

Combining embedded mean-field theory with linear-scaling density-functional theory

Joseph C. A. Prentice, Robert J. Charlton, Arash A. Mostofi, and Peter D. Haynes

Department of Materials, Department of Physics, and the Thomas Young Centre for Theory and Simulation of Materials, Imperial College London, London SW7 2AZ, United Kingdom

Abstract

We demonstrate the capability of embedded mean field theory (EMFT) within the linear-scaling density-functional theory code ONETEP, which enables DFT-in-DFT quantum embedding calculations on systems containing thousands of atoms at a fraction of the cost of a full calculation. We perform simulations on a wide range of systems from molecules to complex nanostructures to demonstrate the performance of our implementation with respect to accuracy and efficiency. This work paves the way for the application of this class of quantum embedding method to large-scale systems that are beyond the reach of existing implementations.

1 Introduction

The behaviour of complex and heterogeneous molecular and materials systems is often predominantly governed by physics and/or chemistry occurring within one or more smaller ‘active’ sub-regions of the system. This does not mean, however, that the rest of the system is not relevant – indeed, it often interacts significantly with the region of interest, with important consequences for the behaviour and properties of the system as a whole. Examples include host-guest systems (e.g., guest molecules within metal-organic frameworks¹⁻³),

molecules in solution,⁴⁻⁶ defects in crystalline materials,⁷⁻⁹ and active sites in enzymes.¹⁰⁻¹² It is therefore important to include the effect of the ‘environment’, or at least large portions of it, when modelling such systems computationally. This raises the challenge of simulating a very large system, consisting of the active region together with its environment, at a level of theory and accuracy that is sufficient to describe the active region correctly – a level likely to be computationally unfeasible on the scale of the whole system.

One solution to this problem is to use an embedding scheme. In such schemes, the active region is treated at a higher, more accurate level of theory, whilst the environment is treated at a lower, less demanding level of theory. This aims to provide both high accuracy and reduced cost when compared to a full system calculation at the higher level. A major issue that must be considered carefully when designing such a scheme, however, is how to treat the interactions between the two regions.

Some well-known embedding approaches, such as QM/MM¹³ and ONIOM,¹⁴ have been used successfully for many years,¹⁵⁻¹⁹ but contain weaknesses that limit their applicability. A significant amount of knowledge about the electronic structure of the system is required to partition the system sensibly, and the number of electrons in each region must be fixed, which limits use to systems where this assignment is physically meaningful and does not change.¹⁹ Both QM/MM and ONIOM require artificial termination of any bonds that traverse the boundary between regions, and both typically treat the environment, and the interactions with the environment, at a classical or empirical level of theory, neglecting the quantum nature of these interactions.^{19,20}

Quantum embedding schemes allow calculations to go beyond a classical treatment of the environment, so that both the active region and the environment are treated using quantum mechanics.²¹ Many quantum embedding schemes are based on density-functional theory (DFT), at least in part, due to the fact that DFT is extremely widely used across many fields, thanks to its combination of relatively low cost and high accuracy.²² One of the main ways to improve the accuracy of a given DFT calculation is by using a different approximation

to the exchange-correlation functional – hybrid functionals, for example, improve agreement with experimental data significantly relative to local and semi-local functionals^{23,24} such as the local density approximation (LDA)^{25,26} or the generalised gradient approximation (GGA).^{27,28} They do this by including a fraction of exact exchange, which makes hybrid functional calculations more accurate, but also more computationally expensive than semi-local DFT calculations. The accuracy of a DFT calculation can also be improved by using a larger or more appropriate set of basis functions, also at increased computational cost. This hierarchy of both functionals and basis sets suggests that DFT-in-DFT embedding schemes, where the high level of theory is hybrid DFT (possibly with a large basis set), and the low level is semi-local DFT (possibly with a smaller basis set), could be powerful computational tools.²¹ Schemes that embed extremely accurate wavefunction-based (WF) methods, such as coupled cluster methods,^{29,30} within DFT (WF-in-DFT) also exist, and are potentially a similarly powerful tool. However, the computational cost of WF methods is typically greater, limiting their use to smaller active regions than those accessible with hybrid DFT. For this reason, in this work we focus only on DFT-in-DFT embedding.

As mentioned above, one of the major difficulties when devising an embedding scheme is dealing with the interactions between the regions. In particular, in quantum embedding it is vital to ensure that the orbitals of the two systems remain orthogonal, to satisfy the Pauli exclusion principle. This requirement is enforced in DFT-in-DFT embedding by the appearance of a non-additive term in the kinetic energy, which is an unknown functional of the density.^{21,31,32} This non-additive term has been treated in several different ways. One option is to approximate it, either using an analytical^{31–33} or a numerically fitted^{34–36} form. Another option is to enforce orthogonality manually to avoid the non-additive term ever appearing: this can be done either by construction,³⁷ or by projecting out the environment orbitals.^{38–40} Each of these options has its own strengths and weaknesses, and some are more suitable than others for a given situation.

Recently, another quantum embedding method that avoids these issues has been pro-

posed, known as embedded mean-field theory (EMFT).²⁰ This method endeavours to avoid some of the complications of other quantum embedding methods by partitioning the atom-centred basis set directly. This means that EMFT is able to remain a mean field theory, like standard DFT, at all times, avoiding the issue of orthogonality altogether and providing a conceptually simple framework for combining different levels of theory. This, together with previous work demonstrating its utility,^{20,41–45} makes it a very appealing scheme to apply for DFT-in-DFT embedding.

In this work, we present a novel implementation of EMFT in a linear-scaling DFT code, ONETEP,⁴⁶ which is distinct from previous codes where EMFT has been implemented. ONETEP uses a minimal atom-centred basis set, where each basis function is optimised within its own local environment. These functions are strictly localised within a given radius, making the relevant matrices sparse and enabling linear scaling of computational cost with the size of the system. This implementation therefore moves EMFT outside of the realm of Gaussian basis sets, and demonstrates its utility for arbitrary localised basis functions and large linear-scaling DFT calculations in general. We test our implementation of EMFT on several example systems, which demonstrate that this formalism can be used on both molecular and solid state systems. In particular, we are able to apply EMFT to a system containing nearly 3000 atoms, significantly larger than any such calculation attempted before, and obtain accurate results at a significantly lower computational cost than the full high-level calculation.

The work is organised as follows. In Section 2, we give an overview of the theory of EMFT and the advances that have been made in previous work. In Section 3, we provide details of the implementation of EMFT in ONETEP. In Section 4, we give the results of testing our implementation on a variety of systems, including both molecular and periodic systems. Finally, in Section 5, we summarise our work and form conclusions.

2 Embedded mean field theory

EMFT rests on the idea of separating the basis set into two subsets, one associated with the active region, A , and one with the environment, B .²⁰ In an atom-centred basis set, such as Gaussians or the more general basis set of ONETEP, this is quite simple, as an atom, and thus their associated basis functions, can easily be assigned to a given region. Extending the EMFT formalism to allow for multiple regions or for separating one of the regions into a further set of subregions is straightforward in principle, but we do not consider it in this work.

This partitioning of the basis set means that we can separate out the single-particle density matrix (or kernel when written in terms of the basis functions) K ⁴⁶ into blocks:

$$K = \begin{pmatrix} K^{AA} & K^{AB} \\ K^{BA} & K^{BB} \end{pmatrix}. \quad (1)$$

We can partition the overlap matrix, which gives the overlaps between the various basis functions, similarly:

$$S = \begin{pmatrix} S^{AA} & S^{AB} \\ S^{BA} & S^{BB} \end{pmatrix}. \quad (2)$$

We now consider the two (mean-field) levels of theory we will be applying to the two regions, where the energy is a function of K . At a general level, we can separate the energy out into parts arising from one-electron and two-electron terms:

$$E[K] = \text{tr}[KH_0] + G[K]. \quad (3)$$

Here the first term contains the one-electron terms arising from the core Hamiltonian H_0 , whilst the second contains all two-electron terms, such as Hartree and exchange-correlation energy contributions. In this work, as we will be focusing on DFT-in-DFT embedding, we will assume that the higher and lower levels of theory only differ in their treatment of the

two-electron terms, giving G^{high} and G^{low} respectively. However, the formalism also applies if there is also a difference in the one-electron terms. We can now write the EMFT energy functional as

$$E^{\text{EMFT}}[K] = \text{tr}[KH_0] + G^{\text{low}}[K] + (G^{\text{high}}[K^{AA}] - G^{\text{low}}[K^{AA}]) . \quad (4)$$

This means that we calculate the energy of the full system at the low level, before then applying a high level correction based only on the density matrix in the active region. The off-diagonal blocks of the density matrix are therefore treated at the lower level of theory. The ground state energy is then obtained by minimising E^{EMFT} with respect to K subject to the constraints of normalisation and idempotency. The theory remains mean-field throughout, and it is simple to calculate forces²⁰ and to extend the theory to obtain excitations within a time-dependent EMFT (TD-EMFT) framework.⁴³

The most obvious way of gaining significant computational advantages from the EMFT framework is by treating the environment with a semi-local functional, but the active region with a hybrid functional – that is, a functional containing a fraction of exact exchange. To do this, the exact exchange contribution to the energy of the active region must be calculated. However, because we have partitioned the system, there are several different ways this could be done, depending on whether exchange between regions is included or not. The simplest and least expensive method is the EX0 method, which only includes exchange interactions within the active region, giving an exchange energy of^{20,42}

$$E_{\text{EX0}} = -\frac{1}{4} \sum_{\alpha\beta\gamma\delta \in A} K_{\alpha\beta}^{AA} (\phi_\alpha\phi_\delta|\phi_\beta\phi_\gamma) K_{\delta\gamma}^{AA} , \quad (5)$$

where $\alpha, \beta, \gamma, \delta$ run over the basis functions $\{\phi_i\}$ associated with region A . $(\phi_\alpha\phi_\delta|\phi_\beta\phi_\gamma)$ represents an electron repulsion integral (ERI), whose implementation in ONETEP is described

elsewhere:⁴⁷

$$(\phi_\alpha \phi_\delta | \phi_\beta \phi_\gamma) = \iint \frac{\phi_\alpha^*(\mathbf{r}) \phi_\delta(\mathbf{r}) \phi_\beta(\mathbf{r}') \phi_\gamma^*(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3\mathbf{r} d^3\mathbf{r}' . \quad (6)$$

Alternatively, the inter-region exchange interaction can be included via the EX1 method,⁴² whereby the interaction is averaged symmetrically between the subsystems, giving an energy of

$$E_{\text{EX1}} = E_{\text{EX0}} - \frac{1}{4} \sum_{\alpha\beta \in A; \gamma\delta \notin A} K_{\alpha\beta}^{AA} (\phi_\alpha \phi_\delta | \phi_\beta \phi_\gamma) K_{\delta\gamma} . \quad (7)$$

The cost of calculating the exact exchange energy is dominated by evaluating the ERIs, which scales quartically with the number of basis functions. Thus the EX0 and EX1 methods scale as $O(N_A^4)$ and $O(N_A^2 N^2)$ respectively, where N_A and N are the number of basis functions in subsystem A and the full system, such that for $N_A \ll N$ EX0 is significantly more efficient than EX1 and previous work has shown that it is of comparable accuracy to EX1.⁴² Including the full inter-region exchange simply amounts to replacing the factor of $\frac{1}{4}$ in Eq. (7) with $\frac{1}{2}$, but this has the same cost as the EX1 scheme, and has been previously found to produce less accurate results.⁴² For these reasons, we use the EX0 method throughout this work.

An issue that has previously been noted with EMFT calculations is that in some cases, where the two levels of theory are significantly different (e.g., hybrid DFT in semi-local DFT embedding), EMFT converges to an unphysical solution with much lower energy than that produced by either level of theory.⁴² In these situations, the trace of the diagonal blocks of the density matrix, $\text{tr}(K^{AA}S^{AA}) + \text{tr}(K^{BB}S^{BB})$ becomes extremely large and positive – implying that, in order to lower the overall energy, more and more of the electron density localises in the two subsystems, even beyond the number of electrons actually present in the system. Conservation of electron number $N_e = \text{tr}(KS)$ is still mathematically ensured by this being balanced by large negative values arising from the off-diagonal blocks, $\text{tr}(K^{BA}S^{AB}) + \text{tr}(K^{AB}S^{BA})$, but this behaviour is clearly unphysical. This phenomenon was reported in previous work,⁴² and was also observed in our own preliminary calculations. One way to fix this issue is to force the off-diagonal blocks of the overlap matrix to be zero,

i.e., orthogonalising the basis functions of environment region with respect to those of the active region. This block orthogonalisation procedure involves applying a transformation matrix U to the basis functions, given by⁴²

$$U = \begin{pmatrix} I^{AA} & -P^{AB} \\ 0^{BA} & I^{BB} \end{pmatrix}. \quad (8)$$

Here, we have partitioned the matrix in the same way as K and S , with I^{IJ} and 0^{IJ} representing identity and zero matrices of appropriate sizes for their blocks. $P^{AB} = (S^{AA})^{-1} S^{AB}$ is the projection matrix that performs the orthogonalisation. Using this transformation solves the problem of these unphysical solutions, at essentially no extra cost, and is included in the implementation we detail here. Note that this block orthogonalisation procedure can be considered as a pre-processing step in the construction of the local orbitals and does not require reference to the molecular orbitals. This is distinct from other embedding techniques that involve explicitly imposing orthogonality between the Kohn-Sham orbitals of the subsystems.^{37–40}

3 Implementation in ONETEP

We now give details of our implementation of EMFT in ONETEP. ONETEP, like Gaussian basis set DFT codes, uses atom-centred functions as a basis in which to express the density matrix; unlike most other DFT codes, however, ONETEP optimises these basis functions to take into account the local environment of the atom they are centred on.⁴⁶ This means that the energy of the system is minimised not only with respect to the density kernel, but also with respect to the form of the basis functions themselves. As these basis functions are not required to be orthogonal to one another, and are strictly localised within a given radius around their atom, they are referred to as non-orthogonal generalised Wannier functions (NGWFs). Using an optimised basis set such as this allows fewer basis functions to be used

for the same accuracy, reducing the size and therefore cost of the calculation. As the number and radius of the NGWFs associated with a particular atom can be easily adjusted, differently sized basis sets can be used for different regions, a form of embedding discussed in Sec. 1. In this work, however, we focus on the effect of using different exchange-correlation functionals for different regions, as the optimisation of the basis set means that only relatively modest basis set sizes are necessary.

In order to describe the NGWFs, they must be expanded in a set of underlying basis functions. In ONETEP, these are periodic cardinal sine (psinc) functions, centred on the points of a grid.⁴⁶ The size of this grid is controlled by a cutoff energy, which is related to the quantity of the same name in the plane-wave pseudopotential method – the larger the cutoff energy, the finer the grid. ONETEP also samples the Brillouin zone of the system at the Γ -point only – an excellent approximation for the large systems ONETEP was designed for.

Equations (1) and (2) show how the density kernel and overlap matrix can be partitioned. In order to use the existing machinery for optimising the kernel and NGWFs within ONETEP, a similar expression is required for the Hamiltonian. This is given by

$$H^{\text{EMFT}} = \begin{pmatrix} H_{\text{high}}^{AA} & H_{\text{low}}^{AB} \\ H_{\text{low}}^{BA} & H_{\text{low}}^{BB} \end{pmatrix}. \quad (9)$$

Each block is created by calculating the expectation value of the appropriate Hamiltonian operator between an appropriate set of basis functions on either side – for example, for the high level Hamiltonian

$$H_{\text{high};\alpha\beta}^{IJ} = \langle \phi_{\alpha}^I | \hat{H}_{\text{high}} | \phi_{\beta}^J \rangle, \quad (10)$$

where I, J label the regions, and $|\phi_{\alpha}^I\rangle$ is a NGWF from region I . A similar equation applies for the low level Hamiltonian. The Hamiltonian operators themselves are defined as

$$\hat{H}_{\text{high}} = \hat{T} + \hat{V}_{\text{ext}} + \hat{V}_{\text{Hartree}} + \hat{V}_{\text{XC}}^{\text{high}}, \quad (11)$$

with a similar definition again for the low level Hamiltonian. The kinetic energy \hat{T} and external potential \hat{V}_{ext} (which includes local and non-local pseudopotential contributions) together correspond to the single-electron core Hamiltonian H_0 in Eq. (3); the Hartree potential \hat{V}_{Hartree} and the exchange-correlation potential \hat{V}_{XC} together correspond to the two-electron term G in Eq. (3). Only \hat{V}_{XC} changes as the level of theory is changed. Exact exchange contributions are a special case – they are calculated as discussed in Sec. 2 and added on to the energy and its gradient afterwards. This allows us to treat \hat{V}_{XC} as local in the following discussion. One other thing to note is that the off-diagonal blocks in the Hamiltonian are treated at the lower level of theory.

The Hamiltonians, defined as in Eq. (11), are then constructed as follows. Firstly, the total electron density $n(\mathbf{r})$ is constructed from the full system NGWFs and kernel, and from this $V_{\text{XC}}^{\text{low}}(\mathbf{r})$ is calculated. Next, the active subsystem density $n^A(\mathbf{r})$ is constructed from the active region NGWFs and kernel block K^{AA} , from which the potentials $V_{\text{XC}}^{\text{high},A}(\mathbf{r})$ and $V_{\text{XC}}^{\text{low},A}(\mathbf{r})$ are calculated. The high-level EMFT potential is then calculated as

$$V_{\text{XC}}^{\text{high}}(\mathbf{r}) = V_{\text{XC}}^{\text{low}}(\mathbf{r}) + \left(V_{\text{XC}}^{\text{high},A}(\mathbf{r}) - V_{\text{XC}}^{\text{low},A}(\mathbf{r}) \right) . \quad (12)$$

Finally, from the potentials $V_{\text{XC}}^{\text{low}}(\mathbf{r})$ and $V_{\text{XC}}^{\text{high}}(\mathbf{r})$, we can calculate \hat{H}_{low} and \hat{H}_{high} respectively.

Given K , S and H in this way, a calculation within ONETEP may proceed as usual, using all the computational machinery already available. In particular, the covariant NGWF gradient, which provides the main information required to minimise the energy with respect to the NGWFs, is given by

$$g_\alpha = \frac{\partial E^{\text{EMFT}}}{\partial \langle \phi_A^\alpha |} = |\phi_i^A\rangle \mathcal{A}_\alpha^i + \hat{H}^{\text{high}} |\phi_i^A\rangle \mathcal{B}_\alpha^i + |\phi_i^B\rangle \mathcal{C}_\alpha^i + \hat{H}^{\text{low}} |\phi_i^B\rangle \mathcal{D}_\alpha^i . \quad (13)$$

This expression gives the gradient for NGWFs within region A ; there is a precisely analogous one for NGWFs in region B . Summation over i is assumed. The gradient is tensorial in nature owing to the non-orthogonality of the NGWFs, meaning that the metric (the overlap matrix)

is non-trivial.

Eq. (13) is similar to the general form of the NGWF gradient for non-embedding calculations,⁴⁸ but with the terms for the two regions separated out. The coefficient matrices in this expression are defined as

$$\mathcal{A} = [3LHL - 2LHLSL - 2LSLHL]^{AA} S^{AA} \quad (14)$$

$$\mathcal{B} = [3LSL - 2LSLSL]^{AA} S^{AA} = K^{AA} S^{AA} \quad (15)$$

$$\mathcal{C} = [3LHL - 2LHLSL - 2LSLHL]^{BA} S^{AA} \quad (16)$$

$$\mathcal{D} = [3LSL - 2LSLSL]^{BA} S^{AA} = K^{BA} S^{AA} , \quad (17)$$

where L is the auxiliary matrix used in the purification of the density kernel, defined as $K = 3LSL - 2LSLSL$.⁴⁹ The $[\dots]^{AA}$ notation means that the expression inside the brackets is calculated using the full matrices, and the AA block of the result is taken. It is important to note that in Eq. (13), \hat{H}^{high} only acts on NGWFs from the active region A , and \hat{H}^{low} only acts on NGWFs from the environment, B .

As mentioned above, when there is a large difference between the levels of theory used, a block orthogonalisation procedure is necessary to prevent unphysical growth of the diagonal blocks of the overlap matrix. The transformation required to orthogonalise the environment basis functions with respect to those in the active region is as shown previously in Eq. (8), but the fact that the basis functions are optimised in ONETEP presents an additional challenge. The block orthogonalisation procedure works well for kernel optimisation, as in previous work,⁴² but severely affects the optimisation of the NGWFs. Block orthogonalisation effectively adds another term into the NGWF gradient that competes with the usual terms present in the gradient. This competition is strong enough that in many cases, the NGWF gradient never converges to zero, instead stalling at a finite value. The energy of the system continues to decrease indefinitely, in a similar way to the issue with the optimisation of the density kernel that block orthogonalisation was designed to prevent. The NGWFs try to

optimise towards the region described by the level of theory with the lowest energy. This is inconsistent with a good description of the system, and as such these results are unphysical. To overcome this issue, we pre-optimize all the NGWFs in the system, as well as the kernel, at the lower level of theory, before block orthogonalising and fixing the NGWFs for an optimisation of the kernel within the EMFT framework. Whilst this implies that the NGWFs are not ideally optimised under the EMFT Hamiltonian H^{EMFT} , our results, which we will discuss in detail in the next section, show that the error introduced by using NGWFs optimised at the lower level of theory is typically only 1% of the energy difference between the low and high levels of theory; in the worst case scenario where the character of the local environment of the NGWFs in the active region significantly changes when the level of theory changes, a larger set of NGWFs may be used, still enabling high quality calculations.

Additional results further demonstrating the effect of using NGWFs optimised at different levels of theory are presented in the Supplementary Information. These results use a mixed set of NGWFs, where the NGWFs in the active region are taken from a full system high-level calculation, and the NGWFs in the environment are taken from a full system low-level calculation. No block orthogonalisation was applied, allowing the smooth transition from PBE to B3LYP to be seen clearly. As performing such calculations requires a full high-level calculation, however, negating any performance gains from embedding, we do not discuss this method further.

4 Results

We have applied our implementation of EMFT in ONETEP to several different systems. In this section, we describe each of these systems, and detail the results that demonstrate the capability of and validate our implementation.

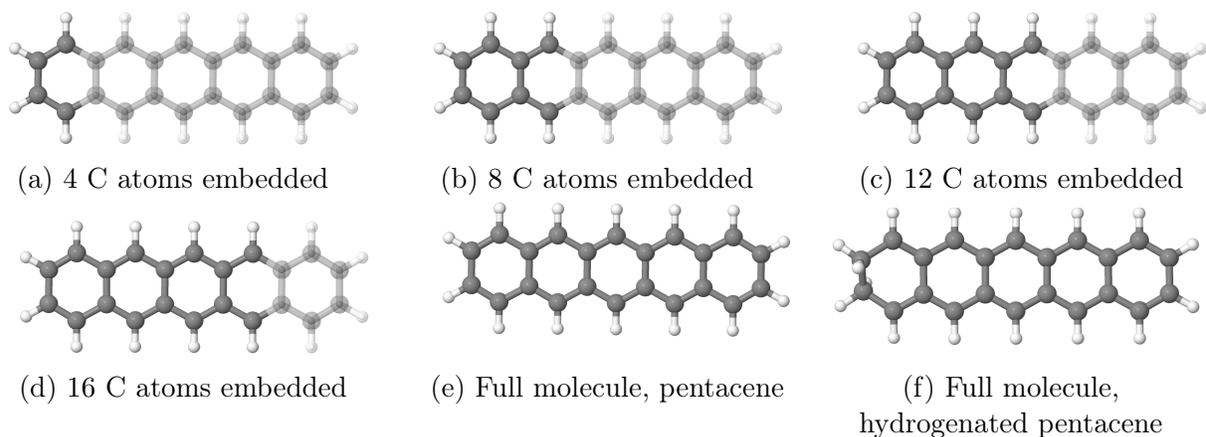


Figure 1: The different embedding regions used for the various calculations on pentacene and terminally-hydrogenated pentacene. The atoms in the active region are opaque and the atoms in the environment are translucent. a)-d) show the embedding regions used for pentacene – the regions used for hydrogenated pentacene are identical, but with the 2 additional hydrogen atoms bonded to the two left-most carbon atoms. e) and f) show the full structure of pentacene and terminally-hydrogenated pentacene, respectively. The individual figures are labelled by the number of carbon atoms included in the active region. H and C atoms are white and grey respectively. The structures are geometry optimised within ONETEP at the PBE level. Figures produced using Jmol.⁵⁰

4.1 Hydrogenation of pentacene

Pentacene is an organic molecule that has been the subject of significant research interest in recent years. As a linear acene, it consists of five fused benzene rings, giving it the chemical formula $C_{22}H_{14}$ (see Fig. 1e). Thanks to its light-absorbing and semiconducting properties, it is of interest in several different fields: it has been used as part of organic photovoltaic systems,⁵¹ organic thin-film transistors,⁵² and when doped into *para*-terphenyl, it forms the basis for a room-temperature maser.⁵³ This makes pentacene an excellent first system for our implementation of EMFT.

As relative energies, rather than total energies, are almost always the quantity of interest in DFT calculations, we will apply our implementation to both pentacene and terminally-hydrogenated pentacene – that is, pentacene where two hydrogen atoms have bonded with the two carbon atoms at one end of the molecule (see Fig. 1f). From the energy difference between these two molecules, we can calculate the terminal hydrogenation energy of pen-

tacene, ΔE_{hyd} . This reaction has been used as a test case for previous implementations of EMFT.^{20,42}

In our calculations, we use the PBE functional²⁸ as the low level of theory and the B3LYP hybrid functional⁵⁴ as the high level. We use a cutoff energy of 800 eV throughout, and all NGWF radii are set to 4.76 Å (9 bohr). The carbon and hydrogen norm-conserving pseudopotentials distributed with ONETEP were used. In all EMFT calculations, we use the procedure outlined in Sec. 3 – we optimise both the NGWFs and kernel at the low level of theory, then freeze the NGWFs and optimise the kernel alone within EMFT, using block orthogonalisation. We initially treat both molecules wholly at the PBE level of theory (a ‘full’ PBE calculation), and then increment the size of the region being treated with B3LYP until we are treating both molecules wholly at this higher level of theory. The embedded regions and their sizes are detailed in Fig. 1. To ensure the comparisons made between the different calculations are meaningful, we calculate the full molecules’ energy at the B3LYP level twice: once using NGWFs that are optimised at the PBE level as in the other calculations, and once using NGWFs that are optimised at the B3LYP level (a ‘full’ B3LYP calculation). This last calculation provides the benchmark against which all others are measured. Once we have the energy of both the pentacene and hydrogenated pentacene molecules for each embedding configuration, we calculate the terminal hydrogenation energy as

$$\Delta E_{\text{hyd}} = E_{\text{HP}} - E_{\text{P}} - E_{\text{H}_2} , \tag{18}$$

where E_{HP} , E_{P} , and E_{H_2} are the energy of terminally-hydrogenated pentacene, pentacene, and an isolated hydrogen molecule respectively. E_{H_2} is always calculated at the same level of theory as the additional hydrogen atoms in the hydrogenated pentacene, i.e. PBE for a ‘full’ PBE calculation, B3LYP with B3LYP-optimised NGWFs for a ‘full’ B3LYP calculation, and B3LYP with PBE-optimised NGWFs otherwise.

The results of these calculations are shown in Fig. 2. The first data point, where there are

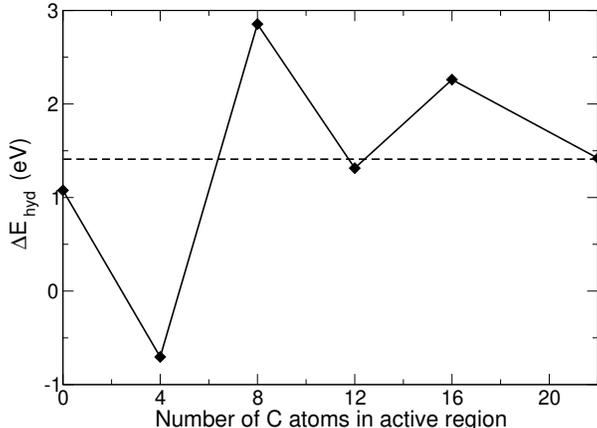


Figure 2: Terminal hydrogenation energy, ΔE_{hyd} , of a pentacene molecule as a function of the number of C atoms treated with B3LYP, with all NGWFs optimised at the PBE level of theory. The energy obtained when both the kernel *and* the NGWFs are optimised using B3LYP is marked with a dotted line. ΔE_{hyd} is calculated as shown in Eq. (18). The precise atoms included within the embedded active region for each calculation are shown in Fig. 1.

no C atoms in the active region, corresponds to a full PBE calculation. The final data point, where all C (and H) atoms are in the active region, corresponds to the B3LYP calculation using NGWFs optimised at the PBE level. The dotted line shows the energy obtained from a full B3LYP calculation, which will act as our reference result.

The first aspect of these results to note is that, when treating the whole molecule at the B3LYP level, there is practically no difference between the energy obtained with B3LYP-optimised NGWFs and PBE-optimised NGWFs. The difference in energy is 0.01 eV, which is at least an order of magnitude smaller than the energy difference between the full B3LYP calculation and any other embedding configuration. This demonstrates that the approximation of using only low level-optimised NGWFs, made as part of our block orthogonalisation implementation, is valid.

It can also be observed that the energy oscillates around the full B3LYP value, as seen in previous implementations of EMFT,²⁰ but that the amplitude of these oscillations reduces as the active region increases in size. Going from four to eight carbon atoms in the active region reduces the absolute value of the error relative to the full B3LYP calculation by 32%, from 2.12 eV to 1.44 eV. With 16 carbon atoms in the active region, the error reduces again,

by 59% relative to the error for eight carbon atoms. (The fact that the value for 12 carbon atoms is anomalously close to the full B3LYP value is almost certainly coincidental.) This demonstrates that as the size of the active region increases, the result tends towards the high-level result for the full system, as would be expected. As increasing the size of the active region will bring the result closer to the full high-level result, this means that we can vary the size of the active region to obtain the desired balance between computational cost and accuracy. Depending on the accuracy required, even relatively small active regions can give an acceptable result. From a practical standpoint, this shows the utility of our implementation of EMFT, as we can potentially obtain high-accuracy results for a fraction of the cost of a full high-level calculation. Although pentacene is a system that has been studied with previous implementations of EMFT,^{20,42} the linear scaling behaviour of ONETEP will enable much larger systems to be studied, such as those presented in the next sections.

4.2 Pentacene in *p*-terphenyl

As noted in Sec. 4.1, if a crystal of *p*-terphenyl is doped with pentacene, it can form the basis of a room-temperature maser – essentially a laser based on microwaves instead of visible light.⁵³ Although the invention of the maser actually predates the laser,^{55,56} most maser systems are only operational in vacuum, under strong magnetic fields, and at very low temperatures.^{57–59} The general introduction of a room-temperature maser will therefore make significant improvements in several fields possible, ranging from medical diagnostics to biological structure determination.⁵³

To obtain the population inversion necessary for masing, the system must be pumped using a visible light laser, which excites the pentacene molecules from their overall (singlet) ground state S_0 into their first-excited singlet state S_1 . These excited molecules then go through an intersystem crossing, via a triplet excited state T_2 , into the triplet ground state T_1 , but in a spin-selective way – the highest energy spin sub-level X is preferentially populated over the lowest energy sub-level Z , leading to a population inversion and masing

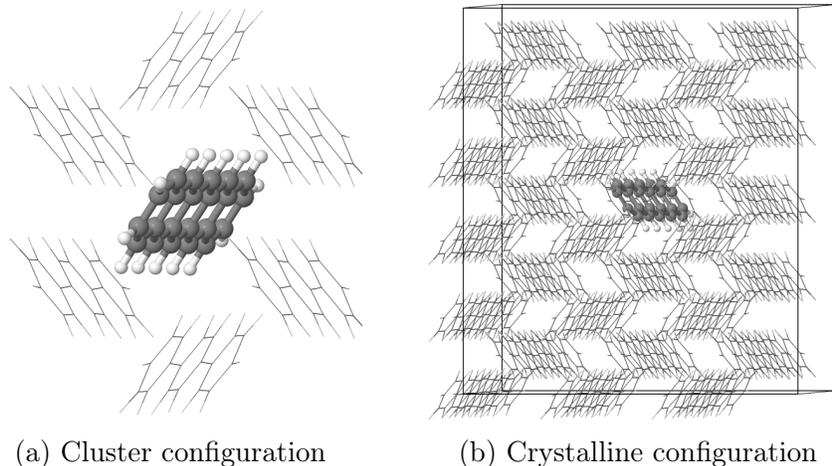


Figure 3: The two configurations used in this study for the pentacene in *p*-terphenyl system. (a) shows the cluster configuration, consisting of a pentacene molecule surrounded by six nearest neighbour *p*-terphenyl molecules. (b) shows the crystalline configuration, consisting of a $3 \times 5 \times 3$ supercell of crystalline *p*-terphenyl, with the central *p*-terphenyl replaced by a pentacene molecule. The supercell is shown at an angle slightly away from the *c*-axis for clarity, but it should be noted that three molecules appear on top of each other along the viewing direction. (b) also shows the periodic unit cell used in the calculations. In both cases, the *p*-terphenyl molecules are shown as wireframes, and the pentacene using a ball-and-stick model for ease of viewing. H and C atoms are white and grey respectively. Figures produced using Jmol.⁵⁰

behaviour.^{53,60} Non-radiative decay or phosphorescence then leads to the molecules decaying back to the overall ground state S_0 . The energy differences between the ground state S_0 and the excited states S_1 and T_1 , $\Delta E_{S_0 \rightarrow S_1}$ and $\Delta E_{S_0 \rightarrow T_1}$ respectively, are therefore of critical importance in understanding this system, the former as the alignment of S_1 with T_2 enhances the intersystem crossing, and the latter as it determines how long it takes for the pentacene molecules to decay from the T_1 state.

Previous experimental work has shown that the excitations of the pentacene molecule depend strongly on the identity of the host molecules.^{61–63} Attempts have been made to include this effect in computational calculations by treating the host as an implicit solvent,⁶⁴ but this did not solve the problem completely. Studies on similar systems in solution have shown that some of the solution/host must be treated quantum mechanically to obtain agreement with experimental data, including the nearest neighbour molecules at a minimum.^{6,65} The

requirement to include some of the host explicitly, therefore, makes this system an excellent example of an embedded system amenable to the EMFT method in ONETEP.

Here, we focus on calculating the values of $\Delta E_{S_0 \rightarrow T_1}$ and $\Delta E_{S_0 \rightarrow S_1}$, first by treating the whole system with the PBE functional, and then by defining the pentacene molecule as the active region, treated with B3LYP, whilst the environment is treated with PBE. The energy $\Delta E_{S_0 \rightarrow T_1}$ can be determined using the Δ SCF method – since the T_1 state is in fact the ground state if the electrons are restricted to be in their triplet configuration, its energy can be calculated as simply the difference between the ground state energies of the $S = 1$ and $S = 0$ spin states.

The $\Delta E_{S_0 \rightarrow S_1}$ transition energy cannot be determined directly using the Δ SCF method since S_1 does not correspond to a ground state in the singlet spin configuration. The S_1 excited state consists of a transition from the highest occupied molecular orbital (HOMO) to the lowest unoccupied molecular orbital (LUMO) of the pentacene, as indeed does T_1 . Consequently, we can calculate the singlet-triplet splitting $\Delta E_{T_1 \rightarrow S_1}$ via the virial exciton method due to Becke,⁶⁶ whereby the splitting is determined exactly as the following electron repulsion integral (ERI) for the frontier orbitals,

$$\Delta E_{T_1 \rightarrow S_1} = (\psi_H \psi_H | \psi_L \psi_L) = \iint \frac{\psi_H^*(\mathbf{r}) \psi_H(\mathbf{r}) \psi_L(\mathbf{r}') \psi_L^*(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3\mathbf{r} d^3\mathbf{r}', \quad (19)$$

where ψ_H and ψ_L are the HOMO and LUMO wavefunctions, respectively. This method has previously been applied to evaluating S_1 for single molecules, producing promising results for a multitude of systems, for example benzene where charge-transfer effects are significant.^{66,67} The molecular orbitals can be readily obtained via a one-off diagonalisation of the Hamiltonian matrix after a ground state optimisation of the NGWFs and density kernel. Then, as a post-processing step, the ERI in Eq. (19) can be computed using the same infrastructure as a regular exact-exchange calculation.⁴⁷

We conduct the Δ SCF calculations at the PBE and B3LYP-in-PBE level for three dif-

ferent geometries: vacuum pentacene, a cluster and a fully periodic crystal (see Fig. 3). In addition, the singlet-triplet splitting is determined for every $S = 1$ triplet state calculation for all geometries. A previous first-principles study of the pentacene-in-*p*-terphenyl system, using both PBE and the range-separated hybrid OT-LC ω PBE functionals, treated a cluster of a pentacene molecule and its six nearest neighbour *p*-terphenyls.⁶⁰ We use the same structure for our cluster geometry, which is shown in Fig. 3a, and we compare the results of the current work to those found in this previous study.

We also go beyond the cluster geometry, by conducting calculations on a fully periodic structure. The experimental structure of *p*-terphenyl⁶⁸ was obtained from the Cambridge Structural Database.⁶⁹ A $3 \times 5 \times 3$ supercell containing 90 molecules was then produced, and a pentacene molecule was substituted for the central *p*-terphenyl molecule, as in the cluster geometry. This structure can be seen in Fig. 3b. This structure contains 2884 atoms, an order of magnitude more than the 228 atoms included in the previous cluster calculations. Performing such a large calculation, which will include longer range interactions between the pentacene guest and the *p*-terphenyl host, demonstrates the utility of our implementation of quantum embedding methods within linear-scaling DFT – we can treat a very large environment at the low-level of theory, before making use of embedding to treat the active region more accurately. Such a calculation would have been impossible using previous implementations of EMFT.

In all calculations, a cutoff energy of 750 eV was used, and all NGWF radii were set to 4.76 Å (9 bohr). As before, we use PBE and B3LYP as the low and high levels of theory, respectively, and use the block orthogonalisation procedure. The carbon and hydrogen norm-conserving pseudopotentials distributed with ONETEP were used. In this system, all calculations used NGWFs optimised at the PBE level – no full B3LYP calculations were performed.

Table 1 shows the results for all three structures at both the PBE and B3LYP-in-PBE levels of theory. In addition to the raw excitation energies from the Δ SCF calculations and

Table 1: Energy differences between the S_0 ground state and the T_1 and S_1 excited states for pentacene, both in vacuum and doped into p -terphenyl, treated using both cluster and crystal geometries shown in Fig. 3. The effect of the p -terphenyl host on the $S_0 \rightarrow S_1$ excitation energy is also shown. All results are compared to experimental data, taken as marked from Refs. 61,70,71. For the cluster and crystal configurations, B3LYP-in-PBE refers to an EMFT calculation, with the pentacene molecule forming an active region treated with B3LYP, and the environment of p -terphenyl molecules being treated with PBE. For the vacuum configuration, B3LYP-in-PBE refers to a calculation where the NGWFs are optimised at the PBE level, before the density kernel is optimised at the B3LYP level, for consistency with the EMFT calculations. All $\Delta E_{S_0 \rightarrow T_1}$ values were calculated using the Δ SCF method. The values for $\Delta E_{S_0 \rightarrow S_1}$ were calculated using the Becke method.⁶⁶ See main text for further details.

Configuration	$\Delta E_{S_0 \rightarrow T_1}$ (eV)			$\Delta E_{S_0 \rightarrow S_1}$ (eV)		
	PBE	B3LYP-in-PBE	Exp. ⁷⁰	PBE	B3LYP-in-PBE	Exp.
Vacuum	0.885	0.942		1.835	1.895	2.31 ⁷¹
Cluster	0.879	0.934	0.86 ± 0.03	1.801	1.853	2.1 ± 0.1 ⁶¹
Crystal	0.891	0.934		1.804	1.756	

Configuration	$\Delta E_{S_0 \rightarrow S_1}$ host redshift (eV)		
	PBE	B3LYP-in-PBE	Exp.
Cluster	0.034	0.042	0.2
Crystal	0.031	0.139	

the singlet-triplet splittings, Table 1 includes the calculated redshift for both the S_1 and T_1 states, plus experimental results for comparison.

It can immediately be seen that the B3LYP-in-PBE calculations slightly overestimate $\Delta E_{S_0 \rightarrow T_1}$ in vacuum and host by 0.07 eV, whilst the PBE results are somewhat closer – this is almost certainly coincidental, however, and due to both calculations missing an effect that would reduce $\Delta E_{S_0 \rightarrow T_1}$. One possibility may be the presence of a Stokes shift in the triplet state, which is reasonably long-lived. It is likely that the energies obtained using PBE are affected by over-delocalisation, a well known issue with semi-local functionals;⁷² as the S_0 state will be more delocalised than the T_1 state due to the Pauli exclusion principle, the S_0 state will be more strongly affected by this, relatively lowering its energy and thus increasing $\Delta E_{S_0 \rightarrow T_1}$. This is borne out by the increase in $\Delta E_{S_0 \rightarrow T_1}$ as we go from the smaller cluster configuration to the larger crystal, where more delocalisation is possible. These results are consistent with previous Δ SCF calculations on the vacuum and cluster configurations,

however, which also found that hybrid functionals overestimated $\Delta E_{S_0 \rightarrow T_1}$, with PBE giving a result closer to experiment.⁶⁰ This shows that our implementation of EMFT is successfully replicating the high-level result, as it is meant to do – the fact that the high level description actually produces a result slightly further from experiment than the low level description in this system does not alter this. As expected, there is a very small redshift for $\Delta E_{S_0 \rightarrow T_1}$ from the vacuum to the cluster and crystal, amounting to only 8 meV in both cases. This indicates that the highly localised T_1 state changes little in the host compared to vacuum, with the cluster of nearest neighbours capable of capturing the impact of the environment.

Turning our attention to the S_1 state, we notice that both PBE and B3LYP-in-PBE significantly underestimate $\Delta E_{S_0 \rightarrow S_1}$ for both the vacuum and *p*-terphenyl host structures. The vacuum energy of 1.895 eV for B3LYP-in-PBE is in line with previous *ab initio* calculations of $\Delta E_{S_0 \rightarrow S_1}$ using time-dependent density functional theory (TDDFT) with B3LYP,⁷³ which predicted an S_1 energy of 1.89 eV. While also underestimated relative to experiment, the PBE vacuum result of 1.835 eV is significantly higher than reference *ab initio* studies of pentacene using PBE with TDDFT, which predict values in the range of 1.58–1.64 eV for S_1 ,^{60,64,74} suggesting that the virial exciton method can reduce the error caused by the failure of semi-local TDDFT to capture excitation energies for linear acenes. These results seem to confirm that the combination of EMFT and the virial exciton approach is working as expected in producing vacuum excitation energies, although these do not match experimental absolute energies. Instead, we now consider the solvatochromic shift from placing the pentacene within a *p*-terphenyl host.

For PBE, we observe a redshift of 34 meV and 31 meV in $\Delta E_{S_0 \rightarrow S_1}$ for the cluster and crystal, respectively, relative to the vacuum energy. The former is in fact lower than a previous study using PBE for the same structure,⁶⁰ though this may reflect the tendency of PBE to encourage charge delocalisation as a result of spurious self-interaction. This causes long-range charge-transfer effects to be underestimated in energy by semi-local TDDFT, such that the S_1 state is no longer a pure HOMO→LUMO transition and instead contains

significant contributions from neighbouring molecules.⁶⁰ The fact that the crystal redshift is in fact lower than the cluster result using Δ SCF implies that the PBE excitation is largely unaffected by the presence of the wider crystal.

For B3LYP-in-PBE, the redshifts amount to 42 meV and 139 meV in $\Delta E_{S_0 \rightarrow S_1}$ for the cluster and crystal, respectively. The cluster result is slightly larger than the PBE redshift of 34 meV, in line with past studies using the cluster structure that showed a greater redshift for $\Delta E_{S_0 \rightarrow S_1}$ using OT-LC ω PBE as compared to PBE.⁶⁰

Comparing the cluster and crystal results, we observe a much larger redshift of 139 meV in the crystal structure with B3LYP-in-PBE, which is much closer to the anticipated redshift of ~ 0.2 eV and, significantly, indicates a different behaviour between the PBE and B3LYP functionals for this system. There are two scenarios that may explain the difference between the cluster and crystal redshifts. First, this suggests that the cluster is not sufficient to capture all the effects of the environment on the pentacene excited states, for example, the cluster only includes a ‘cylinder’ of *p*-terphenyls around the pentacene, with the end-to-end molecules along the *c*-axis in Fig. 3 being excluded. Thus a larger structure such as the periodic crystal utilised here is required to describe the full impact of the environment. The fact that PBE does not demonstrate similar behaviour suggests there are long-range effects that are not accounted for by the use of semi-local functionals for the crystal structure. Alternatively, this may indicate a spuriously low singlet-triplet splitting in the crystal structure due to changes in the S_1 state in the host — if the excited state is not strictly composed of the HOMO \rightarrow LUMO transition, then Eq. (19) is not valid and $\Delta E_{T_1 \rightarrow S_1}$ may be misleading. However, it is not clear why this would only apply to the crystal and not the cluster too, and past studies⁶⁰ of the composition of the $S_0 \rightarrow S_1$ transition in explicit host suggest that HOMO \rightarrow LUMO remains dominant, particularly for hybrid functionals, so it appears that the redshift is meaningful. Further studies of the validity of the virial electron method in host structures would help to clarify this issue.

In summary, our B3LYP-in-PBE calculations produce results that are consistent with

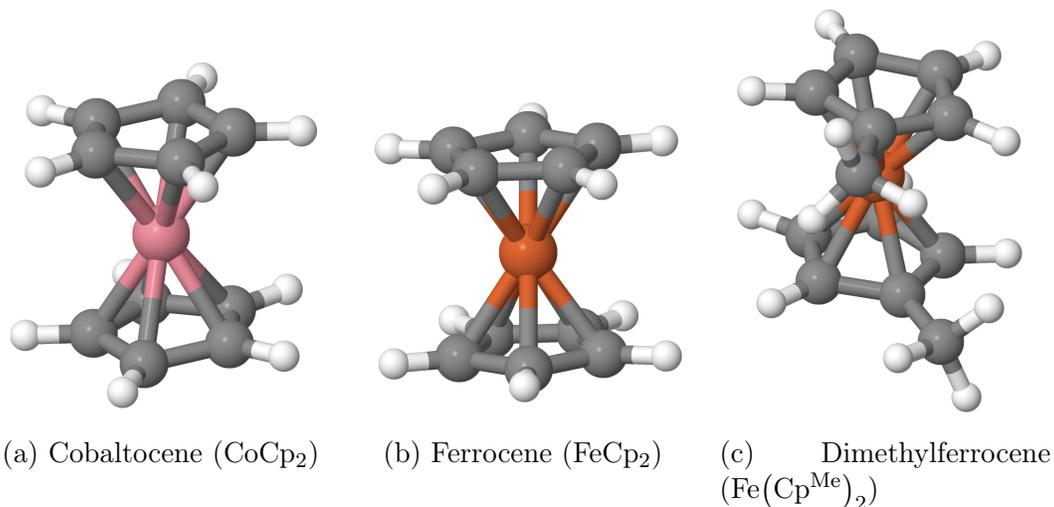


Figure 4: The metallocenes considered in this study. All consist of a transition metal ion, either Co or Fe, sandwiched between two cyclopentadienyl rings, which may have methyl side-groups substituted onto them. Each molecule is labelled by the name it is referred to by in this work. All structures are geometry optimised within ONETEP *in vacuo* at the PBE level. Co, Fe, H, and C atoms are pink, orange, white, and grey respectively. Figures produced using Jmol.⁵⁰

past *ab initio* studies using B3LYP for both the singlet and triplet configurations, and demonstrates divergent behaviour compared to using PBE. Although there are unanswered questions regarding the validity of the redshifts obtained for the explicit host systems, the combination of EMFT and linear-scaling DFT opens new avenues for studying such systems by expanding the range of possible simulations that can be done at the DFT-in-DFT level to several thousands of atoms.

4.3 Metallocenes in carbon nanotubes

The final system we study in this work is that of metallocenes encapsulated in carbon nanotubes. Carbon nanotubes (CNTs) are among the most well-known and studied materials for future technological applications, thanks to their very large tensile strength,^{76,77} high electrical⁷⁸ and thermal conductivity,⁷⁹ and chemical stability. The precise properties of a CNT depend on exactly how it is ‘rolled up’, usually defined by the two integers (n, m) – for example, if $n - m$ is divisible by 3, the CNT will be metallic (although curvature

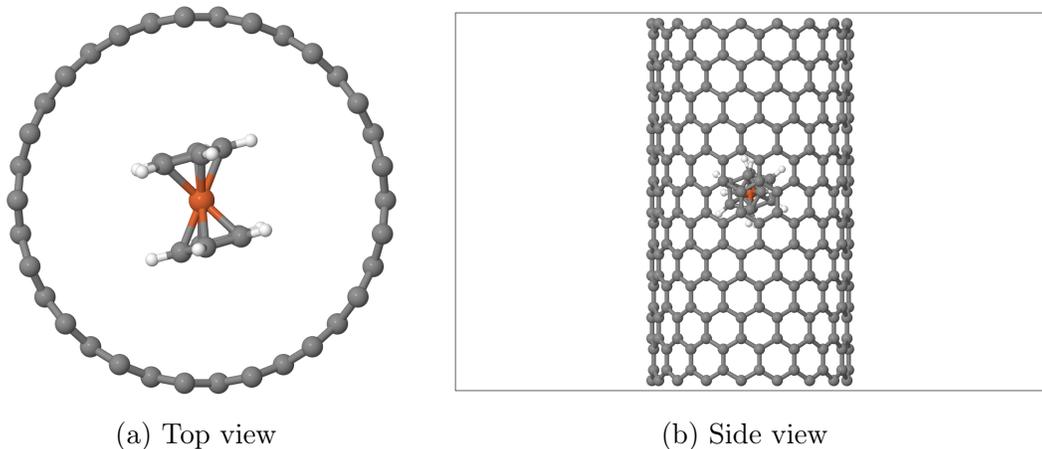


Figure 5: An example geometry for a metallocene encapsulated within a CNT – in this case ferrocene within a semiconducting (17,0) CNT – from two perspectives. (a) shows the structure looking along the axis of the CNT, showing the orientation of the metallocene perpendicular to this axis, following previous work.⁷⁵ (b) shows the structure looking perpendicular to the axis of the CNT, and also shows the periodic box used for the calculations. The box has significant size in the plane perpendicular to the axis of the CNT to prevent the system interacting with its periodic images. Fe, H, and C atoms are orange, white, and grey respectively. Figures produced using Jmol.⁵⁰

effects will open up a gap if n and m are small), but semiconducting otherwise.⁸⁰ Besides the details of their construction, the electronic properties of CNTs can be further tuned by encapsulating guest molecules within them. Recent work has made use of transition metal complexes, including metallocenes, for this.^{81–84} Metallocenes, which consist of a transition metal cation M sandwiched between two cyclopentadienyl anions, have the general formula $M(C_5H_5)_2$, often abbreviated as MCp_2 . Ferrocene ($FeCp_2$) and cobaltocene ($CoCp_2$)⁸⁵ are among the most commonly synthesised metallocenes, and derivative compounds have also been synthesised and studied extensively.⁸⁶

Such metallocene@CNT systems are amenable to investigation using EMFT within ONETEP for several reasons: there is a clear separation between an active region (the guest molecule) and the environment (the CNT); including exact exchange should allow us to describe transition metal complexes more accurately (although this is not guaranteed in all cases^{87,88}), and embedding allows the metallocene to be described in this way whilst still including a large portion of the CNT; and the potential technological importance of these systems means

there is previous work to compare against.⁷⁵

Here, we will focus on the encapsulation of three different complexes – cobaltocene (CoCp_2), ferrocene (FeCp_2), and a derivative of ferrocene here called dimethylferrocene ($\text{Fe}(\text{Cp}^{\text{Me}})_2$). The structure of these compounds is shown in Fig. 4. The CNT considered here is the $(17, 0)$ CNT, which is semiconducting. These components are chosen to allow straightforward comparison to the results of Ref. 75. The quantity of interest is the encapsulation energy of the guest molecule, calculated as $\Delta E_{\text{enc}} = E_{\text{metallocene@CNT}} - E_{\text{metallocene}} - E_{\text{CNT}}$, where care is taken to ensure that the metallocene and the CNT are treated at the same level of theory in the calculation of all terms on the right hand side. If a fixed localised basis set, such as Gaussians, were used to calculate this quantity, a counterpoise correction would typically be required.⁸⁹ In standard ONETEP calculations, this correction is unnecessary, thanks to the optimisation of the NGWF basis set.⁹⁰ However, the NGWFs used for the active region in our EMFT implementation are only optimised at the lower level of theory, meaning that a counterpoise correction is no longer strictly unnecessary. Nevertheless, this change in basis set only has a small effect on our results, as demonstrated above, and therefore we expect any counterpoise correction to also be small.

To create the geometries required for these calculations, firstly the structure of the metallocene molecules was optimised at the PBE level in vacuum. These optimised molecules were then placed in a perpendicular orientation⁷⁵ at the centre of a $(17, 0)$ CNT. No further geometry optimisation was performed; previous PBE calculations on similar systems predict that the metallocenes do not move significantly away from the CNT axis.⁹¹ A test geometry optimisation using PBE on the cobaltocene@CNT system confirmed this. Previous geometry optimisations done using the LDA functional *do* predict that the metallocenes move away from the CNT axis, in better agreement with experiment,⁹¹ but in this work, we choose to adhere to the PBE result to be consistent throughout. Fully periodic boundary conditions were used, but the cell size was sufficient to separate the CNT and its periodic images by at least 40 Å. The only real periodicity, therefore, was that along the CNT itself, with a lattice

Table 2: Error in total energy from using PBE-optimised NGWFs in a B3LYP calculation on several metallocenes, relative to the difference in energy between a full PBE and a full B3LYP calculation.

Molecule	Relative error from NGWFs (%)
CoCp ₂	1.06
FeCp ₂	0.960
Fe(Cp ^{Me}) ₂	0.998

Table 3: Encapsulation energy of several metallocenes within a (17, 0) CNT, calculated at the PBE and at the B3LYP-in-PBE level. In the B3LYP-in-PBE calculations, the metallocene forms the active region treated with B3LYP. The structures of the molecules used are shown in Fig. 4, and an example of an encapsulated configuration is shown in Fig. 5.

Molecule	E_{enc} (eV)	
	PBE	B3LYP-in-PBE
CoCp ₂	0.112	0.654
FeCp ₂	0.051	0.277
Fe(Cp ^{Me}) ₂	0.114	0.756

vector of 25.584 Å. An example of one of the resulting structures is shown in Fig. 5.

In all calculations, a cutoff energy of 1200 eV was used, with all NGWF radii were set to 5.29 Å (10 bohr). PBE and B3LYP are used for the low and high levels of theory respectively. The overall spin of the system is fixed at zero throughout, and the pseudopotentials used were the appropriate on-the-fly norm-conserving pseudopotentials generated by the CASTEP code.⁹² The block orthogonalisation procedure outlined in Section 3 was used for all EMFT calculations. All calculations were done using NGWFs optimised at the PBE level, with the exception of some calculations of the total energy of individual metallocene molecules, which were done fully at the B3LYP level for the purposes of comparison.

Table 2 shows, for each of the metallocenes in vacuo, the difference in energy (calculated at the B3LYP level) arising from using PBE-optimised NGWFs rather than B3LYP-optimised NGWFs, expressed as a percentage of the difference in energy between the full PBE and full B3LYP calculations. As also seen above in Fig. 2, the error arising from using NGWFs optimised at the low level of theory is relatively small, around 1%, again justifying the approximation made in Section 3.

Table 3 shows the value of E_{enc} for each metallocene encapsulated in a (17,0) CNT, calculated at both the PBE and B3LYP-in-PBE level. It can be seen that cobaltocene binds significantly more strongly to the CNT than ferrocene. This is consistent with similar previous work on (16,0) CNTs,⁹¹ and is as expected due to cobaltocene’s much higher tendency to transfer charge to the CNT.⁷⁵ Dimethylferrocene is also seen to bind more strongly than bare ferrocene, as the methyl groups encourage charge transfer to the CNT.⁷⁵ Our results here are not directly comparable to these previous results, as the metallocene is positioned on-axis (as discussed above). However, from a qualitative standpoint, the PBE results agree well with the general trend outlined above, as do the B3LYP-in-PBE results, which shows our implementation of EMFT is successfully describing the system. It should be noted that there are no direct experimental measurements of the encapsulation energies in metallocene@CNT systems that we can compare against, only semi-local DFT calculations, but these results demonstrate the power of our implementation. We are able to describe a large host-guest system fully from first principles using a high level of theory, at a significantly reduced computational cost – such a calculation would not have previously been possible.

5 Conclusions

In conclusion, we have presented the first implementation of embedded mean-field theory, a quantum embedding scheme already successfully applied to molecular systems, into a linear scaling DFT code, ONETEP. This allows the possibility of DFT-in-DFT embedding calculations on systems significantly larger than was previously possible, by making use of the linear scaling behaviour of ONETEP. By doing this, the behaviour of the region of interest can be investigated at a high level of theory, whilst still including the effect of a large environment, at a fraction of the cost of treating the whole system at the high level of theory. Our implementation has been applied to a range of different systems, and was found to perform very well compared to previous first principles calculations and experiment.

In particular, the power of this new work is demonstrated explicitly by the application of EMFT to a pentacene-doped *p*-terphenyl system containing nearly 3000 atoms. In this calculation, the active region was treated with a hybrid functional, and the environment with a semi-local functional, obtaining accurate results at a significantly lower computational cost than an equivalent full hybrid functional calculation. This work therefore opens the door for applying such embedding calculations to even larger systems, enabling the study of a variety of problems of interest in condensed matter physics and chemistry, including solvatochromism, doping of molecular crystals, and defects in semiconductors.

Supporting Information Available

The supplementary information provided alongside this work includes details of the pseudopotentials used in this work, and an auxiliary set of results obtained using a ‘mixed’ set of NGWFs – that is, where some NGWFs are optimised at the lower level of theory and some at the higher level.

Acknowledgement

We acknowledge support from the UK Collaborative Computational Project for the Study of the Electronic Structure of Condensed Matter (CCP9) and the Engineering and Physical Sciences Research Council in the form of Software Infrastructure Grant EP/P02209X/1. We also acknowledge the support of the Thomas Young Centre for Theory and Simulation of Materials, through grant TYC-101. R.J.C. acknowledges the studentship support of the EPSRC Centre for Doctoral Training Centre in Theory and Simulation of Materials (No. EP/L015579/1) and a President’s PhD Scholarship from Imperial College London.

References

- (1) Eisbein, E.; Joswig, J.-O. et al. Proton Conduction in a MIL-53(Al) Metal-Organic Framework: Confinement versus Host/Guest Interaction. J. Phys. Chem. C **2014**, 118, 13035.
- (2) Hirao, H.; Ng, W. K. H. et al. Multiscale Model for a Metal-Organic Framework: High-Spin Rebound Mechanism in the Reaction of the Oxoiron(IV) Species of Fe-MOF-74. ACS Catal. **2015**, 5, 3287.
- (3) Witman, M.; Ling, S. et al. Rational Design of a Low-Cost, High-Performance Metal-Organic Framework for Hydrogen Storage and Carbon Capture. J. Phys. Chem. C **2017**, 121, 1171.
- (4) Altun, A.; Yokoyama, S. et al. Spectral Tuning in Visual Pigments: An ONIOM(QM:MM) Study on Bovine Rhodopsin and its Mutants. J. Phys. Chem. B **2008**, 112, 6814.
- (5) Isborn, C. M.; Götz, A. W. et al. Electronic Absorption Spectra from MM and *ab initio* QM/MM Molecular Dynamics: Environmental Effects on the Absorption Spectrum of Photoactive Yellow Protein. J. Chem. Theory Comput. **2012**, 8, 5092.
- (6) Zuehlsdorff, T. J.; Haynes, P. D. et al. Solvent Effects on Electronic Excitations of an Organic Chromophore. J. Chem. Theory Comput. **2016**, 12, 1853.
- (7) Huber, L.; Grabowski, B. et al. A QM/MM approach for low-symmetry defects in metals. Comput. Mater. Sci. **2016**, 118, 259.
- (8) Chen, H.; Ortner, C. QM/MM Methods for Crystalline Defects. Part 1: Locality of the Tight Binding Model. Multiscale Model. Simul. **2016**, 14, 232.
- (9) Chen, H.; Ortner, C. QM/MM Methods for Crystalline Defects. Part 2: Consistent Energy and Force-Mixing. Multiscale Model. Simul. **2017**, 15, 184.

- (10) Mulholland, A. J. Chemical accuracy in QM/MM calculations on enzyme-catalysed reactions. Chem. Cent. J. **2007**, 1, 19.
- (11) Cole, D. J.; Hine, N. D. M. Applications of large-scale density functional theory in biology. J. Phys.: Condens. Matter **2016**, 28, 393001.
- (12) Kulik, H. J. Large-scale QM/MM free energy simulations of enzyme catalysis reveal the influence of charge transfer. Phys. Chem. Chem. Phys. **2018**, 20, 20650.
- (13) Warshel, A.; Levitt, M. Theoretical studies of enzymic reactions: Dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme. J. Mol. Biol. **1976**, 103, 227.
- (14) Svensson, M.; Humbel, S. et al. ONIOM: A Multilayered Integrated MO + MM Method for Geometry Optimizations and Single Point Energy Predictions. A Test for Diels-Alder Reactions and Pt(P(t-Bu)₃)₂ + H₂ Oxidative Addition. J. Phys. Chem. **1996**, 100, 19357.
- (15) Friesner, R. A.; Guallar, V. Ab initio quantum chemical and mixed quantum mechanics/molecular mechanics (QM/MM) methods for studying enzymatic catalysis. Annu. Rev. Phys. Chem. **2004**, 56, 389.
- (16) Lin, H.; Truhlar, D. G. QM/MM: what have we learned, where are we, and where do we go from here? Theor. Chem. Acc. **2006**, 117, 185.
- (17) Senn, H. M.; Thiel, W. QM/MM studies of enzymes. Curr. Opin. Chem. Biol. **2007**, 11, 182.
- (18) van der Kamp, M. W.; Mulholland, A. J. Combined Quantum Mechanics/Molecular Mechanics (QM/MM) Methods in Computational Enzymology. Biochemistry **2013**, 52, 2708.

- (19) Chung, L. W.; Sameera, W. M. C. et al. The ONIOM Method and Its Applications. Chem. Rev. **2015**, 115, 5678.
- (20) Fornace, M. E.; Lee, J. et al. Embedded Mean-Field Theory. J. Chem. Theory Comput. **2015**, 11, 568.
- (21) Sun, Q.; Chan, G. K.-L. Quantum Embedding Theories. Acc. Chem. Res. **2016**, 49, 2705.
- (22) Jones, R. O. Density functional theory: Its origins, rise to prominence, and future. Rev. Mod. Phys. **2015**, 87, 897.
- (23) Becke, A. D. A new mixing of Hartree-Fock and local density-functional theories. J. Chem. Phys. **1993**, 98, 1372.
- (24) Mardirossian, N.; Head-Gordon, M. Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. Mol. Phys. **2017**, 115, 2315.
- (25) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. Phys. Rev. **1965**, 140, A1133.
- (26) Perdew, J. P.; Zunger, A. Self-interaction correction to density-functional approximations for many-electron systems. Phys. Rev. B **1981**, 23, 5048.
- (27) Perdew, J. P. Accurate Density Functional for the Energy: Real-Space Cutoff of the Gradient Expansion for the Exchange Hole. Phys. Rev. Lett. **1985**, 55, 1665.
- (28) Perdew, J. P.; Burke, K. et al. Generalized Gradient Approximation Made Simple. Phys. Rev. Lett. **1996**, 77, 3865.
- (29) Čížek, J. On the Correlation Problem in Atomic and Molecular Systems. Calculation of Wavefunction Components in Ursell-Type Expansion Using Quantum-Field Theoretical Methods. J. Chem. Phys. **1966**, 45, 4256.

- (30) Bartlett, R. J.; Musiał, M. Coupled-cluster theory in quantum chemistry. Rev. Mod. Phys. **2007**, 79, 291.
- (31) Cortona, P. Self-consistently determined properties of solids without band-structure calculations. Phys. Rev. B **1991**, 44, 8454.
- (32) Wesolowski, T. A.; Warshel, A. Frozen density functional approach for ab initio calculations of solvated molecules. J. Phys. Chem. **1993**, 97, 8050.
- (33) Kaminski, J. W.; Gusarov, S. et al. Modeling Solvatochromic Shifts Using the Orbital-Free Embedding Potential at Statistically Mechanically Averaged Solvent Density. J. Phys. Chem. A **2010**, 114, 6082.
- (34) Staroverov, V. N.; Scuseria, G. E. et al. Optimized effective potentials yielding Hartree-Fock energies and densities. J. Chem. Phys. **2006**, 124, 141103.
- (35) Roncero, O.; de Lara-Castells, M. P. et al. An inversion technique for the calculation of embedding potentials. J. Chem. Phys. **2008**, 129, 184104.
- (36) Goodpaster, J. D.; Ananth, N. et al. Exact nonadditive kinetic potentials for embedded density functional theory. J. Chem. Phys. **2010**, 133, 084103.
- (37) Culpitt, T.; Brorsen, K. R. et al. Communication: Density functional theory embedding with the orthogonality constrained basis set expansion procedure. J. Chem. Phys. **2017**, 146, 211101.
- (38) Huzinaga, S.; Cantu, A. A. Theory of Separability of Many-Electron Systems. J. Chem. Phys. **1971**, 55, 5543.
- (39) Manby, F. R.; Stella, M. et al. A Simple, Exact Density-Functional-Theory Embedding Scheme. J. Chem. Theory Comput. **2012**, 8, 2564.
- (40) Hégyely, B.; Nagy, P. R. et al. Exact density functional and wave function embedding schemes based on orbital localization. J. Chem. Phys. **2016**, 145, 064107.

- (41) Miyamoto, K.; Miller, T. F. et al. Fock-Matrix Corrections in Density Functional Theory and Use in Embedded Mean-Field Theory. J. Chem. Theory Comput. **2016**, 12, 5811.
- (42) Ding, F.; Manby, F. R. et al. Embedded Mean-Field Theory with Block-Orthogonalized Partitioning. J. Chem. Theory Comput. **2017**, 13, 1605–1615.
- (43) Ding, F.; Tsuchiya, T. et al. Linear-Response Time-Dependent Embedded Mean-Field Theory. J. Chem. Theory Comput. **2017**, 13, 4216.
- (44) Koh, K. J.; Nguyen-Beck, T. S. et al. Accelerating Realtime TDDFT with Block-Orthogonalized Manby-Miller Embedding Theory. J. Chem. Theory Comput. **2017**, 13, 4173.
- (45) Jiang, H.; Kammler, M. et al. Imaging covalent bond formation by H atom scattering from graphene. Science **2019**, 364, 379.
- (46) Skylaris, C.-K.; Haynes, P. D. et al. Introducing ONETEP: Linear-scaling density functional simulations on parallel computers. J. Chem. Phys. **2005**, 122, 084119.
- (47) Dziedzic, J.; Hill, Q. et al. Linear-scaling calculation of Hartree-Fock exchange energy with non-orthogonal generalised Wannier functions. J. Chem. Phys. **2013**, 139, 214103.
- (48) Skylaris, C.-K.; Haynes, P. D. et al. Implementation of linear-scaling plane wave density functional theory on parallel computers. Phys. Status Solidi B **2006**, 243, 973.
- (49) Haynes, P. D.; Skylaris, C.-K. et al. ONETEP: linear-scaling density-functional theory with local orbitals and plane waves. Phys. Status Solidi B **2006**, 243, 2489.
- (50) Hanson, R. M. Jmol – a paradigm shift in crystallographic visualization. J. Appl. Cryst. **2010**, 43, 1250.
- (51) Yoo, S.; Domercq, B. et al. Efficient thin-film organic solar cells based on pentacene/C₆₀ heterojunctions. Appl. Phys. Lett. **2004**, 85, 5427.

- (52) Koch, N. Organic Electronic Devices and Their Functional Interfaces. ChemPhysChem **2007**, 8, 1438.
- (53) Oxborrow, M.; Breeze, J. D. et al. Room-temperature solid-state maser. Nature **2012**, 488, 353.
- (54) Becke, A. D. Density-functional thermochemistry. III. The role of exact exchange. J. Chem. Phys. **1993**, 98, 5648.
- (55) Gordon, J. P.; Zeiger, H. J. et al. The Maser – New Type of Microwave Amplifier, Frequency Standard, and Spectrometer. Phys. Rev. **1955**, 99, 1264.
- (56) Schawlow, A. L.; Townes, C. H. Infrared and Optical Masers. Phys. Rev. **1958**, 112, 1940.
- (57) Kleppner, D.; Goldenberg, H. M. et al. Properties of the Hydrogen Maser. Appl. Opt. **1962**, 1, 55.
- (58) Konoplev, I. V.; McGrane, P. et al. Experimental Study of Coaxial Free-Electron Maser Based on Two-Dimensional Distributed Feedback. Phys. Rev. Lett. **2006**, 96, 035002.
- (59) Siegman, A. E. Microwave solid-state masers; McGraw-Hill Electrical and Electronic Engineering series; McGraw-Hill, 1964.
- (60) Charlton, R. J.; Fogarty, R. M. et al. Implicit and explicit host effects on excitons in pentacene derivatives. J. Chem. Phys. **2018**, 148, 104108.
- (61) Köhler, J.; Brouwer, A. C. J. et al. On the intersystem crossing of pentacene in p-terphenyl. Chem. Phys. Lett. **1996**, 250, 137.
- (62) Patterson, F. G.; Lee, H. W. H. et al. Intersystem crossing from singlet states of molecular dimers and monomers in mixed molecular crystals: picosecond stimulated photon echo experiments. Chem. Phys. **1984**, 84, 51.

- (63) Zimmerman, P. M.; Zhang, Z. et al. Singlet fission in pentacene through multi-exciton quantum states. Nat. Chem. **2010**, 2, 648.
- (64) Bogatko, S.; Haynes, P. D. et al. Molecular Design of a Room-Temperature Maser. J. Phys. Chem. C **2016**, 120, 8251.
- (65) Zuehlsdorff, T. J.; Haynes, P. D. et al. Predicting solvatochromic shifts and colours of a solvated organic dye: The example of nile red. J. Chem. Phys. **2017**, 146, 124504.
- (66) Becke, A. D. Singlet-triplet splittings from the virial theorem and single-particle excitation energies. J. Chem. Phys. **2018**, 148, 044112.
- (67) Feng, X.; Becke, A. D. et al. Communication: Becke’s virial exciton model gives accurate charge-transfer excitation energies. J. Chem. Phys. **2018**, 149.
- (68) Rietveld, H. M.; Maslen, E. N. et al. An X-ray and neutron diffraction refinement of the structure of p-terphenyl. Acta Crystallogr. B **1970**, 26, 693.
- (69) Groom, C. R.; Bruno, I. J. et al. The Cambridge Structural Database. Acta. Cryst. B. **2016**, 72, 171.
- (70) Burgos, J.; Pope, M. et al. Heterofission in pentacene-doped tetracene single crystals. Phys. Status Solidi B **1977**, 83, 249.
- (71) Heinecke, E.; Hartmann, D. et al. Laser spectroscopy of free pentacene molecules (I): The rotational structure of the vibrationless S₁←S₀ transition. J. Chem. Phys. **1998**, 109, 906.
- (72) Autschbach, J.; Srebro, M. Delocalization Error and “Functional Tuning” in Kohn-Sham Calculations of Molecular Properties. Acc. Chem. Res. **2014**, 47, 2592.
- (73) Grimme, S.; Parac, M. Substantial errors from time-dependent density functional theory for the calculation of excited states of large pi systems. ChemPhysChem **2003**, 4, 292.

- (74) Kadantsev, E. S.; Stott, M. J. et al. Electronic structure and excitations in oligoacenes from ab initio calculations. Journal of Chemical Physics **2006**, 124.
- (75) McSweeney, R. L.; Chamberlain, T. W. et al. Direct Measurement of Electron Transfer in Nanoscale Host-Guest Systems: Metallocenes in Carbon Nanotubes. Chem: Eur. J. **2016**, 22, 13540.
- (76) Peng, B.; Locascio, M. et al. Measurements of near-ultimate strength for multi-walled carbon nanotubes and irradiation-induced crosslinking improvements. Nat. Nanotechnol. **2008**, 3, 626.
- (77) Janas, D.; Vilatela, A. C. et al. Performance of carbon nanotube wires in extreme conditions. Carbon **2013**, 62, 438.
- (78) Kang, S. J.; Kocabas, C. et al. High-performance electronics using dense, perfectly aligned arrays of single-walled carbon nanotubes. Nat. Nanotechnol. **2007**, 2, 230.
- (79) Koziol, K. K.; Janas, D. et al. Thermal properties of continuously spun carbon nanotube fibres. Physica E **2017**, 88, 104.
- (80) Nanot, S.; Thompson, N. A. et al. Single-Walled Carbon Nanotubes. Springer Handbook of Nanomaterials. 2013.
- (81) Lebedeva, M. A.; Chamberlain, T. W. et al. New Pathway for Heterogenization of Molecular Catalysts by Non-covalent Interactions with Carbon Nanoreactors. Chem. Mater. **2014**, 26, 6461.
- (82) Pan, X.; Bao, X. Reactions over catalysts confined in carbon nanotubes. Chem. Commun. **2008**, 6271.
- (83) Cheng, H.; Qiu, H. et al. Investigation of the electrochemical behavior of dopamine at electrodes modified with ferrocene-filled double-walled carbon nanotubes. Electrochim. Acta **2012**, 63, 83.

- (84) del Carmen Giménez-López, M.; Moro, F. et al. Encapsulation of single-molecule magnets in carbon nanotubes. Nat. Commun. **2011**, 2, 407.
- (85) Birmingham, J. M.; Fischer, A. K. et al. The reduction of bis-cyclopentadienyl compounds. Naturwissenschaften **1955**, 42, 96.
- (86) Chirik, P. J. Group 4 Transition Metal Sandwich Complexes: Still Fresh after Almost 60 Years. Organometallics **2010**, 29, 1500.
- (87) Cramer, C. J.; Truhlar, D. G. Density functional theory for transition metals and transition metal chemistry. Phys. Chem. Chem. Phys. **2009**, 11, 10757.
- (88) Harvey, J. N. On the accuracy of density functional theory in transition metal chemistry. Annu. Rep. Prog. Chem., Sect. C: Phys. Chem. **2006**, 102, 203.
- (89) Boys, S. F.; Bernardi, F. The calculation of small molecular interactions by the differences of separate total energies. Some procedures with reduced errors. Mol. Phys. **1970**, 19, 553.
- (90) Haynes, P. D.; Skylaris, C.-K. et al. Elimination of basis set superposition error in linear-scaling density-functional calculations with local orbitals optimised in situ. Chem. Phys. Lett. **2006**, 422, 345–349.
- (91) Cao, F.; Ren, W. et al. The structural, energetic and electronic properties of doped carbon nanotubes by encapsulation of MCp₂ (M=Fe, Co, Ni): A theoretical investigation. Chem. Phys. Lett. **2011**, 512, 81.
- (92) Clark, S. J.; Segall, M. D. et al. First principles methods using CASTEP. Z. Kristallogr. **2005**, 220, 567.

Graphical TOC Entry

