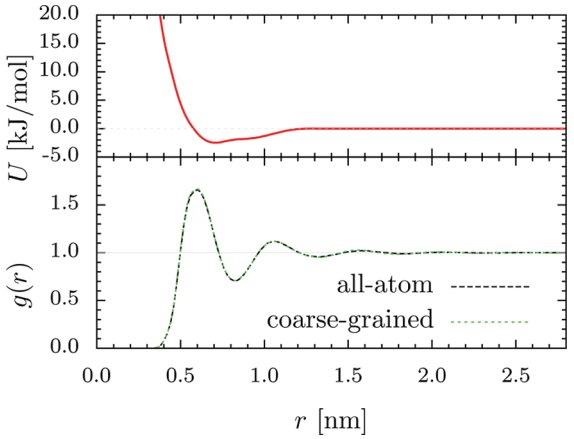

Figure 2. Toluene center of mass rdf in all-atom and coarse-grained representations. The rdfs match within line thickness.

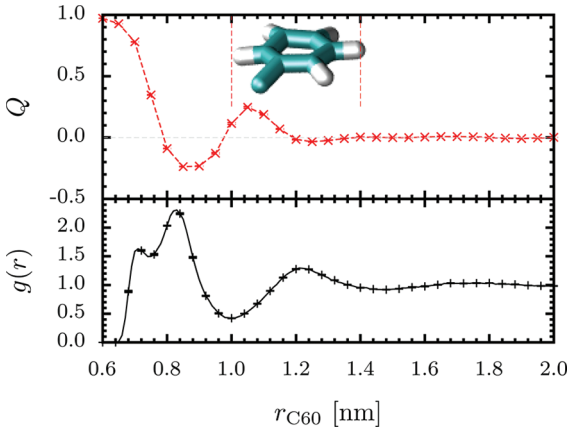

Figure 3. Spherical order parameter Q and spherical center of mass distribution for an atomistic simulation as a function of the relative distance r_{C60} to the center of mass of the C60 molecule. All values are averaged over the simulation trajectory (errorbars are smaller than the line width). The inset shows the toluene molecule in the same length scale.

were taken from the OPLS force field,[23,24] where the CH_3 group is represented by a single interaction site (united-atom model). The C60 force field was taken from Weiss et al.,[25] which uses the same nonbonded interaction parameters for the carbon−carbon interaction as the OPLS force field (σ = 0.355 nm and ε = 0.293 kJ/mol).

Electrostatic interactions needed due to the partially charged atoms in toluene were evaluated using the reaction field[26,27] method with a self-consistent dielectric constant of ε = 1.16735. For thermostatting, a Langevin thermostat[28] with a time constant of τ = 3 ps^{-1} was used. Before the production run in the NVT ensmble, the equilibrium density was determined by 400 ps of equilibration employing a Parrinello−Rahman barostat.[29] This had been determined to be consistent with an independent simulation employing long-range particle mesh Ewald[30] electrostatics.

3.2. Adaptive Simulation. First we studied the influence of the coarse graining of the solvent. For this we fixed the fullerene to its initial position in the middle of the box by using a harmonic restraint potential in the first set of adaptive simulations. With this setup we ran simulations with three different radii of the all atom cavity, namely d_{at} = 0.75 nm, d_{at} = 1.0 nm and d_{at} = 1.5 nm, respectively, which correspond to the density minima of the toluene distribution around C60, shown in Figure 3. The force interpolation region (i.e., the region where $0 < w < 1$) had a fixed size of d_{hy} = 1.4 nm. For the subsequent simulations to study the diffusion constant of the fullerene, the above-mentioned position constraint was released. The thermodynamic force was parametrized in simulations of 1000 pure toluene molecules using 60 iterations (technical details can be found in the manual of the VOTCA package). It can be seen from eq 1 that a molecule in the all-atom zone interacts with hybrid molecules up to the nonbonded cutoff away from the hybrid region. Thus the density distortion originating from the force interpolation extends to beyond the force interpolation region. As a consequence, the thermodynamic force has to be extended by d_{skin} = 0.2 nm. Hence we define, in contrast to previous papers, the all-atom region d_{at} as the region where no thermodynamic force has to be applied ($d_{at} = d'_{at} - d_{skin}$, see eq 2).

4. RESULTS

4.1. Order Parameter. In order to quantify the local order of the solvent around the fullerene, we calculate, similar to the nematic order parameter, an order parameter Q as

$$Q = \left\langle \frac{3}{2}\cos^2(\theta) - \frac{1}{2} \right\rangle \tag{6}$$

where the angle $\cos(\theta) = \hat{\mathbf{n}}\cdot\hat{\mathbf{r}}$ describes the relative orientation of the smallest principal axis $\hat{\mathbf{n}}$ of a toluene molecule pointing away from the center of mass (the smallest principal axis is the eigenvector belonging the smallest eigenvalue of the tensor of inertia) and the direction vector $\hat{\mathbf{r}}$ from the fullerene center of mass to the toluene center of mass. A Q value of 1 indicates ordering of the toluene molecule parallel to the surface of the fullerene, whereas the lowest possible value, Q = −0.5, represents perpendicular orientation. Q = 0 stands for no preferred orientation on average.

Figure 3 clearly shows that the solvent structure is affected by the presence of the fullerene. The shell up to 0.9 nm away from the fullerene consists of a double peak corresponding to orientations preferably perpendicular and orthogonal to the surface. For comparison, the radius of the fullerene is approximately 0.35 nm.

This ordering is also visible in the spherical center of mass distribution around the fullerene (Figure 3 bottom). The first density peak consists of atoms belonging to toluene molecules with preferential parallel ordering, while the second peaks are preferentially perpendicular. The minimum separating the two first peaks is exactly at the location where both orientations cancel on average. At a distance of r = 1 nm from the fullerene, a deep minimum in the density is visible where the solvent molecules form a second layer of parallel orientation. The existence of a minimum at $r \sim 1$ nm is in agreement with experimental findings[11,12] from small-angle X-ray scattering.

4.2. Adaptive Resolution. How well can the adaptive resolutions reproduce the atomistic result? As toluene molecules resemble discs rather than spheres, it is interesting to study how the ordering changes once the surrounding molecules are approximated using spheres. Such a dependency on the chemical details of the surroundings was previously found for water models around fullerenes, when the carbon oxygen interaction was purely repulsive (ideally hydrophobic), while the bulk water structure had almost no influence on the first layers for a very weakly attractive interaction (weakly hydrophobic).[19]

dx.doi.org/10.1021/ct200706f | J. Chem. TheoryComput. 2012, 8, 398−403

To show this, the density distributions of toluene are plotted in Figure 4 for different atomistic zone sizes. For the smallest

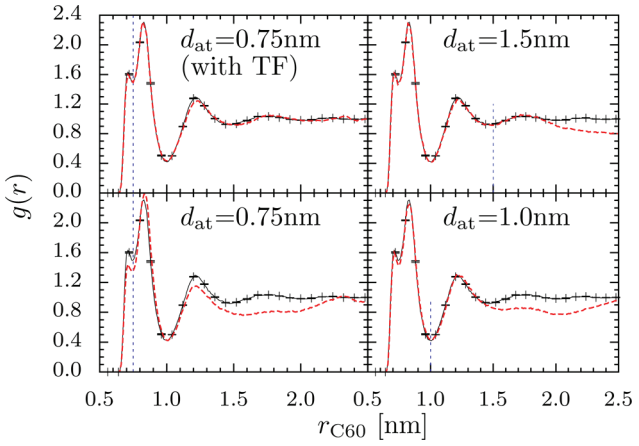

Figure 4. Spherical center of mass distribution relative to the center of mass of the C60. The distribution (red line) is plotted for different sizes of the atomistic region and with the thermodynamic force applied to the hybrid region in comparison to the purely atomistic distribution (black line). The vertical lines indicate the size of the atomistic region.

atomistic zone, the density profile is most perturbed, the deviation at the first two density peaks is clearly visible. The simulations with d_{at} = 1.0 nm and d_{at} = 1.5 nm show no deviations in the atomistic but fail to reproduce the density in the hybrid region, which is an artifact introduced by the force interpolation.[5,15] The simulation employing the thermodynamic force from eq 3 however is perfectly able to reproduce the full density profile even with the smallest atomistic zone d_{at} = 0.75 nm. This result is remarkable as the thermodynamic force was parametrized in pure toluene at equilibrium density (which has negligible near-ordering) and is now used in a situation with strong local density and order variations caused by the presence of the fullerene (see Figures 4 and 5).

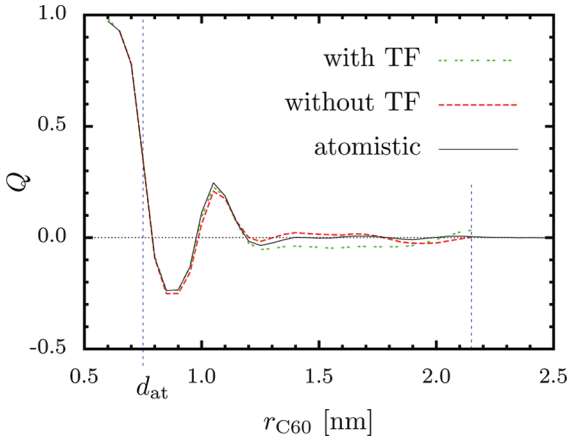

Figure 5. Comparison of order parameter with and without thermodynamic force for an atomistic region with size d_{at} = 0.75 nm. The atomistic value is shown for reference. All values are averaged over the simulation trajectory and normalized to the respective average number of molecules in spherical slabs with a thickness of 0.05 nm. The verticals lines indicate the boundaries of the hybrid region.

For the order parameter describing the orientation of toluene molecules, the situation is slightly different. As can be seen in Figure 5, the order parameter is overall well reproduced by the adaptive simulations. However, in the hybrid zone some weak deviations are visible. The reason is that the hybrid interpolation slightly affects the orientation of the molecules. The thermodynamic force (which radially points away from the AdResS center) neglects orientation of molecules and therefore cannot improve this artifact.

4.3. Freely Diffusing Atomistic Zone. It is known that in the case of the coarse-grained representation, the diffusion is increased in comparison to the atomistic representation due a smoother energy landscape.[21] Thus a natural question occurring in the context of adaptive resolution is how this property is affected by the approximate treatment of the solvent degrees of freedom in the coarse-grained representation. We study this by removing the restraints on the center of mass position of the fullerene and thus letting the atomistic zone diffuse through the liquid along with the molecule.

Before measuring the diffusion properties, we choose a value of the relaxation time of the Langevin thermostat (τ = 3 ps^{-1}) such that weak coupling can be assumed.

The results are shown in Figure 6 and Table 1. In the diffusive regime, the fullerene in adaptive resolution diffuses faster

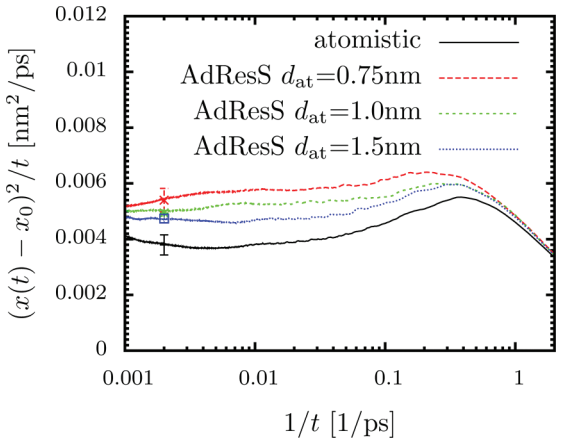

Figure 6. Diffusion behavior of the fullerene simulated atomistically and in adaptive resolutions using different atomistic zone sizes (d_{at}).

Table 1. Diffusion Constants for the C60/Toluene Center of Mass

C_{60}	diffusion constants [nm^2/ps]
all atom	$D = 6.3 \times 10^{-4} \pm 6 \times 10^{-5}$
AdResS	
d_{at} = 0.75 nm	$D = 9.0 \times 10^{-4} \pm 7 \times 10^{-5}$
d_{at} = 1.0 nm	$D = 8.2 \times 10^{-4} \pm 2 \times 10^{-5}$
d_{at} = 1.5 nm	$D = 7.9 \times 10^{-4} \pm 3 \times 10^{-5}$
AdResS with TF	
d_{at} = 0.75 nm	$D = 7.3 \times 10^{-4} \pm 8 \times 10^{-5}$
Pure toluene	
all atom	$D = 2.216 \times 10^{-3} \pm 3 \times 10^{-6}$
coarse grained	$D = 8.629 \times 10^{-3} \pm 1 \times 10^{-5}$

compared to the full all-atom simulation. With increasing atomistic zone size, the diffusion constant decreases again and approaches the all-atom values. This is in consistency with the findings in pure coarse-grained simulations:[21] As degrees of freedom are decreased, the energy landscape is less rugged which leads to increased diffusion (a factor of ≈4 was found for our toluene model). Thus, for the smallest atomistic zone,

dx.doi.org/10.1021/ct200706f | J. Chem. TheoryComput. 2012, 8, 398–403

diffusion is fastest as the least resistance from the surroundings can be expected. Table 1 compares the diffusion constant for the simulation done with thermodynamic force. Without thermodynamic force the diffusion constant is slightly higher, which might be attributed to the artifacts (decreased density) in the hybrid region.

For most practical applications this results means that a reasonably small atomistic/hybrid zone can already reproduce the diffusion with good accuracy. The consequence is that for studying aggregation phenomena, each solute only needs a small atomistic region surrounding it, while the bulk solvent can be treated in computationally inexpensive coarse-grained fashion.

4.4. Computational Performance. Since one goal of adaptive resolutions methods is to reduce the heavy computational costs associated with full details atomistic simulations, we compare the computational performance of our algorithm in Figure 7.

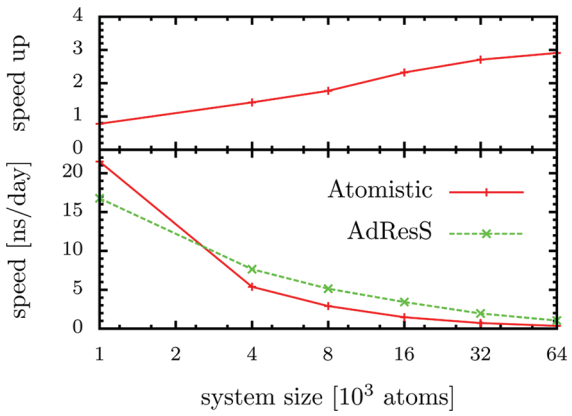

Figure 7. Computational scaling as function of total system size. For the adaptive simulation, the volume of the atomistic zone is kept fixed.

For small systems the adaptive simulation is slower than a conventional simulation as the additional costs from calculating the hybrid interactions outweighs the gain of calculating the coarse-grained interactions. However, if the major part of the system consists of coarse-grained particles, a total speed up of up to a factor ~3 was observed for our current GROMACS implementation. Note that this speed up can be improved in further implementations: The algorithm used for the hybrid interactions is not as heavily optimized as the standard algorithm (written in assembly) implemented in GROMACS. Furthermore, besides the computational speed up, there is a speedup in converging atomistic averages as the reduced degrees of freedom in the coarse-grained regime relax faster. As the coarse-grained potential are much smoother than the atomistic ones, it would be possible to increase the time step in the coarse-grained region and combine it with multiple time-step techniques.

5. CONCLUSIONS

In this paper we studied a C60 fullerene solvated in toluene using the AdResS, which couples a coarse-grained and an atomistic molecular dynamics simulation based on force interpolation. In agreement with experiments, we find from the all-atom simulation a structured layer of about 1 nm in size around the fullerene. For measuring the local ordering, we have defined an order parameter, which takes values between −0.5 and 1 for perpendicular and parallel direction of the toluene to the normal of the C60 surface.

In contrast to a common molecular dynamics simulation where the chemical details can not easily be switched on or off locally, this is possible in the AdResS by varying the size of the atomistic region. Outside the atomistic region, the detailed chemical structure of toluene is replaced by an effective description neglecting the disk-like shape. Using the AdResS method, we determined that the structured layer of toluene around the fullerene is local and needs no support of the chemical structure, the disk-like shape, from the surrounding bulk.

We also measured the diffusion of the C60 for different sizes of the atomistic zone and observed that the C60 diffuses faster for smaller atomistic regions. This is very natural because of the smoother and hence less hindering energy landscape of the coarse-grained region. The all-atom diffusion constant is however reproduced within 20% for a relatively small atomistic zone of d_{at} = 1.5 nm.

Density artifacts in the interpolation region can be corrected by a 'thermodynamic force'. In the current form, the thermodynamic force neglects orientational properties and thus had to be modified in order to correct the small deviations in the interpolated order parameter.

Regarding the computational performance, we have shown that the AdResS leads to a speed-up of approximately a factor 3 for bigger systems. This can still be improved by more optimized kernels. The code used for this publication is freely available from the GROMACS Web site and will be part of GROMACS 4.6.

6. IMPLEMENTATION DETAILS

The implementation of the method poses some technical challenges which have been explained in ref 14. We follow this approach and implement AdResS in a similar fashion, the atomistic details are not removed from the memory once they enter the coarse-grained domain and therefore are still integrated along with the center of mass movement of a molecule. This simplifies the implementation and avoids potential overhead caused by the memory allocation in case of re-entering of a molecule into the atomistic region. The coarse-grained and atomistic representation of a molecule is connected via the virtual site construction. This means the position of the coarse-grained representation is geometrically constructed from the positions of the atomistic ones and forces acting on the coarse-grained representation are distributed to the atomistic representation on which the integrator acts. The region of the force interpolation can be defined by six new parameters:

- `address_type` defines the geometry of the atomistic zone, possibilities are `sphere`, `Xsplit`, `Const`, or `Off`.
- `address_ex_width` defines the width of the atomistic zone, d'_{at} of eq 2.
- `address_hy_width` defines the width of the hybrid zone, d_{hy} of eq 2.
- `address_interface_correction` defines the type of correction to be applied in the hybrid zone, possibilities are `thermoforce` and `no`.
- `address_reference_coords` defines the center of the atomistic zone.
- `address_cg_grp_names` defines which energy group(s) should be treated as coarse-grained.

Nevertheless the adaptive algorithm saves computer time (see Figure 7) with this strategy, as atoms with a weighting factor $w = 0$ can be excluded beforehand from the force calculation loop that usually takes up to 70% of the computer time. In GROMACS, all pairwise nonbonded interactions are computed in separate loops (kernels) for each type of interaction (e.g., Coulomb, van der Waals, combined, etc.). For each of these kernels exists implementations optimized for a specific hardware. We avoid the necessity to rewrite all kernels by implementing AdResS kernels only in C, which can be called independently of the underlying hardware, which however leads to a slower kernel. In the case of a parallel domain decomposition scheme, some computer nodes only calculate purely coarse-grained or all-atom interactions, which then can use the highly optimized kernels. In the domains where both interaction types occur, the nodes calculate the more expensive hybrid region using our hybrid 'AdResS-kernels'. This has been tested to also work with GROMACS dynamic load balancing explained in ref 22, which will lead to big domains for the coarse-grained region and small ones for the hybrid region.

Now we will discuss the reasoning for the computational speedup of AdResS. This does not involve any physical speed-ups, like faster diffusion or quicker equilibration. The total simulation time $t_{\text{all-atom}}$ consists out the time needed to calculate the forces t_{force} and the rest t_{rest}. As in our implementation only affect the force, calculating t_{rest} stays unchanged, which already defined the upper limit of the speedup, $t_{\text{all-atom}}/t_{\text{rest}}$. The approach is still reasonable, as the t_{force} is usually the biggest part of $t_{\text{all-atom}}$. As t_{rest} mostly contains routines, like neighbor search, I/O, and communication, which scale with the degree of coarse grained, e.g., the number of removed atoms per molecule, N_{mol}, a different implementation could still improve the speedup of this part. The t_{force} part scales down with a factor of N_{mol}^2, as all intermolecular interaction get replaced by one effective interaction. It is illustrative to calculate an "ideal speedup" that would be possible with this method and compare it to the measured values. For the case of the big system containing 64 000 molecules, 74.2% of the time is spent for the force calculation in the standard all-atom simulation. Assuming that the volume of the atomistic and hybrid region is small compared to the coarse grained (0.5% of the volume is atomistic or hybrid), the ideal speedup would be $k = (t_{\text{all-atom}})/(t_{\text{AdResS}}) = (t_{\text{force}} + t_{\text{rest}})/(t_{\text{rest}}) \approx 4$. This compares well to the speed-up of ≈3 measured in our simulations. The second biggest part of the simulation time (9.2%) was spent for neighbor search. However, some of the gains get compensated by the costs for the additional coarse-grained interaction in the hybrid zone and for calculating the center of mass and redistributing forces due to the virtual site construction (8.2% of the simulation time in the AdResS simulations). An MD code tailored for the needs of the AdResS simulation (currently under development, Espresso++) could thus reach even larger speed-ups.

AUTHOR INFORMATION

Corresponding Author

*E-mail: fritsch@mpip-mainz.mpg.de.

Present Address

‡Theory Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

We acknowledge Luigi Delle Site for fruitful discussions during the entire project as well as Victor Rühle and Christine Peter for critical reading of the manuscript. C.J. was financially supported by SFB 625 and the multiscale modeling initiative of the Max-Planck Society (M^3). Berk Hess is acknowledged for helping us understanding the insides of GROMACS. Brad Lambeth is acknowledged for bravely testing the first implementation.

REFERENCES

(1) *Coarse-Graining of Condensed Phase and Biomolecular Systems*; Voth, G. A., Ed.; Taylor and Francis: Oxford, U.K., 2008; pp 21−40.
(2) Peter, C.; Kremer, K. *Soft Matter* **2009**, *5*, 4357.
(3) Peter, C.; Kremer, K. *Faraday Discuss.* **2010**, 9−24.
(4) *Monte Carlo and Molecular Dynamics Simulations in Polymer Science*; Binder, K., Ed.; Oxford University Press: Oxford, U.K., 1995.
(5) Praprotnik, M.; Site, L. D.; Kremer, K. *J. Chem. Phys.* **2005**, *123*, 224106−14.
(6) Praprotnik, M.; Delle Site, L.; Kremer, K. *Annu. Rev. Phys. Chem.* **2008**, *59*, 545−571.
(7) Wagoner, J. A.; Pande, V. S. *J. Chem. Phys.* **2011**, *134*, 214103.
(8) Kroto, H. W.; Heath, J. R.; O'Brien, S. C.; Curl, R. F.; Smalley, R. E. *Nature* **1985**, *318*, 162−163.
(9) Wynands, D.; Levichkova, M.; Leo, K.; Uhrich, C.; Schwartz, G.; Hildebrandt, D.; Pfeiffer, M.; Riede, M. *Appl. Phys. Lett.* **2010**, *97*, 073503.
(10) *Nanotechnology in Drug Delivery*; Villiers, M. M., Aramwit, P., Kwon, G. S., Eds.; Springer: New York, 2009; pp 349−384.
(11) Wei, T.; Huang, T.; Wu, T.; Tsai, P.; Lin, M. *Chem. Phys. Lett.* **2000**, *318*, 53−57.
(12) Ginzburg, B. M.; Tuĭchiev, S.; Tabarov, S. K.; Shepelevskiĭ, A. A. *Crystallogr. Rep.* **2005**, *50*, 735−738.
(13) Poma, A. B.; Delle Site, L. *Phys. Rev. Lett.* **2010**, *104*, 250201.
(14) Junghans, C.; Poblete, S. *Comput. Phys. Commun.* **2010**, *181*, 1449−1454.
(15) Poblete, S.; Praprotnik, M.; Kremer, K.; Delle Site, L. *J. Chem. Phys.* **2010**, *132*, 114101.
(16) Fritsch, S.; Poblete, S.; Junghans, C.; Ciccotti, G.; Delle Site, L.; Kremer, K. *arXiv:1112.31512011*
(17) Tschöp, W.; Kremer, K.; Batoulis, J.; Burger, T.; Hahn, O. *Acta Polym.* **1998**, *49*, 61−74.
(18) Reith, D.; Pütz, M.; Müller-Plathe, F. *J. Comput. Chem.* **2003**, *24*, 1624−1636.
(19) Lambeth, B. P.; Junghans, C.; Kremer, K.; Clementi, C.; Delle Site, L. *J. Chem. Phys.* **2010**, *133*, 221101.
(20) Rühle, V.; Junghans, C.; Lukyanov, A.; Kremer, K.; Andrienko, D. *J. Chem. Theory Comput.* **2009**, *5*, 3211−3223.
(21) Wang, H.; Junghans, C.; Kremer, K. *Eur. Phys. J. E* **2009**, *28*, 221−229.
(22) Hess, B.; Kutzner, C.; van der Spoel, D.; Lindahl, E. *J. Chem. Theory Comput.* **2008**, *4*, 435−447.
(23) Jorgensen, W. L.; Nguyen, T. B. *J. Comput. Chem.* **1993**, *14*, 195−205.
(24) Jorgensen, W. L.; Laird, E. R.; Nguyen, T. B.; Tirado-Rives, J. *J. Comput. Chem.* **1993**, *14*, 206−215.
(25) Weiss, D. R.; Raschke, T. M.; Levitt, M. *J. Phys. Chem. B* **2008**, *112*, 2981−2990.
(26) Neumann, M. *Mol. Phys.* **1983**, *50*, 841.
(27) Neumann, M. *J. Chem. Phys.* **1985**, *82*, 5663.
(28) Van Gunsteren, W. F.; Berendsen, H. J. C. *Mol. Sim.* **1988**, *1*, 173.
(29) Parrinello, M.; Rahman, A. *J. Appl. Phys.* **1981**, *52*, 7182.
(30) Essmann, U.; Perera, L.; Berkowitz, M.; Darden, T.; Lee, H.; Pedersen, L. *J. Chem. Phys.* **1995**, *103*, 8577−8593.