Computer Physics Communications 143 (2002) 162–173

Computer Physics Communications

www.elsevier.com/locate/cpc

High-dimensional quantum dynamics of molecules on surfaces: a massively parallel implementation

S. Borowski [a], S. Thiel [a], T. Klüner [a,*], H.-J. Freund [a], R. Tisma [b], H. Lederer [b]

[a] *Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, 14195 Berlin, Germany*
[b] *Rechenzentrum der Max-Planck-Gesellschaft Garching, Max-Planck-Institut für Plasmaphysik, Boltzmannstraße 2, 85748 Garching, Germany*

Received 24 August 2001

Abstract

We present a massively parallel implementation to perform quantum dynamical wave packet calculations of molecules on surfaces. The employed algorithm propagates the wavefunction via the time-dependent Schrödinger equation within a finite basis representation by Split and Chebyshev schemes, respectively. For the parallelization, a problem adapted data decomposition in all dimensions is introduced that ensures an optimal load balancing. In a speedup analysis of the timing and scaling properties, the overall semi-linear scaling of the algorithm is verified. The almost linear speedup up to 512 processing elements indicates our implementation as a powerful tool for high-dimensional calculations. The implementation is applied to laser induced desorption of molecules from surfaces.

PACS: 02.60Cb; 79.20La

Keywords: Quantum dynamics; Wave packet; Parallel computing; Photodesorption

1. Introduction

A rigorous treatment of quantum dynamical phenomena requires the solution of the time-dependent Schrödinger equation. While low-dimensional quantum dynamics can be performed routinely on workstations, high-dimensional calculations still remain a challenge. This is due to the exponential scaling behaviour with the number of degrees of freedom f taken into account. For a general basis representation of the extent N, the overall problem size amounts to N^f. Then, the action of the Hamiltonian upon the wavefunction shows the scaling relation N^{2f} of a matrix–vector multiplication. An iterative action of the Hamiltonian, which is the most time consuming step of many algorithms, results in an exhausting numerical effort. Therefore, accurate quantum dynamical studies are restricted to few dimensions only.

In the last two decades progressive developments in algorithms, propagation methods and numerical implementations have yielded much more economical approaches. Fast transform algorithms [1] in conjunction

* Corresponding author.
E-mail address: kluener@fhi-berlin.mpg.de (T. Klüner).

0010-4655/02/$ – see front matter
PII: S0010-4655(01)00445-3

with coordinate representations and basis expansions tailored to the specific problem [2–4] minimize the problem size and approach linear scaling with the problem size for the action of the Hamiltonian. Very efficient propagation schemes reduce the number of necessary Hamiltonian actions [5] in time-dependent approaches. However, in spite of the recent methodical progress, contemporary applications are only able to tackle three-dimensional exact quantum dynamics [6–8]. For hydrogen studies, the small spectral range of the Hamiltonian due to lightweight system requires only a small-sized basis expansion that allows a six-dimensional investigation [9,10]. By limiting the configuration space, approximative approaches are able to treat more degrees of freedom [11–14]. Because of the limitations, these methods have to refer to exact quantum dynamical studies as a benchmark.

Computing power and memory requirements of such high-dimensional studies exceed the capabilities of single workstations by far. Consequently, a great need for efficient parallel implementations has raised over the past years. Characterized by highly scalable algorithms and a large range of problem size, quantum dynamical studies are well suited for exploitation of parallel computing power. However, recent efforts [15–18] are mainly focused on diagonalization of the Hamiltonian within a time-independent approach. The employed problem decompositions in one or few dimensions are not very flexible resulting in restrictions for the problem shape and finally for the problem size.

The present paper introduces a parallel algorithm for the quantum dynamical wave packet propagation of molecules on surfaces. The paper is organized as follows: In Section 2, we introduce our physical model. The basis representation, the computation of the resulting Hamiltonian matrix elements and the employed propagation schemes are described. Section 3 presents the parallel implementation. The proposed strategy of data decomposition and communication is discussed in a speedup analysis. Section 4 shows the first results from the application of our implementation to a photodesorption process. Finally, we conclude in Section 5.

2. Physical model

2.1. Representation

The nuclear quantum dynamics of a rigid diatomic molecule on a corrugated static surface includes five degrees of freedom as shown in Fig. 1. In this coordinate representation the corresponding Hamiltonian may be written as

$$\widehat{H} = \widehat{T} + \widehat{V} = \frac{\widehat{K}^2}{2M} + \frac{\hat{J}^2}{2\mu r_0^2} + \widehat{V}(\mathbf{R}, \theta, \phi) \tag{1}$$

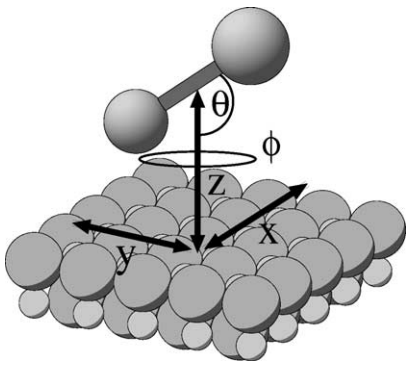

Fig. 1. Coordinate representation of a rigid diatomic molecule adsorbed on a static surface.

with the position $\mathbf{R} = (x, y, z)$ and the momentum operator $\widehat{\mathbf{K}} = (\hat{k}_x, \hat{k}_y, \hat{k}_z)$ of the center of mass, the total mass M, the angular momentum operator $\hat{\mathbf{J}}$, the reduced mass μ and the intramolecular equilibrium distance r_0. The kinetic energy operator

$$\widehat{T} = \frac{\widehat{K}^2}{2M} + \frac{\hat{J}^2}{2\mu r_0^2} = -\frac{\nabla_{\mathbf{R}}^2}{2M} - \frac{1}{2\mu r_0^2}\left(\frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\sin\theta\frac{\partial}{\partial\theta} + \frac{1}{\sin^2\theta}\frac{\partial^2}{\partial\phi^2}\right) \tag{2}$$

is composed of the contributions from the translation of the center of mass M and the rotation of the reduced mass μ. The potential energy surface $\widehat{V}(\mathbf{R}, \theta, \phi)$ is obtained from quantum chemical calculations solving the electronic Schrödinger equation within the Born–Oppenheimer approximation.

According to the structure of the Hamiltonian, we expand the wavefunction

$$\Psi(\mathbf{R}, \theta, \phi, t) = \sum_{ijm} \Psi_{ijm}(t)\Phi_i(\mathbf{R})Y_{jm}(\theta, \phi) \tag{3}$$

in terms of orthogonal eigenfunctions of the kinetic energy operator $\widehat{T}$ with the expansion coefficients Ψ_{ijm}. This finite basis representation (FBR) [19,20] consists of an expansion in plane waves $\{\Phi_i\}$ for the Cartesian coordinates and spherical harmonics $\{Y_{jm}\}$ for the angular coordinates, respectively. The plane wave basis includes $N_{xyz} = N_x N_y N_z$ basis functions with the momenta $\mathbf{K}_i = (k_x^{i_x}, k_y^{i_y}, k_z^{i_z})$ that form an equidistant grid in the Cartesian momentum space. The rotational basis is limited to $N_{\mathrm{rot}} = (j_{\max} + 1)^2$ basis functions with the angular momentum $0 \leqslant j \leqslant j_{\max}$ and its projection $-j_{\max} \leqslant m \leqslant j_{\max}$.

2.2. Matrix elements of the Hamiltonian

The basic operation of any propagation scheme is the action of the Hamiltonian upon the wavefunction [5]. Within a general basis expansion, the Hamiltonian and the wavefunction take the shape of a matrix and a vector, respectively. Accordingly, the Hamiltonian action is a matrix–vector multiplication that scales quadratically with the problem size. This unfavorable scaling behaviour is significantly improved by the special structure of the Hamiltonian matrix elements within the FBR as described in the following paragraphs.

To compute matrix elements of the Hamiltonian, we employ a quadrature scheme on a grid representation that follows the discrete variable representation (DVR) [2]. The Cartesian coordinates $\mathbf{R}$ and the azimuthal angle ϕ are sampled on evenly spaced grids of N_{xyz} and $N_\phi = 2j_{\max} + 1$ grid points, respectively. The polar angle θ is represented by $N_\theta = j_{\max} + 1$ grid points defined by the zeros of the Legendre polynomial $P_{j_{\max}+1}(\theta)$. Corresponding to the grid representation, the matrix elements of the Hamiltonian are integrated over $\mathbf{R}$ and ϕ by the trapezoidal rule whereas the integration over θ is performed by a Gauss–Legendre quadrature with the weight function w [21]. Within this quadrature scheme, the orthogonality of the FBR basis functions

$$\sum_{\alpha\beta\gamma} w_\beta \Phi_i^*(\mathbf{R}_\alpha)Y_{jm}^*(\theta_\beta, \phi_\gamma)\Phi_{i'}(\mathbf{R}_\alpha)Y_{j'm'}(\theta_\beta, \phi_\gamma) = \delta_{ii'}\delta_{jj'}\delta_{mm'} \tag{4}$$

is preserved. Thus, the kinetic energy operator has a diagonal matrix representation and its action is performed by a vector–vector multiplication of the wavefunction by the eigenvalue spectrum [20]. This operation features the linear scaling relation $N_{xyz}N_{\mathrm{rot}}$. As the matrix elements of the potential energy operator

$$V_{ijm,i'j'm'} = \sum_{\alpha\beta\gamma} w_\beta \Phi_i^*(\mathbf{R}_\alpha)Y_{jm}^*(\theta_\beta, \phi_\gamma)V(\mathbf{R}_\alpha, \theta_\beta, \phi_\gamma)\Phi_{i'}(\mathbf{R}_\alpha)Y_{j'm'}(\theta_\beta, \phi_\gamma) \tag{5}$$

are off-diagonal, their straight computation is avoided. Instead, the wavefunction is transformed to the DVR by Fourier transforms in $\mathbf{K}$, a Gauss–Legendre transform in j and a Fourier transform in m. Within the DVR, the action of the potential energy operator upon the wavefunction is performed as a vector–vector multiplication with the linear scaling relation $N_{xyz}N_\theta N_\phi$. Then, the wavefunction is restored to the FBR by the corresponding inverse

transformations [20]. By using a Fast Fourier Transform (FFT) algorithm, the Fourier transforms have the semi-linear scaling relation $N_{xyz}\log(N_{xyz})N_{\mathrm{rot}}$ in $\mathbf{K}$ and $N_{xyz}N_\theta N_\phi \log N_\phi$ in m, respectively. The Gauss–Legendre transform shows the quadratic scaling relation $N_{xyz}N_{\mathrm{rot}}N_\theta$ of a matrix–vector multiplication in j. In practical applications, the semi-linear scaling in $\mathbf{K}$ and m dominates the quadratic scaling in j due to the relatively small extent of the rotational basis (shown in Section 3). Finally, the action of the Hamiltonian upon the wavefunction shows an overall semi-linear scaling within the FBR.

However, this DVR quadrature introduces an inhomogeneous data processing: since the fixed grid in the azimuthal angle ϕ has to resolve the largest projection of the angular momentum, the DVR consists of approximately twice as many grid points $N_\theta N_\phi = (j_{\max}+1)(2j_{\max}+1) \approx 2(j_{\max}+1)^2$ as FBR basis functions $N_{\mathrm{rot}} = (j_{\max}+1)^2$ [2]. The difficulties concerning this issue are addressed in Section 3.

2.3. Dynamics

The time evolution of a quantum mechanical system is described by the time-dependent Schrödinger equation

$$\mathrm{i}\frac{\partial}{\partial t}\Psi(t) = \widehat{H}\Psi(t), \tag{6}$$

which can formally be solved

$$\Psi(t+\mathrm{d}t) = \widehat{U}(\mathrm{d}t)\Psi(t) \quad \text{with } \widehat{U}(t') = \exp(-\mathrm{i}\widehat{H}t') \tag{7}$$

by introducing the time evolution operator $\widehat{U}(t)$. For time-dependent systems, the time evolution operator is restricted to short time steps in order to sample the time dependence of the Hamiltonian. The time-independent Hamiltonian in Eq. (1) enables both short-time and global propagation methods because of its trivial time dependence. Thus, we implemented the Split propagator [22]

$$\widehat{U}_{\mathrm{split}}(\Delta t) \approx \exp\left(-\frac{\mathrm{i}}{2}\widehat{T}\Delta t\right)\exp(-\mathrm{i}\widehat{V}\Delta t)\exp\left(-\frac{\mathrm{i}}{2}\widehat{T}\Delta t\right) \tag{8}$$

for a very efficient short-time propagation providing second-order accuracy and the additional option of time-dependent external fields. Furthermore, we realized an implementation of the Chebyshev propagator [23] for a global and accurate propagation by expanding the evolution operator in terms of Chebyshev polynomials

$$\widehat{U}_{\mathrm{cheb}}(t) \approx \sum_n a_n(t)T_n(\widehat{H}_{\mathrm{norm}}) \tag{9}$$

with the renormalized Hamiltonian $\widehat{H}_{\mathrm{norm}}$ of the spectral range $[-1,1]$. The resulting expansion coefficients $a_n(t)$ are given by Bessel functions of the first kind. The Chebyshev expansion is characterized by exponential convergence and linear scaling in the time domain.

Because of its small computational effort and memory requirement, the fast Split propagator is employed for high-dimensional calculations. Scrutinizing the obtained results, the Split scheme is carefully compared to the numerically exact Chebyshev scheme for certain quantum trajectories.

3. Parallelization

The high-dimensional quantum dynamics presented in Section 2 requires an enormous amount of CPU time and memory. According to these demands, even parallelization techniques on shared memory vector systems (OpenMP, multitasking libraries) do not provide the needed computing power. Only message passing concepts (MPI, PVM, SHMEM) utilizing massively parallel distributed memory architectures prove to be sufficient. Hence, we realized a massively parallel FORTRAN 90 implementation with the SHMEM library, which is superior to comparable message passing libraries in efficiency, especially latency, and ease of usage [24]. Test calculations were made on

the SGI Origin 3400 at the Fritz-Haber-Institut Berlin. Full calculations were performed on the Cray T3E at the Rechenzentrum Garching.

In the following consideration of our parallelization strategy, we freeze one lateral coordinate y since an accurate investigation of the remaining 4D problem is even limited by the resources of currently available supercomputers. Moreover, the computation of ab initio potential energy surfaces for four-dimensional systems is a challenging problem itself. However, the generalization of the concept to more than four dimensions is straightforward.

3.1. Parallelization strategy

The analysis of our physical model with respect to inherent parallelism reveals that only the transformations of the DVR quadrature arrange the data coupling and therefore involve global data. Consequently, the straight parallel realization is a data decomposition resulting in communication collecting and distributing the global data before and after the transformations. Furthermore, reduction communication [24] is required for the computation of observables. According to the symmetry between data distribution and data processing within this scope, a well balanced data decomposition leads to optimal load balancing and minimal communication overhead among the different processing elements (PEs).

[1]To obtain a general implementation for large scale applications, we choose a data distribution among $N^{\mathrm{PE}} = N_x^{\mathrm{PE}} \times N_z^{\mathrm{PE}} \times N_\theta^{\mathrm{PE}} \times N_\phi^{\mathrm{PE}}$ PEs in all degrees of freedom as shown in Fig. 2. Since the Cartesian coordinates have the same data structure within FBR and DVR, a simple block distribution already provides an optimal load balancing (see left panel of Fig. 2). For the distribution among $N_{\{x,z\}}^{\mathrm{PE}}$ PEs, the global dimension of the extent $N_{\{x,z\}}$ is evenly divided into local dimensions of the extent $N_{\{x,z\}}^{\mathrm{L}} = N_{\{x,z\}}/N_{\{x,z\}}^{\mathrm{PE}}$. The relation between the global index $i_{\{x,z\}}$ and the local index $i_{\{x,z\}}^{\mathrm{L}}$ on PE $n_{\{x,z\}}^{\mathrm{PE}}$ [2] is given by

$$
\begin{aligned}
i_{\{x,z\}}^{\mathrm{L}} &= \left(i_{\{x,z\}} - 1\right) \bmod N_{\{x,z\}}^{\mathrm{L}} + 1, \\
n_{\{x,z\}}^{\mathrm{PE}} &= \frac{i_{\{x,z\}} - 1}{N_{\{x,z\}}^{\mathrm{L}}}, \\
i_{\{x,z\}} &= i_{\{x,z\}}^{\mathrm{L}} + n_{\{x,z\}}^{\mathrm{PE}} \cdot N_{\{x,z\}}^{\mathrm{L}}.
\end{aligned}
\tag{10}
$$

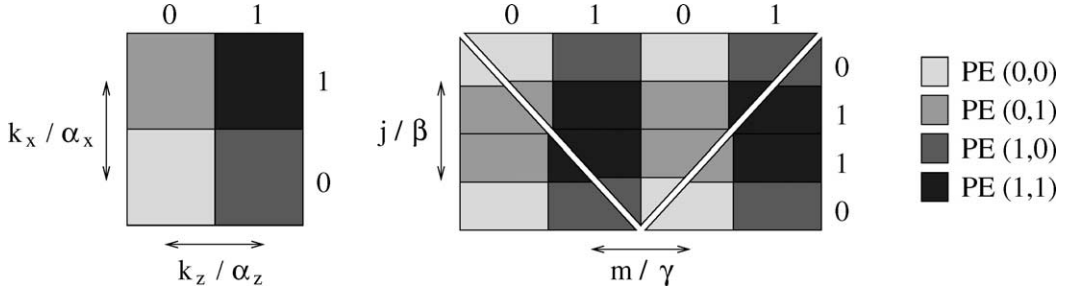

Fig. 2. Equivalence of load and data distribution: 2×2 PE grid for the Cartesian coordinates $\{x, z\}$ (left panel) and the angular coordinates $\{\theta, \phi\}$ (right panel), respectively. Data structures located on a certain PE are labeled by colored areas. The FBR data $\{k_x, k_z, j, m\}$ is limited to the central $\{j, m\}$ triangle only whereas the DVR data $\{\alpha_x, \alpha_z, \beta, \gamma\}$ includes the whole $\{\beta, \gamma\}$ rectangle. The color scheme indicates a perfect load balancing for all dimensions, i.e., all PEs process the same amount of data in both the FBR and the DVR.

[1] In the following paragraphs, the basic arithmetics concerning the data decomposition has to be understood as integer operations.

[2] Global indices denote the element position with respect to the global dimension whereas local indices denote the element position with respect to the local dimension on a certain PE.

For the angular coordinates, a more sophisticated data decomposition is necessary. The data shape of the rotational basis within the FBR is the central $\{j, m\}$ triangle in the right panel of Fig. 2, which is characterized by the multiplicity of the j states with respect to their projection m. However, the data object of the angular grid within the DVR is the whole $\{\beta, \gamma\}$ rectangle due to the fixed high resolution grid in the azimuthal angle ϕ. Hence, we developed a data distribution that ensures optimal load balancing in both the FBR and the DVR data (see right panel of Fig. 2). For the polar angle θ, the global dimension of the extent N_θ is evenly distributed among N_θ^{PE} PEs into local dimensions of the extent $N_\theta^{\mathrm{L}} = N_\theta / N_\theta^{\mathrm{PE}}$. The lower and upper half of the global dimension is pairwisely placed in ascending and descending order into the local dimensions, respectively. For simplicity, we arrange the resulting element pairs into consecutive odd and even local positions. Then, the mapping between the global index j and the local index j^{L} on PE n_θ^{PE}

$$
\begin{aligned}
&j \leqslant N_\theta/2\text{:} && j^{\mathrm{L}} = 2\cdot(j-1) \bmod\bigl(N_\theta^{\mathrm{L}}/2\bigr) + 1, \\
& && n_\theta^{\mathrm{PE}} = \frac{j-1}{N_\theta^{\mathrm{L}}/2}, \\
&j > N_\theta/2\text{:} && j^{\mathrm{L}} = 2\cdot(N_\theta - j) \bmod\bigl(N_\theta^{\mathrm{L}}/2\bigr) + 2, \\
& && n_\theta^{\mathrm{PE}} = \frac{N_\theta - j}{N_\theta^{\mathrm{L}}/2}, \\
&j^{\mathrm{L}} \text{ odd:} && j = j^{\mathrm{L}}/2 + n_\theta^{\mathrm{PE}} \cdot \bigl(N_\theta^{\mathrm{L}}/2\bigr) + 1, \\
&j^{\mathrm{L}} \text{ even:} && j = N_\theta - \bigl(j^{\mathrm{L}}/2 + n_\theta^{\mathrm{PE}} \cdot \bigl(N_\theta^{\mathrm{L}}/2\bigr)\bigr) + 1
\end{aligned}
\tag{11}
$$

requires N_θ^{L} to be even. Similarly, for the azimuthal angle ϕ the global dimension of the extent N_ϕ is evenly divided among N_ϕ^{PE} PEs into local dimensions of the extent $N_\phi^{\mathrm{L}} = N_\phi / N_\phi^{\mathrm{PE}}$. The lower and upper half of the global dimension is pairwisely placed in ascending order into the local dimensions. Again, the element pairs are simply arranged into consecutive odd and even local positions. The resulting transformation between the global index m and the local index m^{L} on PE n_ϕ^{PE}

$$
\begin{aligned}
&m \leqslant N_\phi/2\text{:} && m^{\mathrm{L}} = 2\cdot(m-1) \bmod\bigl(N_\phi^{\mathrm{L}}/2\bigr) + 1, \\
& && n_\phi^{\mathrm{PE}} = \frac{m-1}{N_\phi^{\mathrm{L}}/2}, \\
&m > N_\phi/2\text{:} && m^{\mathrm{L}} = 2\cdot(m - N_\phi/2 - 1) \bmod\bigl(N_\phi^{\mathrm{L}}/2\bigr) + 2, \\
& && n_\phi^{\mathrm{PE}} = \frac{m - N_\phi/2 - 1}{N_\phi^{\mathrm{L}}/2}, \\
&m^{\mathrm{L}} \text{ odd:} && m = m^{\mathrm{L}}/2 + n_\phi^{\mathrm{PE}} \cdot \bigl(N_\phi^{\mathrm{L}}/2\bigr) + 1, \\
&m^{\mathrm{L}} \text{ even:} && m = N_\phi/2 + m^{\mathrm{L}}/2 + n_\phi^{\mathrm{PE}} \cdot \bigl(N_\phi^{\mathrm{L}}/2\bigr)
\end{aligned}
\tag{12}
$$

raises also the restriction that N_ϕ^{L} has to be even. Therefore, we introduced one additional DVR grid point, which is ignored as an FBR basis function, resulting in $N_\phi = 2j_{\max} + 2$ grid points. With the presented data decomposition, we ultimately achieve a perfectly balanced data distribution in all data structures as shown by the color scheme in Fig. 2. Consequently, a perfect load balancing is realized with only few communications by a data processing that is closely related to the data distribution.

The main communication on the data structures amounts to the preparation and redistribution of the global data before and after the transformations of the DVR quadrature. To transform a certain coordinate, the global dimension is collected on all involved PEs while the lost parallelism is recovered by a work decomposition in another

coordinate (gathering). Then, the transformation of the global dimension is performed on the work decomposed structure. By redistributing the transformed global dimension among the involved PEs, the work decomposition is finished (scattering). The consecutive transformation of all coordinates according to this procedure leads to an extremely efficient wavefunction processing:

- gather global data in x (work decomposition in ϕ)
  perform FFT $k_x \to x$
  scatter transformed data in x
- gather global data in z (work decomposition in ϕ)
  perform FFT $k_z \to z$
  scatter transformed data in z
- gather global data in θ (work decomposition in ϕ)
  perform Gauss–Legendre transform $j \to \theta$
  scatter transformed data in θ
- gather global data in ϕ (work decomposition in θ)
  perform FFT $m \to \phi$
  scatter transformed data in ϕ
- action of potential energy operator upon transformed wavefunction
- gather global data in ϕ (work decomposition in θ)
  perform FFT $\phi \to m$
  scatter transformed data in ϕ
- gather global data in θ (work decomposition in ϕ)
  perform Gauss–Legendre transform $\theta \to j$
  scatter transformed data in θ
- gather global data in z (work decomposition in ϕ)
  perform FFT $z \to k_z$
  scatter transformed data in z
- gather global data in x (work decomposition in ϕ)
  perform FFT $x \to k_x$
  scatter transformed data in x.

This transformation sequence narrows the temporarily global data to only one dimension. For further reduction of communication overhead, only contiguous memory segments are communicated to avoid the communication of strided data. In the transformation loops the PE integer arithmetics in Eqs. (10)–(12) is implemented as fast inlined module functions.

In contrast to conventional parallelization schemes we propose a data decomposition in all degrees of freedom in order to aim at large scale applications. Contemporary parallelization strategies in quantum dynamical studies represented, for instance, by the work of Goldfield et al. [25] favor a data decomposition in one or few dimensions only. Such procedure works quite well for medium-sized problems ($N = 3.2 \times 10^7$) as investigated in Goldfield's study but it fails clearly for great challenge applications ($N > 10^9$) with large dimension extents in all degrees of freedom: for a data distribution in one or few coordinates, an increasing problem size in the non-distributed dimensions leads to extremely small extents in the distributed dimensions. Consequently, these dimensions should be assigned to the last indices of the wavefunction FORTRAN array to reduce local strided data access. However, this implies an oversized effort for global data access since the large strides of the last indices have to be reordered before communicating the data [25]. Therefore, a problem decomposition in all dimensions is compulsory for large scale applications. Following this necessity, our data distribution provides optimal local data access and minimizes the reorder effort for global data access.

3.2. Parallel performance

The performance outcome of our quantum dynamics code is demonstrated by a speedup analysis that describes the dependence of the run time $t_{\mathrm{run}}(N, P)$ on both the system size N and the number of PEs P. All timing calculations are performed by using the Split propagator. Since the transformations of the DVR quadrature are the essential communication involving part of the Split propagator as well as of the Chebyshev propagator, the performance results are representative of both propagation schemes.

First of all, we fix the 4D system $z \times x \times \phi \times \theta$ [3] to the medium size $128 \times 128 \times 64 \times 32$ in order to study the behaviour as the number of PEs varies. The resulting fixed size speedup $t_{\mathrm{run}}(N, 1)/t_{\mathrm{run}}(N, P)$ [26] shown in Fig. 3 is almost linear, which indicates a very high computation to communication ratio. An efficiency of 89% for 512 PEs reveals the feasibility of great challenge applications on even more sizable supercomputers.

In addition, we investigate the scaling properties in certain dimensions. By successively doubling the extent and the number of PEs in the dimension of interest, the overall size per PE is fixed to $64 \times 64 \times 32 \times 16$. The corresponding scaled speedup $Pt_{\mathrm{run}}(N, 1)/t_{\mathrm{run}}(PN, P)$ [26] assuming a linear scaling is shown in Fig. 4. For the Cartesian coordinates involving FFTs, the performance is very close to the linear speedup. An almost perfect behaviour is observed for the distance coordinate z because of its contiguous data structure in the first array dimension. The slighter slope for the lateral coordinate x is due to the reordering overhead of the strided data in the second array dimension.

For the closely connected angular coordinates, the initial performance is governed by the semi-linear scaling of the FFT in the azimuthal angle ϕ. With increasing extent of the angular dimensions, the quadratic scaling of the Gauss–Legendre transform in the polar angle θ starts to compete. The resulting speedup slowly drops behind the supposed linear behaviour more and more following the quadratic scaling of the Gauss–Legendre transform. However, the asymptotic quadratic scaling is irrelevant in practical applications since the computational effort of the Gauss–Legendre transform in θ only equals that of the FFT in ϕ. Hence, the quadratic scaling does not even dominate the semi-linear scaling within the transformation of the angular coordinates. This proves the semi-linear total scaling of the Hamiltonian anticipated in Section 2. In conjunction with the outstanding communication efficiency (see Fig. 3), the semi-linear scaling relation lays the foundations for large scale applications by a profitable use of more than 512 PEs.

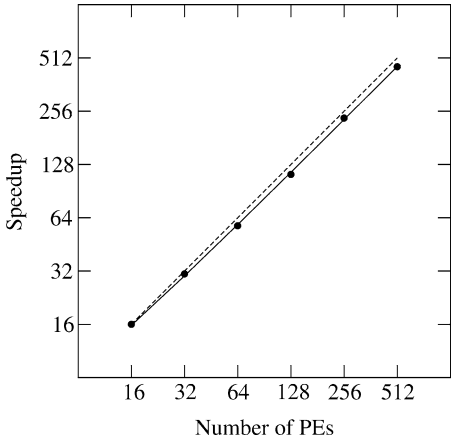

Fig. 3. Fixed size speedup for a 4D system $z \times x \times \phi \times \theta$ of the size $128 \times 128 \times 64 \times 32$ (solid line) compared to the linear speedup (dashed line).

[3] The order of the FORTRAN array dimensions is due to the following guideline: rapidly transforming dimensions of large extent first to guarantee optimal data access.

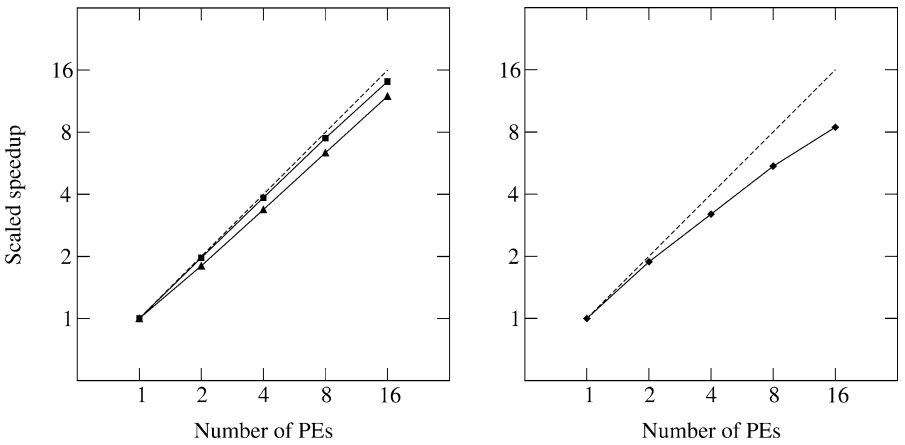

Fig. 4. Scaled speedup for a 4D system $z \times x \times \phi \times \theta$ of the size $64 \times 64 \times 32 \times 16$ per PE in the different dimensions (solid line). The left graph shows the performance for the Cartesian coordinates z (squares) and x (triangles). The behaviour for the angular coordinates $\{\theta, \phi\}$ (diamonds) is illustrated in the right graph. For comparison the linear speedup (dashed line) is given.

In a full 4D calculation on 128 PEs of the Cray T3E, our implementation reaches an average floating point performance of about 100 MFlops per PE. The calculation of a 1.5 ps quantum trajectory takes about 10 hours propagating a 5 GByte wavefunction of the size $N = 3.4 \times 10^8$. Most of the run time, about 90%, is spent on the propagation scheme, 80% on the transformations of the DVR quadrature alone. The remaining time is spent initializing data, computing observables and writing data.

To our knowledge, the only quantum dynamical calculation mastering a similarly sized problem as we do is a six dimensional study of the adsorption dynamics of $H_2/Cu(111)$ by Dai and Light [9]. However, the adsorption dynamics of the lightweight H_2 molecule can be resolved by much smaller extents in the particular dimensions than we use. Accordingly, the non-linear scaling part of the algorithm is decreased in comparison to our calculations. Only the resulting reduction of the computational effort in combination with the short propagation times of about 300 fs makes this six-dimensional investigation feasible. We believe that great challenge applications characterized by large problem sizes and complex dynamics can only be approached by consequently utilizing massively parallel computing power with very efficient parallelization strategies like ours.

4. Application: laser induced desorption of CO from $Cr_2O_3(0001)$

As a large scale application, we simulate the laser induced desorption of CO molecules from a $Cr_2O_3(0001)$ surface that we successfully investigated in 3D studies [27]. In such a process, the laser irradiation generates an electronically excited intermediate. During the short lifetime of the intermediate, the adsorbate changes its nuclear geometry with respect to the substrate collecting kinetic energy in the nuclear degrees of freedom. After relaxation to the electronic ground state, the gained nuclear kinetic energy initiates the desorption of the adsorbate.

This DIET (desorption induced by electronic transitions) [28] process is treated within a two state model using 4D potential energy surfaces (PES) that were obtained from embedded cluster ab initio calculations on configuration interaction (CI) level [29,30]. For the lateral dynamics, the coordinate x is assigned to the short Cr–Cr axis on the Cr terminated surface. The laser induced transition is modeled as a Franck–Condon excitation of the rovibrational groundstate to the electronically excited state that turns out to be an internal CO ($5\sigma \rightarrow 2\pi^*$) excitation of the adsorbed CO molecule [30]. After a time evolution in the excited state for a characteristic lifetime of $\tau = 36$ fs, the wave packet is transferred to the electronic ground state. The subsequent propagation in the ground

state is proceeded up to convergence of the desorbing wave packet part that is consecutively separated by a grid change method [31].

Of course, this simplified model gives only a qualitative insight into the photodesorption process. A more accurate description of the electronic transition is provided by a stochastic wave packet method including lifetime averaging [32]: assuming an exponential decay of the excitation, the observables are averaged with respect to the spectral resonance lifetime of the electronically excited state. The procedure rapidly converges the expensive exact density matrix description [33] but only necessitates the calculation of an ensemble of quantum trajectories with varied residence lifetimes. Thus, a still enormous computational effort allows a correct description of the photodesorption process even for high-dimensional systems, which is intractable by density matrix methods relying on the solution of the Liouville–von Neumann equations. The results from this more sophisticated stochastic wave packet approach will be presented elsewhere [34].

In the analysis of the finally desorbed part of the wave packet, the kinetic properties are investigated. For the desorption coordinate z and the lateral coordinate x, the velocity distributions are shown in Fig. 5. The occurring desorption velocities up to 2000 m/s and the bimodality of their distribution as well as lateral velocities around 500 m/s are in qualitative agreement to experimental results [35]. To characterize the stereodynamics, the quadrupole moment [36]

$$A_0^2(J) = \frac{\langle J|3\widehat{M}^2 - \hat{J}^2|J\rangle}{\hat{J}^2} \tag{13}$$

depending on the rotational excitation J with the angular momentum $\hat{\mathbf{J}}$ and its projection $\widehat{M}$ is shown in Fig. 6. This quantity indicates the rotational alignment of the desorbing molecules. It ranges from -1 for pure cartwheel rotation ($\mathbf{J} \perp z$) to $+2$ for pure helicopter rotation ($\mathbf{J} \parallel z$). The experimental observation of helicopter motion for medium rotational states and cartwheel motion for strong rotational excitation is qualitatively reproduced by the calculated results. However, deviating from the experimental observation the shape of the computed curve shows two pronounced peaks of helicopter motion instead of only a broad one before falling off to cartwheel motion at much higher rotational states.

An exhaustive analysis of the differences to the experimental data requires both further experimental and theoretical investigations. For a detailed discussion and the results from the stochastic wave packet method [32], we refer to the above mentioned studies [34].

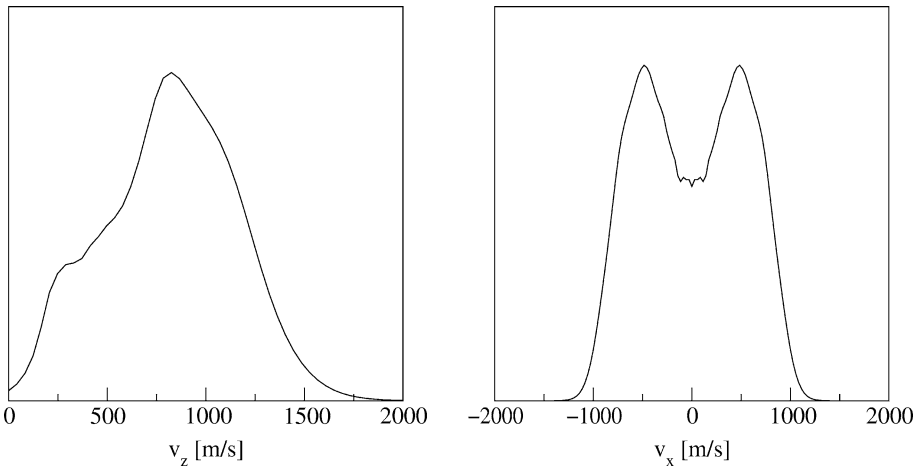

Fig. 5. Velocity distribution for the desorption of CO from $Cr_2O_3(0001)$ in z- and x-direction.

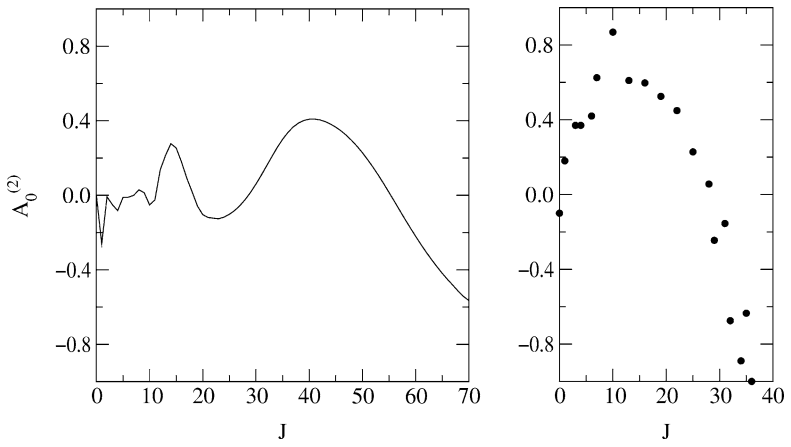

Fig. 6. Quadrupole moment for the desorption of CO from Cr_2O_3(0001) as a function of rotational excitation: comparison of computed results (left graph) and experimental data [35] (right graph).

5. Conclusions

In the present work, we introduce a massively parallel implementation of a quantum dynamical wave packet method for molecules on surfaces. A highly scalable algorithm based on a finite basis representation [19] is optimally realized on massively parallel systems by a homogeneous data decomposition in all dimensions. The very high parallel efficiency up to 512 processing elements shows the feasibility of high-dimensional studies of a size $N > 10^9$. Consequently, our implementation paves the way to successfully simulate complex experimental results on quantum dynamical level, which until now was impossible due to the limited number of degrees of freedom taken into account. First results of 4D calculations investigating a laser induced desorption process are presented.

Acknowledgements

This work has been financially supported by the Deutsche Forschungsgemeinschaft (DFG, SPP1093) and by the German-Israeli Foundation (GIF). We thank Dr. D. Lemoine and Prof. Dr. R. Kosloff for many fruitful discussions.

References

[1] D. Kosloff, R. Kosloff, J. Comput. Phys. 52 (1983) 35;
R. Kosloff, J. Phys. Chem. 92 (1988) 2087.
[2] G.C. Corey, D. Lemoine, J. Chem. Phys. 97 (1992) 4115.
[3] G.C. Corey, J.W. Tromp, D. Lemoine, in: C. Cerjan (Ed.), Numerical Grid Methods and Their Application to Schrödinger's Equation, Kluwer Academic Publishers, Boston, 1993.
[4] G.C. Corey, J.W. Tromp, J. Chem. Phys. 103 (1995) 1812.
[5] C. Leforestier, R.H. Bisseling, C. Cerjan, M.D. Feit, R. Friesner, A. Guldberg, A. Hammerich, G. Jolicard, W. Karrlein, H.-D. Meyer, N. Lipkin, O. Roncero, R. Kosloff, J. Comput. Phys. 94 (1991) 59.
[6] I. Andrianov, V. Bonačić-Koutecký, M. Hartmann, J. Manz, J. Pittner, K. Sundermann, Chem. Phys. Lett. 318 (2000) 256.
[7] S. Li, H. Guo, J. Chem. Phys. 115 (2001) 3330.
[8] Ğ. Barinovs, N. Marković, G. Nyman, J. Chem. Phys. 111 (1999) 6705.
[9] J. Dai, J.C. Light, J. Chem. Phys. 107 (1997) 1676; J. Chem. Phys. 108 (1998) 7816.
[10] W. Zhu, J.Z.H. Zhang, Y.C. Zhang, Y.B. Zhang, L.X. Zhan, S.L. Zhang, D.H. Zhang, J. Chem. Phys. 108 (1998) 3509.

[11] G.A. Worth, H.-D. Meyer, L.S. Cederbaum, J. Chem. Phys. 109 (1998) 3518.
[12] A. Raab, G.A. Worth, H.-D. Meyer, L.S. Cederbaum, J. Chem. Phys. 110 (1999) 936.
[13] R.E. Wyatt, C. Iung, J. Chem. Phys. 98 (1993) 6758.
[14] A. Maynard, R.E. Wyatt, C. Iung, J. Chem. Phys. 106 (1997) 9483.
[15] E.M. Goldfield, Comput. Phys. Commun. 128 (2000) 178.
[16] P. Eggert, A. Viel, C. Leforestier, Comput. Phys. Commun. 128 (2000) 315.
[17] Z. Bačić, Comput. Phys. Commun. 128 (2000) 46.
[18] R.B. Lehoucq, S.K. Gray, D.H. Zhang, J.C. Light, Comput. Phys. Commun. 109 (1998) 15.
[19] D. Lemoine, J. Chem. Phys. 101 (1994) 10 526.
[20] D. Lemoine, Comput. Phys. Commun. 97 (1996) 331.
[21] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes, Cambridge University Press, Cambridge, 1992.
[22] M.D. Feit, J.A. Fleck, A. Steiger, J. Comput. Phys. 47 (1982) 412.
[23] H. Tal-Ezer, R. Kosloff, J. Chem. Phys. 81 (1984) 3967.
[24] Cray T3E Fortran Optimization Guide 004-2518-002.
[25] E.M. Goldfield, S.K. Gray, G.C. Schatz, J. Chem. Phys. 102 (1995) 8807.
[26] J.L. Gustafson, G.R. Montry, R.E. Benner, SIAM J. Sci. Stat. Comput. 9 (1988) 609.
[27] S. Thiel, M. Pykavy, T. Klüner, H.-J. Freund, R. Kosloff, V. Staemmler, Phys. Rev. Lett. 87 (2001) 077601; J. Chem. Phys., submitted.
[28] N.H. Tolk, M.M. Traum, J.C. Tully, T.E. Madey (Eds.), Desorption Induced by Electronic Transitions, DIET, Vol. 1, Springer-Verlag, Berlin, 1983.
[29] M. Pykavy, V. Staemmler, O. Seiferth, H.-J. Freund, Surf. Sci. 479 (2001) 11.
[30] M. Pykavy, S. Thiel, T. Klüner, H.-J. Freund, V. Staemmler, to be published.
[31] R. Heather, H. Metiu, J. Chem. Phys. 86 (1987) 5009.
[32] J.W. Gadzuk, L.J. Richter, S.A. Buntin, D.S. King, R.R. Cavanagh, Surf. Sci. 235 (1990) 317;
J.W. Gadzuk, Surf. Sci. 342 (1995) 345.
[33] P. Saalfrank, Chem. Phys. 211 (1996) 265.
[34] S. Borowski, S. Thiel, T. Klüner, H.-J. Freund, to be published.
[35] I. Beauport, K. Al-Shamery, H.-J. Freund, Chem. Phys. Lett. 256 (1996) 641;
I. Beauport, PhD thesis, Ruhr-Universität Bochum, Bochum, 1996.
[36] C.H. Greene, R.N. Zare, J. Chem. Phys. 78 (1983) 6741.