VU Research Portal

Dynamical pruning of static localized basis sets in time-dependent quantum dynamics

McCormack, D.A.

published in

Journal of Chemical Physics
2006

DOI (link to publisher)

[10.1063/1.2196889](https://doi.org/10.1063/1.2196889)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

McCormack, D. A. (2006). Dynamical pruning of static localized basis sets in time-dependent quantum dynamics. *Journal of Chemical Physics*, 124(20), 204101. <https://doi.org/10.1063/1.2196889>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Dynamical pruning of static localized basis sets in time-dependent quantum dynamics

Drew A. McCormack^{a)}*Theoretische Chemie, Faculteit Exacte Wetenschappen, Vrije Universiteit, De Boelelaan 1083, 1081 HV Amsterdam, The Netherlands*

(Received 23 February 2006; accepted 24 March 2006; published online 22 May 2006)

We investigate the viability of dynamical pruning of localized basis sets in time-dependent quantum wave packet methods. Basis functions that have a very small population at any given time are removed from the active set. The basis functions themselves are time independent, but the set of active functions changes in time. Two different types of localized basis functions are tested: discrete variable representation (DVR) functions, which are localized in position space, and phase-space localized (PSL) functions, which are localized in both position and momentum. The number of functions active at each point in time can be as much as an order of magnitude less for dynamical pruning than for static pruning, in reactive scattering calculations of H_2 on the Pt(211) stepped surface. Scaling of the dynamically pruned PSL (DP-PSL) bases with dimension is considerably more favorable than for either the primitive (direct product) or DVR bases, and the DP-PSL basis set is predicted to be three orders of magnitude smaller than the primitive basis set in the current state-of-the-art six-dimensional reactive scattering calculations. © 2006 American Institute of Physics. [DOI: 10.1063/1.2196889]

I. INTRODUCTION

Time-Dependent wave packet (TDWP) methods are widely used to study quantum molecular dynamics, but are generally only applicable to small chemical systems. A multidimensional wave packet is typically represented by a direct product of reduced-dimensional basis functions or grids; the size of such representations grows exponentially with dimension, limiting TDWP methods to around six degrees of freedom (DOFs) on currently available computer hardware.¹ (Calculations for higher dimensions have been performed in certain favorable cases,² but are not generally tractable.) Increasing computational resources only allow this effective ceiling on dimensionality to rise slowly over time, so algorithmic improvements are all the more important.

There are two ways in which a representation can be made more compact in a wave packet calculation: The first is to adapt it to the problem at hand, by improving the quality of the reduced-dimensional basis functions, or pruning the multidimensional, direct-product basis set.³ Examples of these approaches include using contracted basis functions, as in the potential-optimized discrete variable representation (PODVR),⁴ applying a projection operator formalism in order to utilize different basis functions for different regions of space,⁵ and discarding basis functions that fall outside the energetic range of interest.⁶

The second way to reduce the size of a representation is to make it time dependent, evolving with the changing wave packet. An example is the time-dependent self-consistent-field (TDSCF) method, which represents the wave packet at each point in time by a Hartree product of one-dimensional, time-dependent functions.⁷ A completely different approach

is to utilize an adaptive grid^{8–11} or basis set^{12–19} that moves with the wave packet, following classical trajectories, for example.

Probably the best-scaling quantum wave packet method currently available is the multiconfiguration time-dependent Hartree (MCTDH) method,²⁰ which applies both of the approaches discussed above, representing the wave packet by a direct product of reduced-dimension, time-dependent functions.²¹ The single-particle functions, as they are known, are optimally contracted to represent the wave packet at each point in time. MCTDH has facilitated studies well beyond what has been possible with conventional TDWP methods. In a recent study, several thousands of degrees of freedom were treated with a model potential,²² and studies on realistic potential energy surfaces (PESs) have included up to 12 DOFs.²³

MCTDH is a highly successful method, but not a panacea. It is effective because it exploits structure in the wave packet, namely, weak correlation between DOFs.²¹ If the wave packet becomes strongly correlated, due to trapping, for example, the MCTDH may actually be more expensive than a traditional TDWP method. In recently performed four-DOF calculations with a general PES, Crespos *et al.* found MCTDH to be only slightly more efficient than a standard TDWP method,²⁴ and in six-DOF calculations on a model PES, van Harreveld and Manthe did not observe any of the resonance structure arising from long-time recurrences.²⁵

A compact direct-product basis set can be constructed for systems with weak correlation—as demonstrated by MCTDH—but direct-product bases are generally an inefficient means of representing a multidimensional function. Direct-product representations are usually used in TDWP methods because they lead to simple sparse matrix represen-

^{a)}Electronic mail: da.mccormack@few.vu.nl

tations of quantum operators, which reduce the complexity of the algorithms and programs needed to solve the equations of motion. The price paid for this simplicity is the immensity of the resulting basis set. However, a direct-product basis can be made considerably more compact through “pruning:” multidimensional basis functions that are irrelevant to the problem at hand can be discarded. Dawes and Carrington have recently applied this technique in time-independent quantum calculations.³

In this paper, we investigate dynamical pruning of basis sets comprised of static functions. We exploit both of the tactics discussed earlier in order to reduce the overall size of the basis set: The basis functions used here are localized, with only those that significantly overlap with the wave packet being retained, and the basis sets are time dependent, though the basis functions themselves do not vary. (In other words, the active space varies in time.)

Introducing time dependence into the basis set, but not the individual basis functions, has some advantages over letting the functions themselves vary in time. First, it is not necessary to reformulate equations of motion or utilize specialized propagators; the standard algorithms of TDWP methods can be applied. Methods with time-dependent basis functions, such as MCTDH, usually have more complex, nonlinear equations of motion.²¹ These have to be solved using general-purpose differential equation solvers, which are often less efficient than those purposely developed to solve the time-dependent Schrödinger equation.

The second advantage of using static basis functions is that the basis set is “well behaved.” It cannot become overcomplete or too sparse to adequately represent the wave packet. Methods in which basis functions follow trajectories in time suffer particularly from these problems, and often require workarounds (see, for example, Ref. 16). With fixed basis functions, an orthogonal basis set can be set up that is appropriately distributed throughout the region of relevance, and will remain so throughout the propagation.

The last benefit of the approach adopted here is that no implicit assumptions need to be made about the form of the wave packet in constructing a basis set. By contrast, MCTDH is based on the assumption of weak correlation, and is inefficient in strongly correlated systems. The pruning technique tested here will be less efficient than MCTDH in weakly correlated systems, but should be more generally applicable.

Two different types of localized basis functions are utilized. The first is based on the ubiquitous DVR.²⁶ A primitive basis is formed as the direct product of one-dimensional DVR functions; DVR functions are maximally localized in position space. Functions of the second type are localized in phase space, and are referred to here as phase-space localized (PSL) functions. They are less localized in position space than DVR functions, but are nonetheless well localized in both the position and momentum representations. The benefit of a basis of PSL functions is that potentially fewer functions need to be included, because a PSL basis function can be “culled” based not only on its position (potential energy) but also on its range of momentum (kinetic energy).

PSL functions come in various forms, coherent-state

functions being one well-known example,^{27–29} and *wavelets* being another.^{30–38} Wavelets are localized in both the position and momentum representations, and are commonly used for spectral analysis, but have also found application to quantum dynamics.^{33,34,36–38} Here we use a different approach to generating PSL functions, one that has recently been introduced by Dawes and Carrington.^{3,39} They use the method of simultaneous diagonalization (SD), which seeks a single set of eigenfunctions that diagonalize the position and momentum operator matrices. It can be shown that this is equivalent to simultaneously localizing a set of basis functions in the position and momentum representations. It is not actually possible to completely diagonalize the matrices, but it is possible to nearly diagonalize them, and thereby obtain a basis of localized functions. Dawes and Carrington have shown that a basis set generated in this way can be very effective for time-independent quantum dynamics.^{3,39}

In this study, we test the SD technique of Dawes and Carrington, along with dynamical pruning of the basis set, in reduced-dimensional, time-dependent wave packet calculations for reactive scattering of H₂ on a Pt(211) stepped surface. This system has recently been the subject of classical trajectory^{40–42} and quantum wave packet calculations,⁴³ and is thus a suitable test bed for new quantum dynamics methods.

We will focus here not only on the efficiency of dynamical culling of localized basis functions in quantum wave packet methods but also on the implied inefficiency of standard TDWP methods. We will show that, even for low-dimensional calculations, a pruned basis can be several orders of magnitude more compact than a primitive direct-product basis, implying that of all the information encapsulated in a direct-product basis set, only a fraction of a percent may be relevant to the wave packet at any given time. This realization will hopefully stimulate other researchers to seek new ways of improving upon direct-product representations.

In Sec. II, the SD algorithm used to localize PSL basis functions is introduced, and various aspects of the wave packet method detailed. The approach taken to culling basis functions is also discussed. In Sec. III, the results are presented for wave packet calculations on the H₂+Pt(211) system, with dynamically pruned DVR and PSL basis sets. Section IV concludes.

II. METHOD

A. Potential energy surfaces

Calculations of varying dimensionality were performed on the H₂+Pt(211) PES.⁴² One-dimensional (1D) calculations were performed that only treated the molecule-surface distance coordinate (*Z*), with the molecule parallel to the surface above the top-of-step Pt atom. In these calculations, the bond length (*r*) was fixed, so reaction was not possible. Although this 1D PES is predominantly repulsive, it does include an attractive well.

Two different two-dimensional (2D) systems were tested (Fig. 1). The first was for the same top-of-step site and molecular geometry as in the 1D calculations, but with the bond

FIG. 1. The $\text{H}_2+\text{Pt}(211)$ potential energy surfaces (PES) used in the 2D scattering calculations. (a) Two-dimensional PES cut for the H_2 molecule parallel to the surface, above a step atom, and perpendicular to the step edge. This is a nonactivated site, with no barrier to reaction. (b) Cut for the same orientation as (a), but with the molecule located above the center of the terrace. This site is activated, with a barrier around 0.15 eV high.

length allowed to vary. The second was for the molecule centered over a terrace site atom, rather than atop a step atom. The PES for the step site is barrierless [Fig. 1(a)] leading most molecules to react, while the terrace site PES exhibits an early barrier of around 0.15 eV [Fig. 1(b)], which causes many molecules to be scattered.

Three-dimensional calculations were performed for the same geometry as the 2D step site calculations, but with translational motion also possible along the step edge, which is the X coordinate direction. This PES includes the attractive site in the 2D PES [Fig. 1(a)], but additionally includes repulsive intermediate sites where no reaction is possible.⁴²

Four-dimensional calculations were as for the 2D calculations, but with motion possible in both translational coordinates parallel to the surface, the X and Y coordinate directions. The four-dimensional (4D) PES includes the reactive step and terrace sites, and many unreactive sites. Shallow chemisorption wells also exist, which can lead molecules to become trapped for some time on the surface.

B. Time-dependent wave packet method

A standard TDWP method was used to propagate two wave packets for each calculation: the first was the so-called “exact” wave packet (EWP), which was represented entirely in the primitive direct-product basis set; the second was the wave packet propagated in the pruned basis set, which we will term the pruned-basis wave packet (PBWP).

The Hamiltonians used for the various calculations were very simple, including only the potential energy, and the translational and vibrational terms. For example, the Hamiltonian for the 4D calculations was given by

$$\hat{H} = -\frac{1}{2\mu} \frac{\partial^2}{\partial r^2} - \frac{1}{2M} \left(\frac{\partial^2}{\partial X^2} + \frac{\partial^2}{\partial Y^2} + \frac{\partial^2}{\partial Z^2} \right) + V_{4D}(r, X, Y, Z), \quad (1)$$

where V_{4D} is the 4D potential energy, M is the molecular mass, and μ is the reduced mass corresponding to the mo-

lecular bond length coordinate. Hamiltonians for other dimensionalities were similar to Eq. (1), but with inapplicable terms removed.

The initial state of the wave packet in all calculations was the product of a Gaussian coherent-state function for the scattering coordinate, Z , and the ground state of any other degrees of freedom. For example, for the 4D calculations, the initial state was

$$\psi(r, X, Y, Z; t=0) = \frac{1}{\sqrt{L_X}} \frac{1}{\sqrt{L_Y}} \phi_{v=0}(r) G(Z), \quad (2)$$

with L_X and L_Y the lattice constants for the X and Y coordinates, respectively, $\phi_{v=0}(r)$ the ground vibrational state, and $G(Z)$ the coherent-state function. Parameters for $G(Z)$ were chosen such that the wave packets were initially centered at $Z=8.0$ bohrs, and moving toward the surface with 95% probability of having a kinetic energy in the range of 0.1–0.3 eV.

A 1D DVR derived from sinusoidal functions was used as the primitive representation for each DOF.^{6,44} The grid points in this DVR are equally spaced, but the spacing was chosen differently for each DOF, according to the convergence tests performed recently for a six-dimensional (6D) quantum study.⁴³ For Z , a spacing of 0.150 bohr was used; for r , 0.208 bohr; for X , 0.175 bohr; and for Y , 0.160 bohr.

Each wave packet was propagated with the third-order split operator propagator,⁴⁵ using a time step of 5.0 a.u. Applying the split operator propagator involves continually operating on the wave function with exponential operators for the kinetic energy and potential energy,

$$\psi(t = t_{n+1} = t_n + \Delta t) = e^{-i\hat{K}\Delta t/2} e^{-i\hat{V}\Delta t} e^{-i\hat{K}\Delta t/2} \psi(t = t_n), \quad (3)$$

where n is the time step index, Δt is the propagation time step, $\hat{K}$ is the kinetic energy operator, and $\hat{V}$ is the potential energy operator.

For the wave packets represented in the DVR basis sets, the exponential kinetic energy operator was factorized into 1D operators in the standard way, before being applied to the wave function coefficients. For example, in the 4D calculations, the kinetic exponential operator was factorized as

$$e^{-i\hat{K}\Delta t/2} = e^{-i\hat{K}_X\Delta t/2} e^{-i\hat{K}_Y\Delta t/2} e^{-i\hat{K}_Z\Delta t/2} e^{-i\hat{K}_r\Delta t/2}, \quad (4)$$

with $\hat{K}_X$, $\hat{K}_Y$, $\hat{K}_Z$, and $\hat{K}_r$ the various 1D kinetic operators. Each exponential kinetic energy operator was formed in the kinetic energy matrix eigenspace, where it is diagonal, before being transformed to the DVR representation. The exponential potential operator was simply a pointwise multiplication, because the potential operator is diagonal in the DVR representation.

For the PSL basis sets, the exponential kinetic operator was applied in the same way as for the DVR basis sets, but a different approach was used for the exponential potential operator matrix, because it is not diagonal in the PSL basis set. The wave function coefficients were transformed to the DVR representation, and the operation performed there, before transforming the coefficients back to the PSL representation. Of course, this is very inefficient, and any practical imple-

mentation utilizing PSL basis functions would need to include a better approach to performing the potential operation. This will be discussed further in Sec. III.

Complex absorbing potentials⁴⁶ were used in the Z and r coordinates to absorb the wave packet as it approached the edge of the grid, after being scattered or having reacted. This avoids unphysical reflections from the grid edge.

C. Localized basis functions

PSL functions were generated from the primitive DVR for each DOF using the Jacobi method of SD, adopted from the work of Dawes and Carrington.^{3,39,47} The method was used here to generate near diagonal matrices, rather than completely diagonal matrices, because full diagonalization is not possible for noncommuting matrices.

Various operators can be used in the near diagonalization, but in this study we chose the position operator and the kinetic energy operator, because both are real valued. With this choice, near diagonalization of the operator matrices involves maximizing the following function of diagonal matrix elements:

$$\sum_i [U^T \mathbf{x} U]_{ii}^2 + \sum_i [U^T \mathbf{K}_x U]_{ii}^2, \quad (5)$$

where $\mathbf{U}$ is a unitary transformation, $\mathbf{x}$ is the position operator matrix, and $\mathbf{K}_x$ is the kinetic energy matrix for coordinate x . (The problem can also be formulated as a minimization of the square of the off-diagonal matrix elements.) The unitary transformation matrix $\mathbf{U}$ is the product of Jacobi rotation matrices $\mathbf{J}(i, j)$, which can be determined from expressions given in Ref. 48.

The choice of the matrices $\mathbf{x}$ and $\mathbf{K}_x$ is somewhat arbitrary. Multiplying either matrix by a constant factor leads to an equally valid choice. In this study, we chose a standard scaling for the matrices: First, each matrix was scaled such that the magnitude of the maximum matrix element was unity, giving the matrices $\bar{\mathbf{x}}$ and $\bar{\mathbf{K}}_x$. Next, a constant factor was introduced to control the contribution of each matrix to the optimization function,

$$\alpha \sum_i [U^T \bar{\mathbf{x}} U]_{ii}^2 + (1 - \alpha) \sum_i [U^T \bar{\mathbf{K}}_x U]_{ii}^2. \quad (6)$$

The bias factor α ranges from 0 to 1, and can be used to change the priority of each operator in the diagonalization process. If α is 0, the position operator term plays no part, and only the kinetic energy matrix is diagonalized. This extreme gives the sinusoidal basis functions, which are eigenfunctions of the kinetic energy operator. If α is 1, the kinetic energy term in Eq. (6) becomes zero, and only the position operator is diagonalized, giving the DVR basis. For values of α in between, each matrix will become near diagonal, with the value of α determining how close each matrix is to full diagonality.

Optimization of the function in Eq. (6) was achieved with the Jacobi method.⁴⁹ This entails performing “sweeps” of the matrices, visiting each element in turn, and performing unitary Jacobi rotations in order to reduce the magnitude of off-diagonal elements. Unfortunately, this process does not

always converge rapidly, and can be adversely affected by so-called “fixed points.” Others have utilized a Newton-Raphson (NR) algorithm to improve convergence in cases where fixed points arise.⁵⁰

In this study, we utilized a rescaling scheme similar to that implemented by Dawes and Carrington,^{3,47} in order to address the fixed-point problem, but this did not always result in good convergence. Nonetheless, we took no further steps to improve the optimization method, because the convergence was found to be adequate for our purposes, and other authors have noted that poor convergence has very little impact on the quality of the localized basis functions.⁵⁰

D. Basis set pruning

We have not attempted to develop a particular algorithm for dynamical pruning in this study, nor have we developed specialized software that efficiently implements dynamical pruning in wave packet calculations. Instead, a standard TDWP method has been used to investigate the potential gains of dynamical pruning, without providing a practical implementation.

Rather than removing basis functions from a basis set, we propagated each wave packet in the full-unpruned basis set, but set small coefficients to zero after each time step. This has a similar effect to removing the corresponding basis functions from the calculation for that time step. The number of coefficients that were zeroed was monitored as a function of time.

Two different types of *culling threshold* were used in determining what constituted a “small” coefficient. The first involved a fixed, absolute minimum on the absolute square of any coefficient,

$$c_i^* c_i \geq \varepsilon, \quad (7)$$

where c_i is a complex coefficient, and ε the culling threshold. Any coefficient for which Eq. (7) was not true was set to zero, effectively culling the corresponding basis function from the basis set.

The second type of threshold used was time variant. This culling threshold was adopted where coefficients could vary significantly in magnitude from one wave packet to the next, as arises when comparing systems with different dimensionalities. In this case, the threshold in Eq. (7) was recalculated each time step according to

$$\varepsilon(t) = \varepsilon_{\text{rel}} c_{\text{max}}^*(t) c_{\text{max}}(t), \quad (8)$$

where $c_{\text{max}}(t)$ is the coefficient with the largest absolute value at time t , and ε_{rel} is a relative threshold. Using this time-dependent threshold also means that the general decrease in coefficient values that occurs due to the spread of a wave packet does not lead to excessive culling of basis functions toward the end of a propagation.

It is important to recognize that as ε goes to zero, the number of culled basis functions also goes to zero, and the pruned-basis calculation becomes as accurate as the primitive basis set calculation. The localization of the basis functions in the PSL basis set introduces no extra approximations: the PSL basis and the primitive DVR basis are

FIG. 2. PSL basis set pruning for scattering of a 1D wave packet. [(a)–(c)] The location of basis functions in the kinetic energy-position space. Solid red dots represent culled basis functions, and open dots are basis functions in the active set. A static culling threshold $\varepsilon=10^{-6}$ was used. [(d)–(f)] The evolution of the absolute square of the wave packet (blue lines) is superimposed on the potential energy surface (gray lines). In both sets of plots, time increases down the page. Times are indicated in the charts on the right in atomic units [(d)–(f)].

isomorphic. The only approximation introduced in the pruning method is that arising from pruning itself.

III. RESULTS AND DISCUSSION

A. One-dimensional scattering

We begin our analysis by considering how the set of culled PSL basis functions evolves in time in 1D scattering. Figure 2 shows the evolution of the wave packet in time, and the PSL basis functions used to represent it. The position of each PSL basis function is assumed to be given by the pair of diagonal matrix elements corresponding to the basis function in the near diagonal kinetic and position matrices arising from the SD procedure. In this case, the SD bias parameter α was set to 0.5, favoring neither the kinetic matrix nor position matrix in the near diagonalization.

The wave packet approaches the interaction region, gathering momentum in the attractive well, before reflecting back to the asymptotic region. Oscillations in the wave function can be seen in Fig. 2(e), which are due to interference between the incoming and outgoing components. The wave packet is significantly broadened when it leaves the interaction region in Fig. 2(f).

Figure 2(a) shows that only a small portion of the PSL basis functions are needed to represent the initial wave packet. The region mapped out by the basis functions in the kinetic energy-position space is approximately elliptical. Interestingly, functions corresponding to low kinetic energies are excluded near the edges of the wave packet. This is due to the fact that the wave packet is not stationary, but moving toward the surface; this translates the elliptical region occupied by the wave packet upward in kinetic energy. Figure 3 clarifies this effect diagrammatically.

In Fig. 2(e), the wave packet is located in the attractive well, causing an increase in kinetic energy. This is mirrored in the basis set pruning of Fig. 2(b): the set of used basis functions extends higher up the kinetic energy scale than it does in Fig. 2(a), and is localized in position near the potential well.

In Fig. 2(c), with the wave packet scattering back [Fig. 2(f)], only the basis functions falling in a narrow range of kinetic energies are occupied in the asymptotic region. The reflected wave packet is broader in position space, and narrower in momentum, leading to the long thin “dagger” form seen in Fig. 2(c).

The results in Fig. 2 are for a PSL basis set generated

FIG. 3. Diagram depicting the region of phase space occupied by a moving 1D coherent-state wave packet. The motion of the wave packet results in the occupied region, which is elliptical, being translated along the momentum axis, away from the position axis. This causes low-momentum regions near the edges of the coordinate range to be unoccupied, which can be observed in the pattern of basis set pruning in Fig. 2(a).

with a bias parameter of $\alpha=0.5$. Figure 4 demonstrates the effect of varying α on basis set pruning. The time dependence of the percentage of culled functions for $\alpha=0.5$ remains relatively constant at just below 50%. Increasing α to 1, which corresponds to a DVR basis set, causes more basis functions to be culled at the beginning of the propagation, but fewer later, as the wave packet spreads in the position space. The other extreme, $\alpha=0$, which corresponds to a basis set of delocalized sinusoidal functions, is even more variable in time, with virtually no culling at all occurring while the wave packet is in the interaction region (i.e., between 2000 and 3000 a.u.).

The stability of pruning with $\alpha=0.5$ is favorable, because the likelihood of extremes is reduced. In practical computations, if the culling percentage were to drop to zero, as it does for $\alpha=0$, a computational bottleneck could arise, with increased memory and CPU time requirements. Choosing $\alpha=0.5$ also results in basis functions that are well localized in both position and momentum, which increases the effec-

FIG. 4. The time dependence of the percentage of basis functions culled for different values of the simultaneous diagonalization bias parameter α , in 1D wave packet calculations. Curves are labeled by the value of α used in the derivation of the PSL basis set. A static culling threshold of $\varepsilon=10^{-8}$ was used in the calculations.

FIG. 5. Accuracy of 1D wave packets propagated in pruned PSL basis sets. (a) Time dependence of the relative overlap of the exact and approximate wave packets for various values of the culling threshold, ε . (b) Time dependence of the percentage of PSL basis functions culled for the same culling thresholds used in (a).

tiveness of pruning in multidimensional systems, as we will show later. Given these considerations, we have chosen to use a SD bias of $\alpha=0.5$ for all PSL basis sets from this point onwards.

With the SD bias parameter established, only the choice of culling threshold (ε) still remains to be addressed. Figure 5 shows the time dependence of the normalized overlap of the exact and approximate wave packets, for various values of the threshold. The normalized overlap is defined here as

$$\rho(t) = \frac{\langle \psi_{\text{approx}}(t) | \psi_{\text{exact}}(t) \rangle}{\langle \psi_{\text{exact}}(t) | \psi_{\text{exact}}(t) \rangle}, \quad (9)$$

where ψ_{approx} is the wave function propagated in the pruned basis set, and ψ_{exact} is the wave packet propagated in the primitive direct-product basis set.

Figure 5 shows that the threshold must be quite low for accuracy to be maintained; wave packets are surprisingly sensitive to changes in small coefficients. For a threshold of 10^{-5} , the overlap drops off quite rapidly to less than 0.9. High accuracy is really only achieved with a threshold of 10^{-7} , for which the overlap remains well above 0.99 for most of the propagation, only dropping below at the end, when the wave packet has been mostly absorbed by the optical potential at the edge of the grid.

Even with low culling thresholds, a considerable portion of the PSL basis can still be pruned at each point in time [Fig. 5(b)]. For a threshold of 10^{-7} , around 50% of functions are culled, and for 10^{-8} , about 40%. As was also observed in

FIG. 6. Time dependence of PSL basis set pruning for 2D scattering. (a) Semilog plot of the normalized overlap for H₂ scattering at the step and terrace sites of Pt(211). Also shown is the norm of the wave packet propagated in the primitive direct-product basis, for each surface site. (b) The percentage of culled PSL basis functions vs time for scattering from the step and terrace sites. A static culling threshold of 10^{-8} was used in the calculations.

Fig. 4, the portion of the basis set pruned at each time is reasonably constant for any one value of the culling threshold.

B. Two-dimensional reactive scattering

Significant benefits of pruned PSL basis sets are really only realized in multidimensional calculations, where the basis set is expected to scale more favorably with dimension than a direct-product basis. We begin our exploration of multidimensional systems by considering the two different 2D scattering systems described earlier: one for the nonactivated step site, and one for the activated terrace site (Fig. 1).

Figure 6 shows the time dependence of the PSL basis set pruning for the two surface sites. The nonactivated step site is very reactive, with the wave packet absorbed relatively quickly [Fig. 6(a)]. The norm becomes very small after only around 4000 a.u. The normalized overlap [Eq. (9)] of the step wave packet remains close to 1 up to this point, and then begins to decrease as coefficients become small.

Scattering at the terrace site proceeds much more slowly, because a significant portion of the wave packet is scattered back. The norm of the wave packet reacting at the terrace only drops to around 10^{-1} over the course of the propagation, and the normalized overlap remains close to unity [Fig. 6(a)].

The percentage of basis functions culled remains relatively constant for the activated terrace site, but varies some-

FIG. 7. Time dependence of basis set pruning for 3D scattering of H₂ from the step edge of Pt(211). (a) The normalized overlaps of wave packets propagated in the DVR and PSL dynamically pruned basis sets, and the norm of the wave packet propagated in the primitive direct-product basis set. (b) The number of basis functions used to represent the wave packets in the DVR and PSL basis sets as a function of time. The size of the unpruned basis sets was 108 000. A static culling threshold of 10^{-10} was used in each calculation.

what at the step [Fig. 6(b)]. As the wave packet gets accelerated toward the surface by the attractive potential, it spreads out in position space, and includes higher components of momentum. This apparently leads to a net increase in the number of basis functions needed to describe the wave packet. The number of culled basis functions begins to increase again as the high momentum components start to be absorbed at around 2000 a.u. of time [Fig. 6(a)]. By the end of the propagation, the pruned basis set for the wave packet at the step includes very few functions [Fig. 6(b)], because most of the wave packet has been absorbed, and the remaining coefficients are small, most falling below the culling threshold.

C. Three-dimensional reactive scattering

The three-dimensional (3D) PES, for reaction at the step of the Pt(211) surface, includes both nonactivated reaction sites and repulsive sites where no reaction is possible. Figure 7 presents the results for dynamical pruning of PSL and DVR basis sets for this system. Figure 7(a) shows that both basis sets accurately represent the wave packet, with the normalized overlap only deviating significantly from unity after the norm is small.

Both basis sets are around an order of magnitude smaller at each point in time than the primitive direct-product basis set [Fig. 7(b)]. In this case, the DVR and PSL basis sets are

FIG. 8. Scaling of basis set sizes with dimensionality. The results are presented for dynamically pruned DVR bases (solid red circles), statically pruned DVR bases (open red circles), dynamically pruned PSL bases (solid blue squares), statically-pruned PSL bases (open blue squares), and the primitive direct-product bases (black crosses). Each series of points has been fit with a straight line, corresponding to exponential scaling. The method used to establish basis set sizes is described in the text.

very similar in size, though the DVR basis does include significantly more basis functions at longer times, when the wave packet is more spread out in position space. In the PSL basis, this spread is apparently compensated by narrowing in momentum, which arises due to high kinetic energy components of the wave packet getting absorbed.

D. Scaling with dimension

In the previous section, the DVR and PSL dynamically pruned (DP) basis sets were similar in size, posing the question of whether the PSL basis set has any inherent advantages. Given that the potential energy operation is easier to perform in the DVR basis, the PSL basis needs to be significantly smaller to warrant its use. For the 3D calculations, that was not the case [Fig. 7(b)].

To address this issue, we now consider the scaling of basis size with dimension in more detail, including up to four dimensions. Figure 8 shows time-averaged basis set sizes as a function of dimensionality. In order to make it possible to fairly compare different basis sets and dimensionalities, a dynamic culling threshold was used, as described in Sec. II. Moreover, the relative culling threshold ϵ_{rel} was adjusted such that the time-averaged error in the wave packet was approximately the same in each calculation. The error was defined as

$$e = \frac{\int_0^T |1 - \rho(t)| dt}{T}, \quad (10)$$

with T , the time over which the averaging occurs, taken to be 3500 a.u. For each basis set and dimensionality, a series of calculations was performed for different values of ϵ_{rel} . The values of ϵ_{rel} used were 10^{-5} , 10^{-6} , 10^{-7} , 10^{-8} , and 10^{-9} . The results of the calculation for which the error (e) was closest to 5×10^{-4} are presented in Fig. 8.

TABLE I. Parameters used in the curve fits shown in Fig. 8. The parameters are defined by Eq. (11).

Basis set	Prefactor (a)	Base (n)
Primitive	1.43	48.99
Statically pruned DVR	0.94	35.76
Dynamically pruned DVR	0.23	37.41
Statically pruned PSL	14.81	14.06
Dynamically pruned PSL	11.85	9.41

Figure 8 shows that the DP bases are significantly smaller than the corresponding primitive direct-product bases. For example, for the 4D calculations, the DP-DVR basis set is only around 5% of the size of the primitive basis, and the DP-PSL basis is around 1% the size. Put another way, even in a system with only four DOFs, at least 99% of the functions in the direct-product basis set are irrelevant to the wave packet at any point in time.

Figure 8 also presents the results for static pruning schemes. Static pruning involves culling basis functions for the duration of the calculation, usually according to an energy-based criterion. For example, it is common to discard DVR points for which the potential energy falls way outside the range of interest. In this particular case, we have taken each statically pruned (SP) basis set to be the union of all functions active in the corresponding DP set at any point during a propagation. In other words, any basis function that appears in the DP basis set is added to the SP set. This scheme could be considered an optimal means of generating a static basis, because only functions that are relevant to the wave packet at some point in time are included. (This is only possible with the benefit of hindsight.)

Figure 8 shows that dynamical pruning can be as much as an order of magnitude more effective at reducing basis set size than static pruning. The plots for the SP-PSL and DP-PSL bases run almost parallel, with the static curve shifted upward. The same is true of the DVR bases.

The plots in Fig. 8 are fit with lines corresponding to exponential growth,

$$N(D) = an^D, \quad (11)$$

where N is the number of functions in the basis set, D is the dimensionality, and a and n are constants. Table I gives the values of the prefactor (a) and base (n) for each plot.

It should be clear from the values in Table I, and the slopes of the fits in Fig. 8, that the PSL basis sets scale much more favorably than either the primitive or DVR bases. Even though the DP-DVR basis is approximately the same size as the DP-PSL basis for the 3D system, as noted earlier, the scaling of the DP-PSL basis means that in the 4D system it is already less than a fifth the size of the DP-DVR basis set. Increasing the dimensionality beyond 4D would only accentuate this inequality.

The largest wave packet computations performed today typically include no more than six DOFs. Recent 6D wave packet calculations on the $\text{H}_2 + \text{Pt}(211)$ PES used here required on the order of 10^9 basis functions.⁴³ Using the parameters from Table I to extrapolate to six dimensions, we

arrive at a primitive basis size of 20×10^9 , which is somewhat high. The 6D wave packet calculations adopted various measures to reduce basis size, including using 2D spherical harmonic basis functions for rotational degrees of freedom, and implementing a projection operator formalism to better tailor the basis to particular regions of space, such as the asymptotic region.⁴³ These measures explain the smaller than expected number of basis functions utilized in the 6D computations.

The extrapolated size of the DP-PSL basis set for six DOFs is 8.2×10^6 , only 0.041% as large as the extrapolated primitive basis set. Restated, less than one in every thousand primitive basis functions is relevant to the wave packet at any given point in time. This is quite a shocking statistic, and demonstrates how much there remains to be gained in time-dependent wave packet methods through algorithmic improvements. The DP-DVR basis set is less favorable than the DP-PSL set, but extrapolation predicts that it would only be around 3.2% the size of the primitive basis.

Close attention to the data in Fig. 8 reveals that the DP-PSL may actually scale less than exponentially in this case. In particular, the point representing the basis size in the 4D system lies considerably under a line drawn through the points for 2D and 3D systems. This is despite the fact that the primitive and DVR basis sets exhibit the opposite trend.

Recent studies of Poirier³⁴ and Poirier and Salam^{36,37} have established that subexponential scaling is possible in quantum dynamics when using PSL basis functions. (Dawes and Carrington have also observed subexponential scaling with PSL basis sets.^{3,39}) However, their analysis was based on the assumption that the maximum energy for which the results are sought is independent of dimensionality, which is only reasonable in a time-independent eigensystem calculation. In time-dependent scattering calculations, it is more reasonable to assume that the translational collision energy is constant with respect to dimensionality, and that the total energy thus increases with dimension. On this basis, we had anticipated exponential scaling, and Fig. 8 thus comes as something of a surprise.

One possible explanation is that the proportion of initial translational energy available to each degree of freedom decreases on average with dimension. For example, for a 2D system with a vibration, half of the translational energy could be expected to redistribute into the vibrational mode at some point during the reaction. For a 3D system, only a third of the translational energy could be expected to distribute into each nonscattering coordinate mode. Less energy for any given node corresponds to fewer basis functions, so the scaling may be subexponential for low dimensions.

An equation that takes these considerations into account is

$$N(D) = a(n + 1/n)^D. \quad (12)$$

It is not that difficult to arrive at this form. If we assume that E_0 is the initial energy of a vibrational mode, and E_s is the energy initially in the scattering DOF, then the energy available to the vibration is given by $E_0 + E_s/n$. The accessible area in phase space of a vibrational mode is proportional to its energy, so the number of basis functions required to rep-

resent the wave packet in that mode may actually decrease with dimension roughly in proportion to $1 + 1/n$. Accounting for all modes leads to the scaling law given in Eq. (12).

Although Eq. (12) is subexponential, it approaches exponential scaling for large n , albeit with a smaller exponent than would be measured for low dimensionality. It thus seems possible that the results for the DP-PSL basis set in Fig. 8 are representative of true subexponential scaling, though without more data this cannot be confirmed. And even if it were, it would seem likely that for higher dimensionality the scaling would become exponential.

E. Possible pruning schemes

In this paper, we have not endeavored to present a practical scheme for dynamical pruning, but have instead chosen to test the effectiveness of the approach in realistic scattering calculations. The implementation used for these tests was not efficient, and could only cull basis functions with the benefit of hindsight, i.e., at each step the wave packet was propagated in an unpruned basis set, and only then were small coefficients zeroed. This is, of course, not a practical algorithm for dynamical pruning of basis sets.

A number of possible approaches could be taken to determine which localized basis functions should be culled *before* a given time step. One would be to set up a pruned basis set for the known initial state, and then to include an extra “shell” of basis functions around the occupied region before each time step. This shell would allow unoccupied functions to become occupied during the course of a propagation step, and a scan of the basis set at the conclusion of the step could remove any functions that become unoccupied. The thickness of the shell would need to be such that the wave packet could not move beyond its boundary in a single time step.

Another family of methods could fall under the heading of “guided pruning.” In this scenario, a computationally cheaper and less accurate method, such as classical trajectory propagation—or even unconverged quantum wave packet propagation—could be used to determine the occupied region of phase space at each point in time. The proximity of an ensemble of classical trajectories to the set of all PSL basis functions at any point in time would determine which of those functions were included in the active basis set.

F. The potential energy operation

The last issue we would like to address is that of the potential energy operation. This is one of the more difficult aspects of quantum dynamics calculations, because the potential typically cannot be factorized into reduced-dimensional operators, as can the kinetic energy terms, and can thus become very expensive. One implementation of the MCTDH method, for example, requires that the PES be refitted with a particular product form, which then can be factorized and evaluated more cheaply.²¹ However, this approach does have its limitations.

DVRs are popular because the potential energy operator is diagonal, and can be evaluated via a cheap pointwise multiplication. In contrast, the PSL functions are less spatially localized, and the potential operator matrix is not diagonal.

Nonetheless, the PSL functions are quite localized, and the potential matrix should be sparse in practice. It should be possible either to evaluate the potential matrix in the PSL basis by transforming it from the DVR basis representation where it is diagonal, or by transforming the wave packet coefficients to the DVR representation each time step and applying the potential operator before transforming back to the PSL representation again. The matrix and coefficient transforms in question would be sparse, with each PSL function only overlapping a handful of DVR points. Lombardini and Poirier have recently investigated a similar quadrature-based scheme in time-independent calculations with wavelet bases.³⁸

IV. CONCLUSIONS

We have investigated the effectiveness of dynamically pruning multidimensional basis sets of localized functions in time-dependent wave packet calculations. Included were functions optimally localized in position space (i.e., DVR functions) and phase space (i.e., PSL functions). PSL basis functions were generated using the technique of simultaneous diagonalization.

Time-dependent wave packet scattering calculations of varying dimension were performed for H₂ dissociatively adsorbing on the Pt(211) stepped surface. The PES used was chemically realistic, based on a fit to density functional calculations. Wave packet calculations showed that a direct-product basis set, as is commonly used, can be extremely inefficient. For example, it was predicted that in a six-dimensional wave packet calculation, less than one in every thousand basis functions is relevant to the propagation of the wave packet at any given point in time.

Dynamical pruning, in which the active set of basis functions used at each point in time can vary, was found to reduce basis set size by as much as an order of magnitude over statically pruned sets, where the selection of basis functions does not vary. Although the dynamically pruned DVR basis set was smaller than the phase-space localized (PSL) basis in 2D calculations, the PSL basis set scaled much more favorably with dimension, and in 4D calculations was around a factor of 5 smaller. We even observed that the scaling of the PSL basis may be subexponential for low dimensions, and offered a possible explanation based on energy redistribution arguments.

In this study, we have not presented an algorithm for dynamical pruning of localized basis sets, but have quantified the potential benefits should a practical method be developed. We hope that this reemphasizes the inefficiency of direct-product basis sets, and motivates the search for alternatives.

ACKNOWLEDGMENTS

We would like to thank Bill Poirier and Tucker Carrington, Jr. for helpful correspondence, and Roar Olsen for constructing the potential energy surface of H₂+Pt(211).

- ¹G. J. Kroes and M. F. Somers, *J. Theor. Comput. Chem.* **4**, 493 (2005).
- ²D. Y. Wang, *J. Chem. Phys.* **123**, 194302 (2005).
- ³R. Dawes and T. Carrington, *J. Chem. Phys.* **122**, 134101 (2005).
- ⁴J. Echave and D. C. Clary, *Chem. Phys. Lett.* **190**, 225 (1992).
- ⁵D. Neuhauser and M. Baer, *J. Chem. Phys.* **91**, 4651 (1989).
- ⁶D. T. Colbert and W. H. Miller, *J. Chem. Phys.* **96**, 1982 (1992).
- ⁷R. B. Gerber, V. Buch, and M. A. Ratner, *J. Chem. Phys.* **77**, 3022 (1982).
- ⁸R. E. Wyatt, *J. Chem. Phys.* **117**, 9569 (2002).
- ⁹K. H. Hughes and R. E. Wyatt, *Chem. Phys. Lett.* **366**, 336 (2002).
- ¹⁰K. H. Hughes and R. E. Wyatt, *Phys. Chem. Chem. Phys.* **5**, 3905 (2003).
- ¹¹K. H. Hughes, *J. Chem. Phys.* **122**, 074106 (2005).
- ¹²D. Huber and E. J. Heller, *J. Chem. Phys.* **89**, 4752 (1988).
- ¹³D. Huber, S. Ling, D. G. Imre, and E. J. Heller, *J. Chem. Phys.* **90**, 7317 (1989).
- ¹⁴D. V. Shalashilin and M. S. Child, *J. Chem. Phys.* **113**, 10028 (2000).
- ¹⁵D. V. Shalashilin and M. S. Child, *J. Chem. Phys.* **115**, 5367 (2001).
- ¹⁶D. V. Shalashilin and M. S. Child, *J. Chem. Phys.* **114**, 9296 (2001).
- ¹⁷D. V. Shalashilin and M. S. Child, *J. Chem. Phys.* **119**, 1961 (2003).
- ¹⁸D. V. Shalashilin and M. S. Child, *Chem. Phys.* **304**, 103 (2004).
- ¹⁹D. V. Shalashilin and M. S. Child, *J. Chem. Phys.* **121**, 3563 (2004).
- ²⁰H. D. Meyer, U. Manthe, and L. S. Cederbaum, *Chem. Phys. Lett.* **165**, 73 (1990).
- ²¹H. D. Meyer and G. A. Worth, *Theor. Chem. Acc.* **109**, 251 (2003).
- ²²H. B. Wang and M. Thoss, *J. Chem. Phys.* **119**, 1289 (2003).
- ²³T. Wu, H. J. Werner, and U. Manthe, *Science* **306**, 2227 (2004).
- ²⁴C. Crespos, H. D. Meyer, R. C. Mowrey, and G. J. Kroes, *J. Chem. Phys.* **124**, 074706 (2006).
- ²⁵R. van Harreveld and U. Manthe, *J. Chem. Phys.* **121**, 3829 (2004).
- ²⁶J. Light, I. Hamilton, and J. Lill, *J. Chem. Phys.* **82**, 1400 (1985).
- ²⁷S. T. Ali, J. P. Antoine, J. P. Gazeau, and U. A. Mueller, *Rev. Math. Phys.* **7**, 1013 (1995).
- ²⁸L. M. Andersson, *J. Chem. Phys.* **115**, 1158 (2001).
- ²⁹L. M. Andersson, J. Aberg, H. O. Karlsson, and O. Goscinski, *J. Phys. A* **35**, 7787 (2002).
- ³⁰A. Graps, *IEEE Comput. Sci. Eng.* **2**, 50 (1995).
- ³¹J. L. Calais, *Int. J. Quantum Chem.* **58**, 541 (1996).
- ³²J. P. Modisette, P. Nordlander, J. L. Kinsey, and B. R. Johnson, *Chem. Phys. Lett.* **250**, 485 (1996).
- ³³M. Kovacic, K. Najzar, and J. Horacek, *Czech. J. Phys.* **52**(1), 41 (2002).
- ³⁴B. Poirier, *J. Theor. Comput. Chem.* **2**, 65 (2003).
- ³⁵D. J. Kouri, M. Papadakis, I. Kakadiaris, and D. K. Hoffman, *J. Phys. Chem. A* **107**, 7318 (2003).
- ³⁶B. Poirier and A. Salam, *J. Chem. Phys.* **121**, 1690 (2004).
- ³⁷B. Poirier and A. Salam, *J. Chem. Phys.* **121**, 1704 (2004).
- ³⁸R. Lombardini and B. Poirier (unpublished).
- ³⁹R. Dawes and T. Carrington, *J. Chem. Phys.* **124**, 054102 (2006).
- ⁴⁰R. A. Olsen, D. A. McCormack, and E. J. Baerends, *Surf. Sci.* **571**, L325 (2004).
- ⁴¹M. Luppi, D. A. McCormack, R. A. Olsen, and E. J. Baerends, *J. Chem. Phys.* **123**, 164702 (2005).
- ⁴²D. A. McCormack, R. A. Olsen, and E. J. Baerends, *J. Chem. Phys.* **122**, 194708 (2005).
- ⁴³D. A. McCormack, R. A. Olsen, and E. J. Baerends (unpublished).
- ⁴⁴O. Sharafeddin and J. Z. H. Zhang, *Chem. Phys. Lett.* **204**, 190 (1993).
- ⁴⁵M. D. Feit, J. A. Fleck, and A. Steiger, *J. Comput. Phys.* **47**, 412 (1982).
- ⁴⁶A. Vibok and G. G. Balint-Kurti, *J. Phys. Chem.* **96**, 8712 (1992).
- ⁴⁷R. Dawes and T. Carrington, *J. Chem. Phys.* **121**, 726 (2004).
- ⁴⁸J. F. Cardoso and A. Souloumiac, *SIAM J. Matrix Anal. Appl.* **17**, 161 (1996).
- ⁴⁹A. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in Fortran 77*, 2nd ed. (Cambridge University Press, New York, 1992).
- ⁵⁰R. van Harreveld and U. Manthe, *J. Chem. Phys.* **123**, 124706 (2005).