5 DMRG as Impurity Solver

Ulrich Schollwöck
Department of Physics, University of Munich
Theresienstrasse 37, 80333 München

Contents

E. Pavarini and E. Koch (eds.)
Correlated Materials: Methods and Applications
Modeling and Simulation Vol. 16, Forschungszentrum Jülich, 2026
ISBN 978-3-95806-909-1 http://www.cond-mat.de/events/correl26

1 Tensor networks and embedding techniques

In this lecture, I will bring together two widely used approaches to the quantum many-body problem which take quite unrelated perspectives. After a long prehistory, tensor networks firmly entered the realm of strongly correlated systems as described by the Heisenberg and Hubbard model Hamiltonians in the form of White's density-matrix renormalization-group (DMRG) in 1992 [1]. This method for determining ground states and low-lying excited states was quickly recognized to be, ignoring some subtle details, nothing but an efficient way of variationally determining the lowest-lying eigenstates of a Hamiltonian within a very special class of ansatz states, so-called matrix-product states (MPS). The Hamiltonian is treated exactly, the state space is strongly restricted. Powerful generalizations and improvements have been brought forward over the last 30 years [2, 3].

At the same time, earlier insights into simplifications of the quantum many-body problem in large dimensions (more precisely, in systems with large coordination numbers) [4] led Georges and Kotliar to the formulation of dynamical mean-field theory (DMFT) as an embedding approach to the many-body problem [5–9]. Unlike tensor network methods, it approximates the interaction part of the Hamiltonian by treating a manageable subsystem exactly while modeling the rest of the system by an effective non-interacting "bath" and computes physical observables without explicitly calculating the many-body wave function. DMFT crucially relies on a self-consistency condition for the determination of the bath. Its earliest predecessor is the well-known Weiss mean-field theory of ferromagnetism, where self-consistency is achieved at the level of spin magnetization: The interaction of one particular spin to the others is modeled by a coupling to an effective magnetic field constant in space and time, and the self-consistency condition is that the magnetization of that spin in the effective field must be the same as the magnetization leading to that effective field; after all, the chosen spin was not special. In DMFT for electronic problems, one electronic orbital (site) is picked out and couples through electronic hopping to other orbitals (sites). The "bath" is now time-dependent (dynamical) as opposed to the effective field in Weiss theory, and it turns out that for many problems this is the decisive step forward. Just like the Weiss field, which implies that we consider the other spins to be non-interacting, the bath consists of non-interacting electrons, making the problem more treatable. The self-consistency condition of DMFT, to be detailed below, is that the local self-energy of the chosen orbital is the same as the local self-energy of the lattice; again, the chosen orbital was not special. As in the case of DMRG, original DMFT has been the first step towards highly useful generalizations of the idea of dynamical embedding, and in both cases these generalizations are the subject of intense ongoing research [7–9]. What is particularly interesting about DMFT is that the approach also works for real systems, with the band structure of the material under study calculated by density functional theory (DFT). The DFT band structure has issues with strong correlations due to local interactions, but this part is dealt with by DMFT. The DFT+DMFT method, as we may call it, has emerged as a very powerful method to study real strongly correlated materials beyond simple Hubbard and Heisenberg Hamiltonians [9].

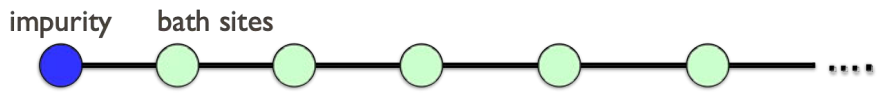

Fig. 1: *Impurity and bath in the NRG framework: The interacting impurity (blue) sits at the end of a one-dimensional chain. All other sites correspond to orbitals with non-interacting electrons, which have an on-site energy and hop to the nearest-neighbor; they form the bath.*

These seemingly unrelated approaches to the quantum many-body problem come together because DMFT (and generalizations) requires for the self-consistency condition the calculation of the local Green function of the chosen interacting orbital, embedded in the non-interacting bath. While this is a much simpler problem than the full many-body problem, this calculation still represents the core numerical challenge and limitation of DMFT, where actual DMFT calculations spend most of their time. A substantial fraction of methodological progress has been related to improving schemes for calculating that Green function. As the setup resembles that of an interacting impurity in a solid with bands of (non-interacting) electrons, well-known from the famous Kondo problem, one speaks colloquially of the "impurity problem" that has to be solved within DMFT and of the "impurity Green function"; the methods are called "impurity solvers". We will use DMRG and more generally tensor network techniques, of which it is one, as impurity solver, hence the title of the lecture.

Most numerical techniques for solving quantum many-body problems have been tried and are used regularly for the impurity problem. Arguably the most versatile technique at the moment is continuous-time quantum Monte Carlo, which comes in two main incarnations, treating kinetic terms as a correction to the local interaction, or the local interaction as a correction to kinetic terms (CT-HYB and CT-INT) [10, 11]. We will speak of QMC-DMFT. Other methods are ED-DMFT and NRG-DMFT. In the former [12], the bath is represented by a short chain of non-interacting electronic sites, with dynamical quantities expressed through Matsubara (imaginary) frequencies (for details, see below); calculations are based on exact diagonalization (ED). The use of DMRG replacing ED is suggested by the fact that DMRG is able to treat larger systems than ED. Alternatively, one can express bath properties by using real frequencies and use the framework of the numerical renormalization group (NRG), which was specifically developed for the Kondo problem, to solve the impurity problem [13]. The use of DMRG replacing NRG is suggested by the fact that NRG leads to an effectively one-dimensional problem (Figure 1), where DMRG excels. The computational scaling of DMRG is also more advantageous for more complicated (e.g. multi-orbital) problems.

The idea of using tensor network (TN) algorithms, in particular DMRG, as an impurity solver, is therefore unsurprisingly not recent, but goes back to the early 2000s (see e.g. [14] for an early example; the next bigger steps came around 10 years later, see [15–18]). The two main approaches to TN-DMFT (real-time DMFT [15–17] and imaginary-time TN-DMFT [18]) are direct adaptations of NRG-DMFT and ED-DMFT. Despite this conceptual simplicity, TN-DMFT has faced numerous challenges and it is only recently that it has emerged not just as a serious competitor to other approaches, but even having a certain lead, in particular when used in cooperation with other DMFT impurity solvers.

I will set out by discussing some of the essentials of DMRG, MPS and tensor networks, to set the nomenclature and also in order to develop an understanding where the methodological bottlenecks lie. I will do the same for DMFT, but in a more compressed form as the progress I want to describe is essentially on the tensor network side; in fact, in our calculations we use the highly developed and refined TRIQS library [19, 20] more or less as a black box for DMFT which in a self-consistency loop both generates the precise impurity problem to be solved and then processes the solution we provide. There is no change to DMFT involved except for the impurity problem module.

We will then discuss crucial improvements to earlier tensor network setups for impurity problems and more generally the calculation of spectral (Green) functions. They are at the heart of this lecture, and in fact topics of ongoing research.

Next, I will set these improvements to work, using the example of LiV_2O_4 [21]. The low-temperature properties of this material have puzzled the community for almost three decades, and numerous explanations have been put forward. All were hampered by the fact that these properties have no high-temperature counterpart and only emerge after a low-temperature crossover where several very small temperature scales are involved. In the past, DMFT (or rather, the impurity solvers) were not able to access the crossover regime and lower temperatures. This is now possible using TN solvers; at the higher temperatures also accessible to QMC solvers, agreement is excellent. Most importantly, we now have an understanding of the low-temperature behavior of LiV_2O_4, and a methodology which should be able to solve similar problems quite easily.

I conclude by ongoing research which has not yet fully been integrated into TN-DMFT, but will be so very soon: can we make calculations of spectral functions and Green functions using real frequencies numerically less limited in reachable times, hence frequency resolution, by a smart use of time-evolution contours in the complex plane? Current results seem to be very promising [22].

2 Tensor networks in a nutshell

Before we dive into details, the key messages which will be required for the rest of the lecture, are: tensor networks yield low-entanglement approximations of quantum states; ground states and mixed thermal states are mostly weakly entangled, hence the power of tensor networks; real-time evolutions take us quickly out of this cosy corner, making long time evolutions, hence fine spectral resolution (which we require), a key challenge in the field.

Let us illustrate tensor networks by starting with the simplest such network, a matrix product state (MPS): let us consider the state space of a (one-dimensional) lattice with L sites, where the local Hilbert space of site i is spanned by site-local states $\{|\sigma_i\rangle\}$. The simplest local Hilbert spaces are the two degrees of a spin-1/2, $\{|\uparrow\rangle, |\downarrow\rangle\}$, or the four degrees of freedom of a local electronic orbital, $\{|\emptyset\rangle, |\uparrow\rangle, |\downarrow\rangle, |\uparrow\downarrow\rangle\}$. Let us call the number of local degrees of freedom d, here 2 or 4. The Hilbert space of the entire system, given by the tensor product of L local

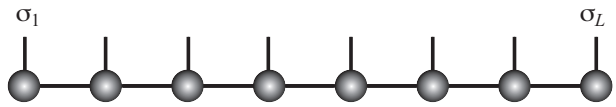

Fig. 2: *A matrix product state (MPS), the simplest tensor network state (TNS): each site (orbital) is represented by a tensor (spherical blob), here with three indices, $M^{\sigma}_{\alpha\beta}$. The physical leg, indicated by a vertical line, represents the local state $|\sigma\rangle$; the two auxiliary legs, indicated by horizontal lines, are contracted over when linked between tensors. The remaining scalar is the expansion coefficient $c^{\sigma_1\sigma_2\cdots\sigma_L}$ in Eq. 1.*

spaces, then has dimension d^L, and the most generic pure quantum state reads

$$|\Psi\rangle = \sum_{\sigma_1,\ldots,\sigma_L} c^{\sigma_1\sigma_2\ldots\sigma_L}|\sigma_1\sigma_2\ldots\sigma_L\rangle \tag{1}$$

with complex coefficients $c^{\sigma_1\sigma_2\ldots\sigma_L}$. In bulk systems we are interested in the thermodynamic limit $L\to\infty$, making an exact treatment unfeasible in that limit, as d^L diverges exponentially. Matrix product states (MPS) form a subset of the full Hilbert space, where the expansion coefficients are given by the product of matrices M^{σ_i}: at each site i, there is one matrix for each local state $|\sigma_i\rangle$. The most general MPS then reads

$$|\Psi\rangle = \sum_{\sigma_1,\ldots,\sigma_L} M^{\sigma_1}M^{\sigma_2}\ldots M^{\sigma_L}|\sigma_1\sigma_2\ldots\sigma_L\rangle. \tag{2}$$

As the matrices are multiplied, their dimensions must match between neighboring matrices; the first and last matrix must be row and column vectors, such that the total product is a scalar $c^{\sigma_1\sigma_2\ldots\sigma_L}$. We can immediately think about generalizations: we read the local objects as (three-legged) tensors $M^{\sigma}_{\alpha\beta}$ where the "auxiliary indices" α and β are contracted with neighboring tensors. We can have, e.g., five-legged tensors $M^{\sigma}_{\alpha\beta\gamma\delta}$, with four auxiliary indices, which link to four other tensors, sitting on the nearest-neighbor sites of a square lattice; ultimately the state coefficients result from contracting over all of the auxiliary indices. An MPS is shown in Figure 2.

It can be shown easily [3] by applying singular value decompositions (SVDs) to coefficients $c^{\sigma_1\sigma_2\ldots\sigma_L}$ that each state can be expressed exactly as an MPS. The price to pay is that towards the center, the matrices become exponentially large, with the largest dimension being $d^{L/2}$, so nothing seems to be gained. The interesting observation is that many quantum states can be expressed almost exactly by much smaller matrices that are limited to some dimension of at most $D\times D$, where D is somewhere in $O(100)$ to $O(1000)$, for which matrix multiplications are fast and efficient, as they scale as D^3. The effective number of coefficients is then $O(dD^2L)$, no longer exponential in system size.

Why is this? It turns out (we only need the result) that if the matrices are chosen to obey certain normalization conditions (always possible), each chain can be partitioned between site ℓ and $\ell{+}1$ into parts A and B and the MPS be written as

$$|\Psi\rangle = \sum_{\alpha} s_\alpha|\alpha\rangle_A|\alpha\rangle_B \tag{3}$$

where $s_\alpha \geq 0$, $\sum_\alpha s_\alpha^2 = 1$ and the states

$$|\alpha\rangle_A = \sum_{\sigma_1 \ldots \sigma_\ell} (M^{\sigma_1} M^{\sigma_2} \ldots M^{\sigma_\ell})_{1,\alpha} |\sigma_1 \ldots \sigma_\ell\rangle \tag{4}$$

$$|\alpha\rangle_B = \sum_{\sigma_{\ell+1} \ldots \sigma_L} (M^{\sigma_{\ell+1}} M^{\sigma_{\ell+2}} \ldots M^{\sigma_L})_{\alpha,1} |\sigma_{\ell+1} \ldots \sigma_L\rangle \tag{5}$$

form orthonormal sets within A and B. This is a so-called Schmidt decomposition. Imagine you have an exact MPS, then the sum will run over an exponential number of values α. It turns out, however, that for many states of interest (ground states and low excitations of Hamiltonians with short-ranged interactions) the values of the (ordered) s_α decay very quickly to exponentially small values, and only a minuscule amount actually contributes to the state. In fact, the reduced density matrix for A reads

$$\hat{\rho}_A = \sum_\alpha s_\alpha^2 |\alpha\rangle_{AA} \langle\alpha|, \tag{6}$$

such that errors in observables become negligible if we restrict ourselves to the D largest values s_α, which imposes a compression condition on the matrices M^σ with very little actual loss. As the von Neumann entropy of entanglement of A and B is given by

$$S_{\text{vN}} = -\text{tr}_A \, \hat{\rho}_A \ln \hat{\rho}_A = -\sum_\alpha s_\alpha^2 \ln s_\alpha^2, \tag{7}$$

a quickly decaying spectrum of the s_α corresponds to weak entanglement between A and B. This link is not mathematically rigorous (it is easy to construct counterexamples [23]), but effectively holds for all states we are encountering in actual systems. The quickly decaying spectrum then reflects the so-called area law: while thermal entropy scales as the volume of a system, the von Neumann entropy of low-lying states (of gapful Hamiltonians) scales only as the surface area between A and B [24]. In one dimension, the surface is a point, and entanglement and the computational complexity of the state (which is roughly exponential in entanglement) do not increase, hence the extreme power of MPS for such systems. In two dimensions, entanglement scales as L for a system of size $L \times L$; the required matrix dimensions then grow exponentially in L, strongly limiting performance of MPS in two dimensions (other TNs do not have this issue, but are otherwise computationally much more demanding).

In this lecture, we are concerned with time-dependent Green functions, which require a time-evolution of a quantum state. It is also known that the states obeying the area law form a (non-empty!) set of measure zero in Hilbert space in the thermodynamic limit, which happens to contain the most interesting states for us. Time evolution will take us inevitably out of this comfortable, but minuscule corner of Hilbert space. In the worst case, entanglement will grow linearly in time, hence required matrix dimensions D grow exponentially in time. This very strongly limits the reachable times. In the case of spectral functions, the entanglement growth is more benign (logarithmic), but still dimensions D will grow with some power law in time, imposing a strong restriction on reachable times. This limitation, of course, limits the resolution we may achieve in frequency space and is therefore the central bottleneck of the following discussion.

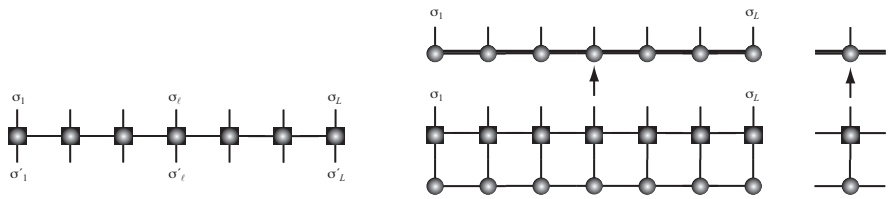

Fig. 3: *A matrix product operator (MPO): on each site (orbital), we have a tensor (square blob), now with four indices, $W^{\sigma\sigma'}_{\alpha\beta}$. The physical legs, indicated by vertical lines, represents the ingoing and outgoing local states $|\sigma^{(\prime)}\rangle$; the two auxiliary legs, indicated by horizontal lines, are contracted over when linked between tensors. When applying an MPO (auxiliary leg dimension D_O) to an MPS, a new MPS results, but with auxiliary leg dimension $D_O D$.*

Finding ground states within the ansatz class of MPS for a given Hamiltonian is what DMRG does, and we can treat this as a black box here. Similarly, there is a large number of methods for the time evolution of MPS [25], which are contained as black boxes in available tensor network codes. As we will push those to the limit, at least a minimal understanding is useful. We want to calculate $|\Psi(t)\rangle = \mathrm{e}^{-\mathrm{i}\hat{H}t}|\Psi(0)\rangle$ with results also at intermediate times.

The earliest algorithms [26–29] all built on the idea of the Trotter decomposition. In its simplest incarnation, it runs as follows: we decompose the time evolution into small timesteps Δt. If the Hamiltonian decomposes into local Hamiltonians acting on a few neighboring sites as $\hat{H} = \sum_j \hat{h}_j$, at least the neighboring local Hamiltonians will not commute in non-trivial problems, $[\hat{h}_j, \hat{h}_{j+1}] \neq 0$. A decomposition of the exponential is therefore not exact,

$$\mathrm{e}^{-\mathrm{i}\hat{H}\Delta t} = \mathrm{e}^{-\mathrm{i}\sum_j \hat{h}_j \Delta t} \approx \prod_j \mathrm{e}^{-\mathrm{i}\hat{h}_j \Delta t}, \tag{8}$$

but the error is of order $(\Delta t)^2$ and can be made arbitrarily small (and smarter Trotter decompositions exist).

The more relevant error, however, comes from the increasing entanglement of the time-evolving state. Technically, this manifests itself as follows (for details, see [3, 25]): the factorized time-evolution operator can be expressed as a matrix product operator (MPO), which looks like an MPS, but with two physical legs, one for the ingoing physical state and one for the outgoing state. It can be applied to an MPS by contracting over physical legs, and a new MPS results with matrix dimension $D_O D$, where D_O is the matrix dimension of the MPO. This state (its MPS representation) has to be compressed back to some matrix dimension D', at an acceptable loss of accuracy, to keep it numerically tractable. Growing entanglement manifests itself by growing D' if we want to keep a certain accuracy, or declining accuracy if we limit D' to some maximum value. Either way, this is the key bottleneck and it represents physics, not some algorithmic short-coming.

When using more recent time-evolution methods such as TDVP (time-dependent variational principle), which has become pretty much the standard by now [25], both the time-stepping error and the entanglement growth problem persists, but TDVP is much better at treating longer-ranged interactions and overall substantially faster. Another approach, the Krylov space method,

which we will discuss below, will allow us to circumvent certain unexpected problems arising from time-stepping; the entanglement issue persists. We will discuss one way out of this generic problem, linear prediction, further below; it is quite specific to Green functions and spectral functions, but this is what we want to calculate.

3 DMFT in a nutshell

Let us now turn to a similarly concise presentation of DMFT, as it has been presented numerous times in this lecture series over the years (most recently in [8, 9]). Even more importantly, the fine details do not matter to us and are best taken care of by the mentioned standard open-source packages. We mainly discuss where and how TNS (mostly MPS) come into play.

Imaginary-time TN-DMFT. This approach simply adapts the approach of [12]. There, frequency-dependent quantities are expressed as a discrete set of values at the imaginary Matsubara frequencies $\mathrm{i}\omega_n$, where $\omega_n = \pi(2n+1)/\beta$ with $\beta = (k_B T)^{-1}$; they become dense as we approach $T = 0$.

Let us start from the Green function of the non-interacting part of the Hubbard Hamiltonian which has a dispersion relation $\varepsilon_{\mathbf{k}}$ (which we may have obtained from band structure, as given by DFT). The chemical potential μ is chosen to give the correct filling. Then $G^0_{\text{lattice}}(\mathbf{k}, \mathrm{i}\omega_n) = (\mathrm{i}\omega_n + \mu - \varepsilon_{\mathbf{k}})^{-1}$, a standard result. When switching on the interaction, this modifies the Green function through the self-energy $\Sigma_{\text{lattice}}(\mathbf{k}, \mathrm{i}\omega_n)$. Let us assume we know it, then $G_{\text{lattice}}(\mathbf{k}, \mathrm{i}\omega_n) = \big(\mathrm{i}\omega_n + \mu - \varepsilon_{\mathbf{k}} - \Sigma_{\text{lattice}}(\mathbf{k}, \mathrm{i}\omega_n)\big)^{-1}$. By a Fourier transformation, we can find the local lattice Green function at position 0 (which we take to be the location of the impurity), which we call G_{loc}, as

$$G_{\text{loc}}(\mathrm{i}\omega_n) = \int \mathrm{d}\mathbf{k}\, \frac{1}{\mathrm{i}\omega_n + \mu - \varepsilon_{\mathbf{k}} - \Sigma_{\text{lattice}}(\mathbf{k}, \mathrm{i}\omega_n)}\,. \tag{9}$$

The central approximation of DMFT is that Σ_{lattice} does *not* depend on $\mathbf{k}$, $\Sigma_{\text{lattice}} = \Sigma_{\text{lattice}}(\mathrm{i}\omega_n)$, and that the momentum dependence of the Green function only enters through the non-interacting dispersion relation $\varepsilon_{\mathbf{k}}$. Then

$$G_{\text{loc}}(\mathrm{i}\omega_n) = \int \mathrm{d}\mathbf{k}\, \frac{1}{\mathrm{i}\omega_n + \mu - \varepsilon_{\mathbf{k}} - \Sigma_{\text{lattice}}(\mathrm{i}\omega_n)}\,, \tag{10}$$

where we still have to find $\Sigma_{\text{lattice}}(\mathrm{i}\omega_n)$, which is done through self-consistency: the local self-energy of the lattice must be the same as the local self-energy of the interacting impurity whose electrons are exchanged with the bath,

$$\Sigma_{\text{lattice}}(\mathrm{i}\omega_n) \stackrel{!}{=} \Sigma_{\text{imp}}(\mathrm{i}\omega_n) \tag{11}$$

(Of course, we have to have $\Sigma_{\text{imp}}(\mathrm{i}\omega_n)$; at some point, we must start with an educated guess.)

Now the local lattice Green function must be close to the local Green function of the impurity, which is no special site, just one we picked, $G_{\text{imp}}(\mathrm{i}\omega_n) \stackrel{!}{=} G_{\text{loc}}(\mathrm{i}\omega_n)$. The impurity Green function reads

$$G_{\text{imp}}(\mathrm{i}\omega_n) = \frac{1}{\mathrm{i}\omega_n + \mu - \Lambda(\mathrm{i}\omega_n) - \Sigma_{\text{imp}}(\mathrm{i}\omega_n)}\,, \tag{12}$$

where the hybridization $\Lambda(\mathrm{i}\omega_n)$ encodes the hopping of electrons between impurity and bath. As we have both $G_{\text{imp}}(\mathrm{i}\omega_n)$ and $\Sigma_{\text{imp}}(\mathrm{i}\omega_n)$, we can determine $\Lambda(\mathrm{i}\omega_n)$ from Equation (12).
It can be shown that for L_b non-interacting bath sites with on-site energies ε_l and hopping amplitude V_l between impurity and bath site l, the hybridization function is given by

$$\Lambda_{\text{discr}}(\mathrm{i}\omega_n) = \sum_{l=1}^{L_b} \frac{|V_l|^2}{\mathrm{i}\omega_n - \varepsilon_l}. \tag{13}$$

We turn this relationship around and find for given L_b (typically small, $O(10)$ or so) V_l and ε_l such that $\Lambda_{\text{discr}}(\mathrm{i}\omega_n)$ approximates $\Lambda(\mathrm{i}\omega_n)$ as closely as possible, i.e., minimizes

$$\chi = \frac{1}{N} \sum_{n=1}^{N} \left| \Lambda(\mathrm{i}\omega_n) - \Lambda_{\text{discr}}(\mathrm{i}\omega_n) \right|^2 \tag{14}$$

or some similar cost function (one may place different emphasis on the Matsubara frequencies, typically only the smaller Matsubara frequencies are considered, as the large-frequency behavior is universal).
Now comes the moment of the TN impurity solver. We consider the Hamiltonian $\hat{H}$ formed from the impurity site with its local microscopic interaction at position 0 and the bath sites at positions 1 through L_b, with on-site energies ε_l and hopping V_l between 0 and l. For this Hamiltonian, we use DMRG to find either the ground state $|\Psi_{\text{GS}}\rangle$ (for a $T = 0$ calculation) or the thermal density operator (for $T > 0$) by an imaginary time evolution, as described below.
Starting from this (mixed) state, we determine the retarded Green function of the impurity in imaginary time as (here written for the ground state; the thermal expression is a simple adaptation, because we will represent mixed states as pure states, see below; we also suppress spin indices)

$$G^p(\tau) = -\Theta(\tau) \left\langle \Psi_{\text{GS}} \right| \hat{c}_0 \mathrm{e}^{-(\hat{H}-E_0)\tau} \hat{c}_0^\dagger \left| \Psi_{\text{GS}} \right\rangle , \quad G^h(\tau) = \Theta(-\tau) \left\langle \Psi_{\text{GS}} \right| \hat{c}_0^\dagger \mathrm{e}^{(\hat{H}-E_0)\tau} \hat{c}_0 \left| \Psi_{\text{GS}} \right\rangle \tag{15}$$

such that

$$G_{\text{imp}}(\tau) = G^p(\tau) + G^h(\tau), \tag{16}$$

where $\Theta(\tau)$ is the Heaviside step function. At this moment, the use of the TN impurity solver is complete. We obtain the impurity Green function in Matsubara frequencies as

$$G_{\text{imp}}(\mathrm{i}\omega_n) = \int_{-\infty}^{\infty} \mathrm{e}^{-\mathrm{i}\omega_n \tau} G_{\text{imp}}(\tau). \tag{17}$$

We now obtain the self-energy $\Sigma_{\text{imp}}(\mathrm{i}\omega_n)$ from Dyson's equation

$$\Sigma_{\text{imp}}(\mathrm{i}\omega_n) = (G^0_{\text{imp}})^{-1}(\mathrm{i}\omega_n) - G^{-1}_{\text{imp}}(\mathrm{i}\omega_n), \tag{18}$$

where $G^0_{\text{imp}}(\mathrm{i}\omega_n)$ is the impurity Green function if we switch the interaction on the impurity site off. We have closed the self-consistency loop! It turns out that we need $O(10)$ iterations of the loop to converge the result. Imaginary-time TN-DMFT was pioneered in [18], and has been

applied in various cases, culminating in [21], which will be detailed as a test case below. The key advantage of imaginary-time TN-DMFT is that imaginary time-evolution is much more benign than real-time evolution and that the bath sizes are quite small. The downside is that any quantity on the real-frequency axis can only be obtained by analytical continuation from the imaginary to the real axis, which is numerically unstable and also the key issue with QMC-DMFT, in particular for very low frequencies.

Real-time TN-DMFT. Historically this approach, which is in fact numerically more challenging, predates imaginary-time TN-DMFT. It was inspired by NRG-DMFT, which works on the real-frequency (energy) axis and does not suffer from the drawbacks of analytical continuation. NRG-DMFT also excels at the lowest temperatures; however, it is handicapped by structural constraints that lead to an explosion of computation time for the study of interesting materials and is not good at larger frequencies. As NRG can be formulated as a very special MPS algorithm, the use of real-time TN-DMFT promises better performance.
Compared to imaginary-time TN-DMFT, the main changes are that local bath and impurity Green functions, the local bath and impurity self-energies, and the hybridization function are functions of ω instead of $\mathrm{i}\omega_n$. The hybridization function is represented by a bath in a different way: let us take an electronic band of some bandwidth $2D$ (D not to be confused with the matrix dimension D), and discretize it into N_b intervals I_l which in the simplest case have equal size, but different choices are possible (in NRG for instance, intervals become exponentially small close to the Fermi edge, a key ingredient for the success of NRG). The local site energy and the hopping between impurity and bath energy interval is given by

$$|V_l|^2 = -\int_{I_l} \mathrm{d}\omega \, \frac{1}{\pi} \mathrm{Im}\Delta(\omega) \tag{19}$$

$$\varepsilon_l = -\frac{1}{|V_l|^2} \int_{I_l} \mathrm{d}\omega \, \frac{1}{\pi} \omega \, \mathrm{Im}\Delta(\omega) \tag{20}$$

The required bath sizes N_b are *much* larger, of the order of a few hundred sites, because we have to discretize energies on a very fine grid. We cannot really expect information lost by coarse-graining in energy (frequency) space to miraculously reappear. Also, the spectral function $A(\omega)$, essentially the imaginary part of the Green function, contains a lot of fine structure to be resolved: in principle, it consists of a dense set of δ-peaks for $T \to 0$. (All this information is smoothened out and hidden across all imaginary frequencies when analytically continuing from the real to the imaginary axis, which makes the reverse procedure so difficult.)
From the pure TN perspective, the key difference is that time evolutions are now in real time, e.g.,

$$G^p(t) = \left\langle \Psi_{\mathrm{GS}} \right| \hat{c}_0 \mathrm{e}^{-\mathrm{i}(\hat{H}-E_0)t} \hat{c}_0^\dagger \left| \Psi_{\mathrm{GS}} \right\rangle. \tag{21}$$

Numerically, this is very challenging due to strong entanglement growth, hence achievable simulation times are quite limited. Nevertheless, the method has been successfully applied to two-orbital problems [15–17] and materials with three and more orbitals, starting with [30], albeit with limited frequency resolution.

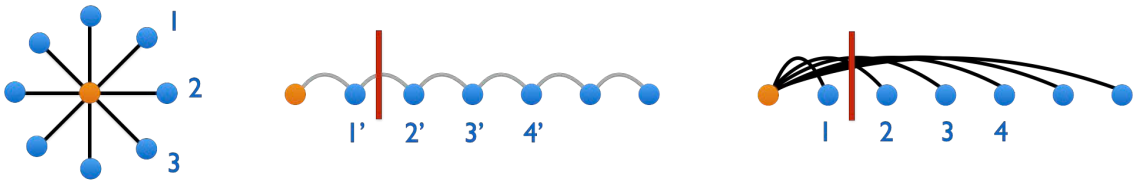

Fig. 4: *Mapping the bath: On the left is the natural star-like structure of the problem: an impurity (orange) is linked by hopping to bath orbitals (blue). This can be mapped on to a chain with nearest-neighbor hopping by tridiagonalization (center); we call this the "chain geometry". It can also be simply translated into long-ranged hoppings by flattening the star (right); we call this the "star geometry".*

4 Bath representations: chains, stars, forks and trees

The label "tensor network" reflects that we consider states on graphs to whose vertices we attach local Hilbert spaces; they can be site or orbital, momentum or energy local; the first is the most frequent case. This leaves a large amount of freedom: what graphs do we consider? How do we arrange the orbitals on the graph? In some cases, like when we are interested in a one-dimensional nearest-neighbor hopping Hubbard model, the choice is pretty obvious: we choose a one-dimensional chain, and put adjacent orbitals on adjacent vertices; then we have on-site interactions and hopping to the neighboring vertices. In many other cases, like ours, the answer is not so clear; rigorous results hardly exist.

Tensor network algorithms like DMRG exploit gauge degrees of freedom for the local tensors which allow us to give them very convenient (generalized) orthonormality properties [3] which both accelerate and numerically stabilize algorithms: many calculations can be avoided and replaced by exact analytical results. These orthonormality properties rely on the absence of loops in the graphs. Hence, quite generally, we prefer to work with chain-like tensor networks, which we already discussed (Figure 2), and tree tensor networks to be discussed below (Figure 6 shows two cases important for us). This lecture will entirely focus on loop-free networks.

Single-orbital setup. For a single-orbital problem (one relevant electronic band), the choice is between a chain or a tree-like structure. Almost all experience so far is on chains in this case, and we will only consider those. There is then the question how we arrange the bath orbitals (sites). The simple impurity problem is given by one impurity orbital with interacting electrons; these electrons can hop back and forth between the impurity orbital and one of many bath orbitals; there is no hopping between the bath orbitals themselves and the electrons on the bath orbitals are non-interacting (these are highly delocalized bath states), so the physics suggests a star-like structure as shown on the left of Figure 4.

How do we map this to a linear tensor network? We leave the impurity orbital untouched and put it on one end. Standard wisdom is that hopping entangles parts of a system, so it should be convenient to map the bath to a linear object with short-ranged hopping. In fact, any non-interacting problem can be brought into tridiagonal form, where the diagonal represents on-site

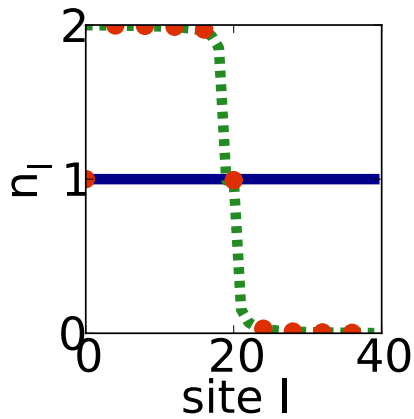

Fig. 5: *Occupation numbers in a $T = 0$ bath with an impurity.*

energies and the two off-diagonals a nearest-neighbor hopping. This tridiagonalization step is routinely employed in NRG [31], and the resulting Hamiltonian formally resembles a Hubbard chain with on-site energies and nearest-neighbor hopping (shown in the center of Figure 4). On the other hand, we could simply place the bath orbitals one after another on the chain, leaving the star-structure unaffected (shown on the right of Figure 4). The price to pay are long-ranged hopping terms, which are harder to implement and should lead to faster entanglement. Or so it seems. In fact, it turns out that the entanglement growth is much faster in the chain geometry, hence the required computational resources are much larger.

Why is this? Figure 5 has the answer. Unlike in a typical Hubbard model problem, on which most intuition has been developed, and in fact unlike the chain geometry setup, where there is the same (or no) on-site energy on each site after tridiagonalization, in our case, the on-site energies are spread widely. In a real-frequency TN-DMFT, one encounters at $T = 0$ occupation numbers as shown in Figure 5 if we order the sites l from small to big on-site energy ε_l: the low-lying orbitals contain 2 electrons, the high-lying orbitals none, and the edge would be sharp in the absence of the impurity. Coupling to the impurity broadens the step. This blocks the dynamics on almost all sites except close to the step: there are either no electrons that can hop or electrons cannot hop into the full orbital. This dynamical blocking without counterpart in the "normal" Hubbard model strongly suppresses entanglement growth.

Multi-orbital setups. In many cases of interest, the band structure around the Fermi level is dominated by multiple bands. For instance, in transition metals, there are 5 $3d$-bands; effects like crystal symmetry breaking often separate groups of $3d$-bands by energies such that only one group needs to be considered, often leading to 3-band or 2-band situations. In such cases, each of the orbitals of the "impurity" (n orbitals of interacting electrons for a n-band setup) connects to one (sometimes more, e.g., in the presence of spin-orbit coupling) of the n bands (represented by one bath each). Of course, one may represent n baths by n linear tensor networks which are joined into one huge linear structure. This means, however, that cumbersome, extremely long-ranged hoppings appear, in particular in real-time TN-DMFT with its large baths.

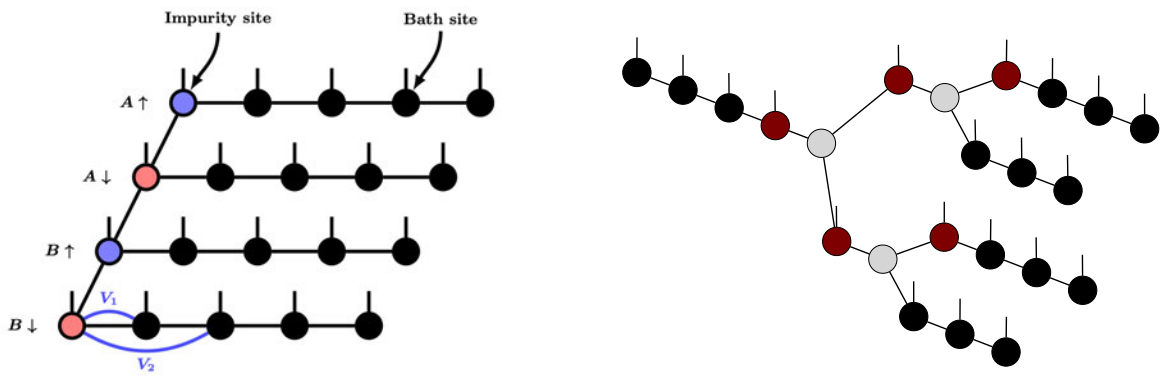

Fig. 6: *Fork tensor network (left; from [30]) and T3NS (three-legged tree tensor network; right). Note that the fork shows a two-orbital problem where the spin-up and spin-down electrons have been put into separate baths for each band, which is often numerically advantageous. In the T3NS we distinguish two types of tensors, "physical" (red or black), with two auxiliary and one physical leg, as in an MPS, and "geometrical" (grey), which have three auxiliary and no physical leg, carrying the structure of the tree. The red tensors show where the impurities could sit in a five-band problem.*

A natural answer is provided by fork tensor networks [30]: One arranges the 3 or 5 interacting orbitals "vertically", and the bath orbitals which represent the 3 or 5 bands fan out "horizontally", which resembles a fork (left part of Figure 6). The big advantage of this setup is that it contains no loops and places corresponding bath and impurity sites in proximity; it is a special case of a tree tensor network. This setup has allowed important progress in DMFT calculations for transition metal oxides. There is, however, one big disadvantage: while all bath sites and the first and last impurity orbitals are represented by three-legged tensors, as in an MPS, with computational cost scaling as D^3, the central impurity orbital(s) yield four-legged tensors (one physical and three auxiliary legs), such that computation cost scales as D^4. Given that D can be pretty large, $O(100)$ or even $O(1000)$, the calculation cost for this single tensor will become dominant at some point.

A way out, pioneered by [32] (and first used in [33], slightly modified, in the present context), are T3NS (three-legged tree tensor network states); right part of Figure 6. They sit on a tree-tensor network where *every* tensor is three-legged. This is achieved by introducing two different types of vertices on a graph, physical vertices which connect to two other vertices and have one physical leg, and "geometrical" vertices which carry the branching structure: they connect to three other vertices and have no physical leg. In the worst case of extreme branching, the number of vertices is doubled, but this factor 2 in computational cost vastly outweighs the disadvantage of four-legged tensors, which actually increases with the number of bands. In a T3NS one has the additional advantage that one can represent the bath sites efficiently by trees instead of chains; for very large baths in real-time TN-DMFT this may hold considerable advantage.

5 Time evolution: Finite temperatures, linear prediction, time steppers

Finite temperatures. Most tensor network calculations are concerned with ground states or the time-evolution of pure states prepared in, say, cold atom experiments. On the other hand, many materials present interesting cross-overs in their properties, as there are competing energy scales; the smaller scales, which often correspond to interesting collective behavior, become more prominent as temperature is lowered and thermal effects no longer "smear out" these scales. It is therefore relevant to be able to calculate spectral functions at any non-zero temperature. There are several techniques for that. The historically oldest technique is called purification [29] and transforms any mixed state into a pure state, to which conventional tensor network techniques can be applied. Consider some density operator $\hat{\rho}$ on a system S. Its eigendecomposition reads

$$\hat{\rho}_S = \sum_\alpha p_\alpha |\alpha\rangle_{SS}\langle\alpha|, \tag{22}$$

where the (real) p_α sum to 1, and where the states $\{|\alpha\rangle_S\}$ live in the Hilbert space $\mathcal{H}_S$ of the physical system S. We now "double" the physical system by creating an environment E which has the same Hilbert space as S, $\mathcal{H}_E = \mathcal{H}_S$, so we can identify 1-to-1 states $|\alpha\rangle_S$ and $|\alpha\rangle_E$ in S and E. In this larger Universe SE, we introduce a pure state

$$|\Psi\rangle_{SE} := \sum_\alpha \sqrt{p_\alpha}|\alpha\rangle_S|\alpha\rangle_E; \tag{23}$$

then the mixed state of S is given by the partial trace over E of the pure state projector given by $|\Psi\rangle_{SE}$:

$$\hat{\rho}_S = \mathrm{tr}_E\big(|\Psi\rangle_{SESE}\langle\Psi|\big). \tag{24}$$

This allows us to replace *all* calculations using mixed states by familiar pure-state calculations, at the cost of working in a universe double the size of the system, for instance

$$\langle\hat{O}_S\rangle_S = \mathrm{tr}_S\hat{O}_S\hat{\rho}_S = \mathrm{tr}_S\hat{O}_S\mathrm{tr}_E\big(|\Psi\rangle_{SESE}\langle\Psi|\big) = \mathrm{tr}_{SE}\hat{O}_S|\Psi\rangle_{SESE}\langle\Psi| = {}_{SE}\langle\Psi|\hat{O}_S|\Psi\rangle_{SE}. \tag{25}$$

Similarly, and relevant in our case, the real-time evolution of a density operator,

$$\hat{\rho}_S(t) = \mathrm{e}^{-\mathrm{i}\hat{H}_S t}\hat{\rho}_S(0)\mathrm{e}^{+\mathrm{i}\hat{H}_S t}, \tag{26}$$

becomes a pure-state real-time evolution of $|\Psi\rangle_{SE}$,

$$|\Psi\rangle_{SE}(t) = \mathrm{e}^{-\mathrm{i}\hat{H}_S t}\mathrm{e}^{-\mathrm{i}\hat{H}_E t}|\Psi(0)\rangle_{SE}, \tag{27}$$

where $\hat{H}_E$ can be any hermitian operator. It turns out that the choice of $\hat{H}_E$ strongly influences the amount of entanglement, hence the dimensions of the MPS and the computational cost. Without a rigorous mathematical proof, the best choice is to set $\hat{H}_E = -\hat{H}_S$. A hand-waving argument is to consider the time-evolution of a purified projector on an eigenstate of $\hat{H}_S$: it

is exactly for this choice of $\hat{H}_E$ that the purification stays invariant, hence maintains constant entanglement as it ideally should for an energy eigenstate evolving in time.
The problem is of course that we usually don't know the eigenrepresentation of $\hat{\rho}_S$; if we did, we could solve everything directly. In the case of a thermal density operator $\hat{\rho}_S(\beta) = \mathrm{e}^{-\beta\hat{H}_S}$ there is a simple constructive procedure: we can write

$$\hat{\rho}_S(\beta) = \mathrm{e}^{-\beta\hat{H}_S/2} \cdot \hat{I} \cdot \mathrm{e}^{-\beta\hat{H}_S/2} \tag{28}$$

where (up to normalization!) the identity operator $\hat{I}$ is the infinite-temperature density operator $\hat{\rho}_S(0)$. An imaginary time-evolution up to times $-\mathrm{i}\beta/2$ can be carried out by standard time-evolution tools and is in fact numerically advantageous: all numerical errors are damped out exponentially. It therefore remains to purify the identity. But this is very simple: assuming that $\mathcal{H}_S$ is a tensor product of local state spaces $\mathcal{H}_S^1$ of dimension d (works in our use cases), the identity, normalized to 1, is purified by

$$|\Psi(0)\rangle_{SE}^1 = \frac{1}{\sqrt{d}} \sum_\sigma |\sigma\rangle_S |\sigma\rangle_E, \tag{29}$$

where we couple states in S and E pairwise with equal weight. This can be represented as a MPS of dimension $D = 1$ (build this as an exercise; or check for instance [3]) and can be written down by hand. In fact, smart choices of pairs enables the introduction of good quantum numbers for a drastic speed-up: for instance, for a spin-1/2,

$$|\Psi(0)\rangle_{SE}^1 = \frac{1}{\sqrt{2}}\big(|\uparrow\rangle_S|\downarrow\rangle_E - |\downarrow\rangle_S|\uparrow\rangle_E\big) \tag{30}$$

is a purification which conserves both magnetization and total spin. The purification of the normalized identity on all of S is then simply the tensor product of these states on all sites.

Linear prediction. Either due to entanglement growth (as the dominant problem) meeting limited computational resources or due to accumulating numerical errors, high-accuracy calculations are strongly limited in time. This is a crucial problem: take a correlated metal. The main energy scale is set by the bandwidth, which is often of the order of 1 eV (corresponding to a temperature of 11 600 K). At the same time, many correlation effects are on energy scales of the order of 1 meV (or about 10 K). This means that we have to resolve frequencies down to 0.001 or less of the typical frequency set by the bandwidth, or times that are about 1000 times the typical time scale (like the inverse hopping in the Hubbard model). Similarly, we have to access temperatures being a similar fraction of the dominant temperature scale.
At the moment, no tensor network time-evolution for non-trivial problems gets anywhere near that. A powerful, but also quite simple technique that allows to extend the time range by a factor of 10 or so for specific time-dependent quantities is provided by linear prediction [34]. Let us assume that we have a time series of values x_n at times $n\Delta t$, $n = 1, \ldots N$. We make the bold assumption that the values x_n are for any n a *fixed* linear combination of p previous values,

$$x_n \stackrel{!}{=} -\sum_{k=1}^{p} a_k x_{n-k}. \tag{31}$$

We obtain the $\{a_k\}$ by using available data x_n and choosing the $\{a_k\}$ to provide the optimal fit over some training set of $\{x_n\}$, typically excluding short-time behavior. Once we have the $\{a_k\}$, we can predict $x_{N+1} = \sum_{k=1}^{p} a_k x_{N+1-k}$ from data up to x_N, and extend this prediction iteratively as long as we wish. Formally, we can introduce a column vector with $\mathbf{x}_n = [x_{n-1} \dots x_{n-p}]^T$ and a matrix

$$\mathbf{A} = \begin{bmatrix} -a_1 & -a_2 & \dots & -a_p \\ 1 & 0 & & \\ 0 & 1 & \ddots & \\ & \ddots & \ddots & 0 \\ & & 0 & 1 \end{bmatrix} \tag{32}$$

such that $\mathbf{x}_{n+1} = \mathbf{A} \cdot \mathbf{x}_n$ and $\mathbf{x}_{n+m} = \mathbf{A}^m \cdot \mathbf{x}_n$. The last equation indicates that the time series is given by a superposition of p oscillating exponential decays,

$$x_{n+m} = \sum_{\nu=1}^{p} c_\nu \mathrm{e}^{-\mathrm{i}(\omega_\nu - \mathrm{i}\eta_\nu)m} x_n. \tag{33}$$

But this is ideal for the calculation of time-dependent Green functions: they usually have the form

$$G(k,\omega) = \frac{1}{\omega - \varepsilon_k - \Sigma(k,\omega)} \tag{34}$$

where the complex self-energy $\Sigma(k,\omega)$ determines where the poles of $G(k,\omega)$ sit. Assuming we have a single pole at $\omega_1 - \mathrm{i}\eta_1$, the Fourier transform from frequency to time yields

$$G(k,t) = c_1 \mathrm{e}^{-\mathrm{i}(\omega_1 - \mathrm{i}\eta_1)t}, \tag{35}$$

and a superposition of such terms for multiple poles. But this is exactly the form of the time series of linear prediction; it is hence perfectly suited for our purpose.

Time steppers, beware! Perhaps the biggest surprise (which should not have been one) encountered in imaginary-time TN-DMFT was the problematic use of time-steppers: as discussed, most time-evolution methods rely on partitioning the object $\mathrm{e}^{-\hat{H}\tau}$ into small timesteps $\mathrm{e}^{-\hat{H}\Delta\tau}$ and factorizing ("Trotterizing") this expression. The Trotter error is given by some power in $\Delta\tau$, which we therefore control very well; it can be made very small.
The Fourier transformation of the Matsubara Green function $G(\omega) = -\langle \mathcal{T} \hat{c}(0) \hat{c}^\dagger(\tau) \rangle$, where $\mathcal{T}$ is the time-ordering operator, reads for large frequencies *exactly* $1/\mathrm{i}|\omega|$. This reflects the jump of 1 at $t{=}0$ due to time-ordering and the fact that fermionic operators anticommute, hence a fundamental quantum property. Yet, using time stepper methods naively, one finds that the numerator deviates substantially from 1. This is *not* a consequence of the Trotter errors as such. Assume that the value of $G(\tau)$ is exact at each timestep $\tau_n = n\Delta\tau$. How do we interpolate between these values? This is important because the statement that "short times correspond to large frequencies" in a Fourier transformation is not precise: large frequencies correspond to short time *intervals* in sampling; in fact, this is the origin of the wrong large-ω behavior.

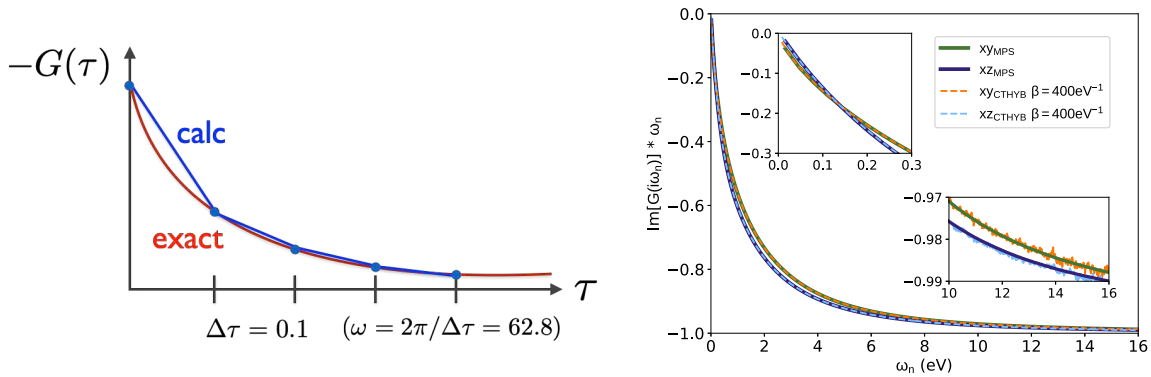

Fig. 7: *Time stepper problem (left): even if we know $G(\tau)$ exactly at certain points in time, how do we interpolate? After proper time evolution (right): for some Green function calculated both by MPS and QMC, we find that Im $G(i\omega_n)\,\omega_n$ approaches -1 correctly for large ω_n; the QMC results exhibit statistical noise around the MPS results (see second inset).*

The changes in $G(\tau)$ are largest for small times, and therefore reducing timesteps is important for small t. Naive reductions of $\Delta\tau$ in time-stepping methods increase computation time drastically and numerical errors somewhat. An alternative approach, which we use for short times before switching to time-steppers is the Krylov method, which we formulate here for real-times (so that there is no misunderstanding that it applies to real and imaginary time-evolutions): assume we want the time evolution of some state, $|\psi(\delta t)\rangle = \mathrm{e}^{-\mathrm{i}\hat{H}\delta t}|\psi(0)\rangle$; inspired by a Taylor series expansion of the evolution operator (which is not advisable), we iteratively build a Krylov-basis of Hilbert space by performing an iterative Gram-Schmidt orthogonalization of the sequence $\{|\psi(0)\rangle, \hat{H}|\psi(0)\rangle, \hat{H}^2|\psi(0)\rangle, \ldots\}$, leading to Krylov basis vectors $|q_0\rangle, \ldots, |q_{m-1}\rangle$. This becomes computationally very expensive with increasing powers of $\hat{H}$, but m less than 10 easily suffices. We then get highly exact results by evaluating $\mathrm{e}^{-\mathrm{i}\hat{H}\delta t}|\psi(0)\rangle$ in the subspace spanned by the (incomplete) Krylov basis: by the Gram-Schmidt construction, the Hamiltonian is tridiagonal in the Krylov basis, $\alpha_n = \langle q_n|\hat{H}|q_n\rangle$ and $\beta_{n-1} = \langle q_n|\hat{H}|q_{n-1}\rangle$. If we introduce the matrix $Q_m = [|q_0\rangle|q_1\rangle \ldots |q_{m-1}\rangle]$ with the Krylov vectors as columns, the projected Hamiltonian matrix $T_m = Q_m^\dagger H Q_m$ has dimension $m \times m$ and is tridiagonal. It can then be shown that

$$|\psi(\delta t)\rangle \approx Q_m \exp(-\mathrm{i}T_m\delta t)|e_0\rangle, \tag{36}$$

where $|e_0\rangle = [1\,0\,0\,\ldots]^T$. Now, as m is small and T_m almost trivial to diagonalize numerically, we can calculate at essentially no cost any time-evolution not just for δt, but also smaller times, at the same or even higher accuracy; in other words, we get arbitrarily small time steps at almost no extra cost. There is then no need to guess an interpolation of $G(t)$, and results become excellent.

It turns out that the recipe of using Krylov time evolution for short times before switching to faster methods for longer times immediately resolves the large-frequency problem in $G(\mathrm{i}\omega_n)$.

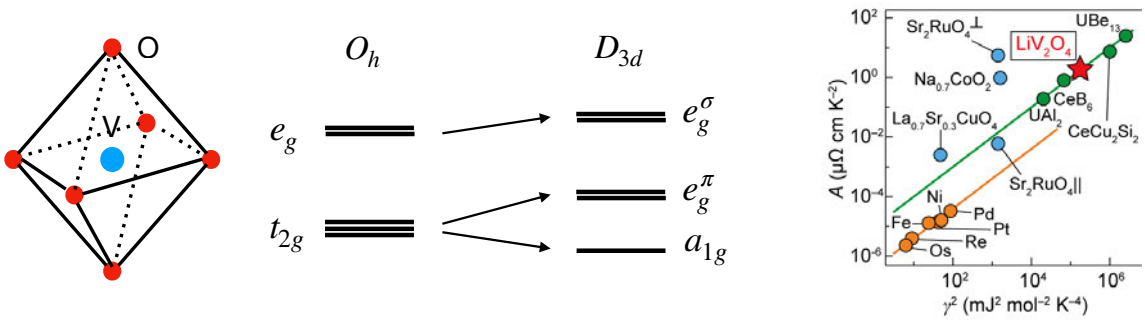

Fig. 8: *Lithium vanadate LiV_2O_4: (left) Due to its crystalline structure, the degeneracy of the 5 $3d$ orbitals of Vanadium is lifted; the Fermi surface crosses two degenerate e_g^π and one a_{1g} orbital. (right) The surprising feature of lithium vanadate is that in a plot of the Kadowaki-Wood ratio of A, the T^2-coefficient in electrical resistivity, and γ^2, the (square of the) Sommerfeld coefficient in the electronic specific heat capacity, it sits right on the line defined by heavy fermion materials (from [35]).*

6 Test case: LiV_2O_4

Let us illustrate the previous discussion by a material, where in my opinion, tensor networks have for the first dime decisively outperformed competing impurity solvers and been crucial in elucidating (new) physics [21], as opposed to previous calculations, where quantitative improvements in performance and accuracy were achieved without making a real qualitative difference. At the same time, the cooperative concurrent use of competing impurity solvers served as an excellent mutual benchmark in regimes where several methods work, providing confidence in the validity of the approaches.

The material in question is LiV_2O_4, which we will call lithium vanadate in the following, although there are further lithium vanadates with other stoichiometric coefficients. The relevant orbitals around the Fermi edge are the 5 d-orbitals of vanadium. Due to the local symmetry provided by the crystalline structure (Figure 8), these 5 orbitals are not degenerate, but split into 2 e_g^σ orbitals, 2 e_g^π orbitals and one a_{1g} orbital. The resulting bands are such that the 2 bands from the e_g^σ orbitals are sufficiently high above the Fermi energy that they can be neglected here. The lower three bands are all cut by the Fermi energy, and we have an effective three-band problem. Furthermore, an electron counting yields that there are 1.5 electrons per V in these 3 orbitals, which are hence on average quarter-filled. Also, the lattice structure is such that there is geometric frustration of (antiferro)magnetic interactions, such that this material does not show magnetic order down to the lowest temperatures. It was also observed that the Fermi liquid scale, where $\rho \propto T^2$, is only reached well below 10 K.

The reason why this material has created a lot of interest since around 1997 [36–38] is an anomalously large effective mass enhancement below about 25 K; experimentally, one finds $m^*/m \sim 25$, pointing to very strong collective effects. A further reason of interest was that in a Kadowaki-Wood plot of A, the T^2-coefficient in electrical resistivity, vs. γ^2, the (squared) Sommerfeld coefficient in the electronic specific heat capacity (Figure 8), lithium vanadate is

right on a line populated by "heavy fermion" materials, where large effective masses are found in rare earth compounds. If one claimed that lithium vanadate was a heavy fermion material, it were the first of its kind in a transition metal oxide.

The problem is that the standard explanation of heavy fermion behavior cannot work: in rare earths, there are very narrow f-bands, corresponding to almost localized electrons, and d-bands with itinerant electrons, which carry conductivity. Similar to the Kondo effect, where localized impurity electrons form a singlet with free electrons in an antiferromagnetic interaction, the localized f-band electrons form singlets with the d-band itinerant electrons, which are therefore much less mobile, yielding a very large effective mass. In lithium vanadate there are no f-band electrons. Yet, one finds that the e_g^π bands have a width of about 2 eV, and the a_{1g} band is much narrower at about 0.5 eV or so, so an early attempt at an explanation was that the e_g^π electrons play the role of the itinerant electrons, and the a_{1g} electrons that of the localized electrons [39]. Yet, this explanation does not hold up to closer inspection: the dominant magnetic interaction is a ferromagnetic Hund coupling, various experimental details do not match this explanation, and finally a band width of 0.5 eV is not excessively small.

Many other explanations were attempted, and it seems that most people were ultimately satisfied with an explanation based on extensive numerics [40], where it was found that the interactions in lithium vanadate repopulate bands such that there is a density of $n \approx 0.98$ in the a_{1g} band, in close proximity of a density 1, with an associated transition to a Mott insulator. The closeness of the transition should induce the observed strong mass renormalization; the e_g^π bands are then mere bystanders.

The problem of all explanations was, to the extent that numerics were used, that the typical temperature scale set by the bands is around 11 000 K (1 eV), and that starting from there, numerical methods could not get below around 300 K, whereas the physics of interest is on the scale of 10 K or so. This requires methods that can reach such temperatures and provide energetic resolution better than 1/1000 of the typical energy scale! This is now possible.

In our calculations, we modeled the interactions on the Vanadium atoms by a standard three-band Hubbard-Kanamori model given by

$$\begin{aligned}\hat{H}_K = U \sum_m \hat{n}_{m\uparrow}\hat{n}_{m\downarrow} + U' \sum_{m\neq m'} \hat{n}_{m\uparrow}\hat{n}_{m'\downarrow} + (U'-J) \sum_{m<m',\sigma} \hat{n}_{m\sigma}\hat{n}_{m'\sigma} - J \sum_{m\neq m'} \hat{d}^\dagger_{m\uparrow}\hat{d}_{m\downarrow}\hat{d}^\dagger_{m'\downarrow}\hat{d}_{m'\uparrow} \\ + J \sum_{m\neq m'} \hat{d}^\dagger_{m\uparrow}\hat{d}^\dagger_{m\downarrow}\hat{d}_{m'\downarrow}\hat{d}_{m'\uparrow}\,.\end{aligned}$$

Microscopic calculations yield values $U = 3.94$ eV, $U' = 2.83$ eV, and $J = 0.56$ eV, where $U' = U - 2J$, which are generally agreed on.

We applied (i) our TN-DMFT solver using purified MPS at non-zero temperatures down to 2.9 K ($\beta = 4000$), and (ii) a high-performance QMC-DMFT solver down to the lowest temperatures possible, and (iii) an NRG-DMFT solver, which is very reliable at temperatures close to zero, but is only able to treat a simplified version of the Hubbard-Kanamori model, the Dvorin-Narath model, where the last line of the Hubbard-Kanamori model containing pair hopping terms is dropped.

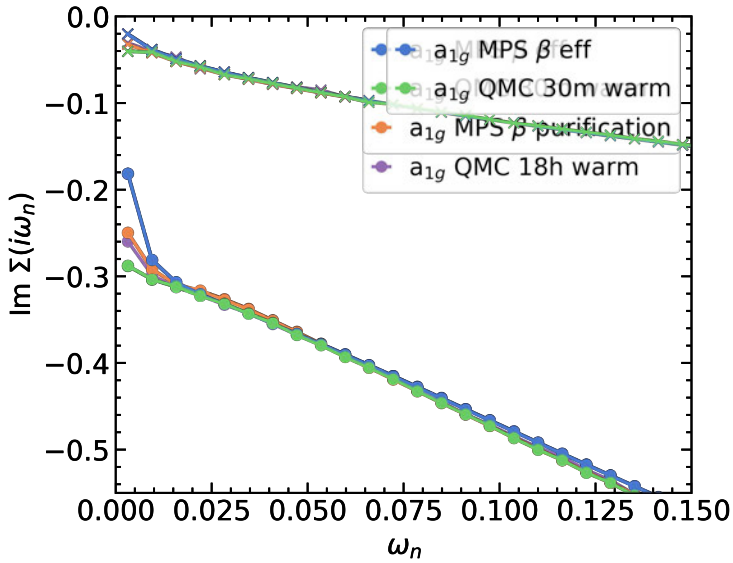

Fig. 9: *Limits of accuracy: We show the most challenging quantity calculated, the imaginary part of the self-energy. $\beta = 1000$, the lowest temperature reached by QMC. On the top are results for the e_g^π bands without specific issues. On the bottom are results for the a_{1g} band. Blue: a $T = 0$ calculation for TN-DMFT; orange: the $\beta = 1000$ result for TN-DMFT. QMC agrees well even at the difficult smallest Matsubara frequency, provided that the warm-up time is ramped up from more typical 30 minutes to extreme 18 hours.*

Let us be clear up front regarding the limitations. Figure 9 shows that QMC-DMFT recognizes the bend of the difficult Im $\Sigma_{a_{1g}}$ towards zero at the lowest frequencies, required by Fermi liquid physics, only after unusually long warm-up times; acceptance rates are extremely low. It should be noted, though, that here QMC-DMFT has been pushed to β=1000 (11.6 K), beyond the usual state of the art. TN-DMFT maps Im Σ properly, but runs into difficulties at its lowest temperature β=4000 (2.9 K), where we observe an odd-even effect in bath length for the 1–2 lowest Matsubara frequencies. Physics dictates the correct choice, but a question mark remains.

Let us now turn to the physics. In a DFT calculation, where interactions are treated in a very approximate manner, the (low-temperature) occupations are found to be $n(a_{1g}) = 0.427$ and $n(e_g^\pi) = 0.537$. Proper consideration of the local interactions shifts them to $n(a_{1g}) \approx 0.9$ and $n(e_g^\pi) \approx 0.3$. This already excludes the explanation based on a vicinity to a Mott transition; $n(a_{1g}) \approx 0.9$ is too small for that.

A crucial observation is now, on the other hand, given by the spectral function $A(\mathbf{k}, \omega)$. Figure 10 shows in the left panel as lines the DFT bands around the Fermi energy; the colors indicate the spectral function $A(\mathbf{k}, \omega)$ at T=11.6 K. The crucial feature is the very narrow horizontal structure, a flat band, right at the Fermi energy. The physics is easier to see if we consider the momentum-integrated $A(\omega)$ at various temperatures (right panel), focusing on the contribution from the a_{1g} band. As we lower temperature, a very sharp peak emerges at temperatures of a few 10 K, still rounded out at around 40 K, very pronounced and peaked at around 10 K. This very narrow peak (not visible in previous calculations due to temperature) corresponds to an emerging flat band with quite localized, hence very heavy electrons, and is therefore at the bottom of the experimental phenomenon to be explained.

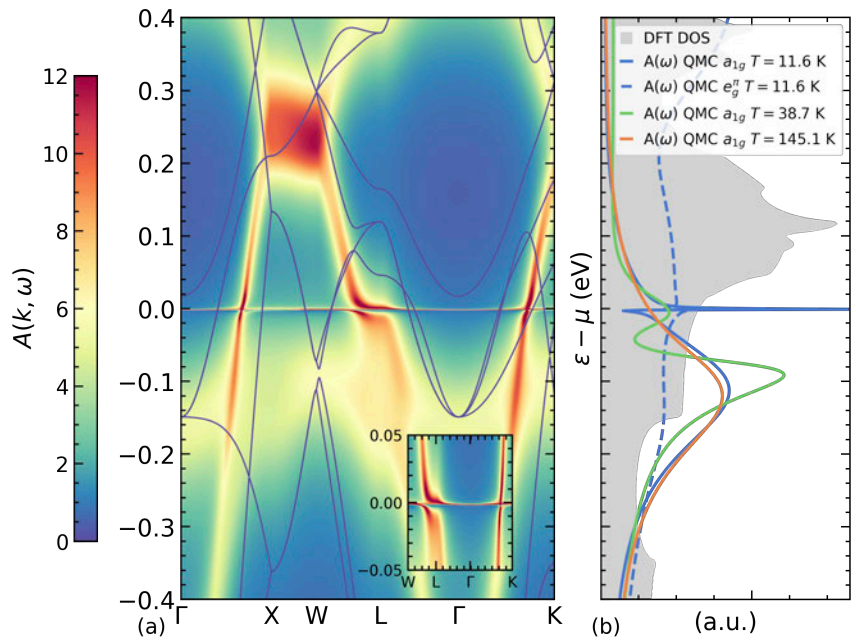

Fig. 10: *Spectral functions $A(\mathbf{k},\omega)$ at T=11.6 K (left) and $A(\omega)$ at various T (right). The lines on the left are the DFT band structure. On the right, one sees that the a_{1g} part of $A(\omega)$ develops a very sharp peak at very low T, previously unreached. From [21].*

It turns out that the effect is due to Hund's coupling J, involving all three bands: the two e_g^π are *not* bystanders. We can show that for fixed U, U' the splitting of the levels (hence the redistribution of the densities to the values reported above) is driven by J, but for $J \gtrsim 0.3$eV essentially independent of J. The Hund coupling at the same time leads to the formation of high-spin states between electrons in all three bands (almost the entire weight of states with 2 or 3 electrons on the V d-orbitals is in S=1 and S=3/2 states). This is the underlying mechanism of the low-energy flat band, because this inhibits the low-energy motion in the a_{1g} band, which we can also read as a high effective mass. We might call the effect "Hund-driven orbital-selective Mott physics".

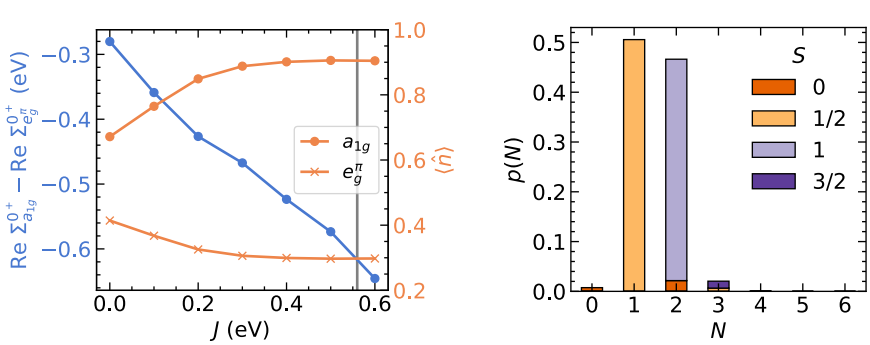

Fig. 11: *Hund's physics in lithium vanadate. Left panel: difference in the real parts of the zero-frequency self-energy of the e_g^π and a_{1g} bands, showing their relative shift due to Hund's coupling J, and the resulting changes in electronic densities. Above a certain threshold, densities are robust under changes of J. Right panel: total spin vs. number of electrons on the impurity orbitals at J=0.56 eV. High-spin states dominate. From [21].*

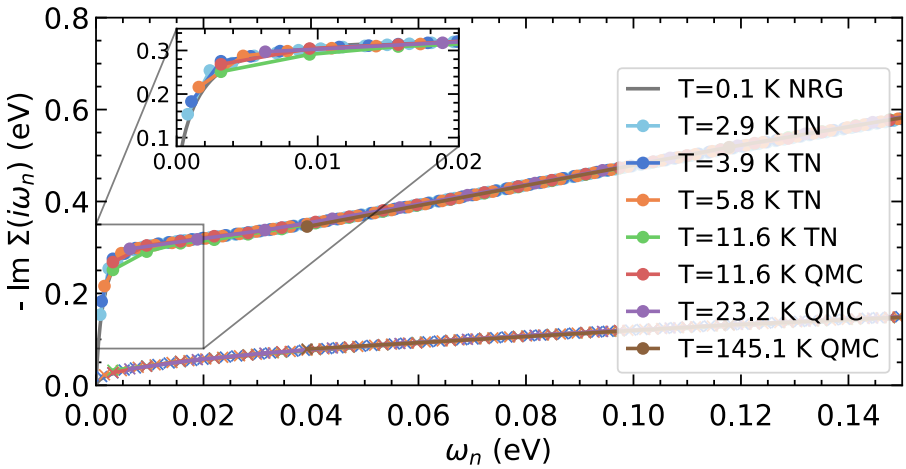

Fig. 12: *Imaginary part of the low-frequency self-energy of the e_g^π bands (bottom) and the a_{1g} band. The results of all impurity solvers at various T collapse on the same curve; in particular, the sharp bend appearing for the a_{1g} band self-energy is properly captured. From [21].*

Let us return to the question of the reliability of these simulations. It can be assessed from the imaginary part of the self-energy on the Matsubara axis $\Sigma(\mathrm{i}\omega_n)$, which is not the physically most intuitive quantity, but numerically particularly demanding. As one can see, all three methods – NRG at the lowest T, tensor networks at very low T, and QMC at higher T seamlessly overlap (Figure 12). In particular, the strong bend emerging at the lowest ω_n for very low T is resolved consistently. A numerically easier, but physically more obvious quantity is the spectral weight right at the Fermi edge. It should become large for the a_{1g} band at the lowest temperatures, and linear in T in the Fermi liquid regime. This is observed, but more importantly, the results of NRG, TN and QMC impurity solvers connect with high accuracy (Figure 13). In fact, the narrow flat band observed in theory was at the same time independently seen in ARPES measurements of lithium vanadate at 6 K [35].

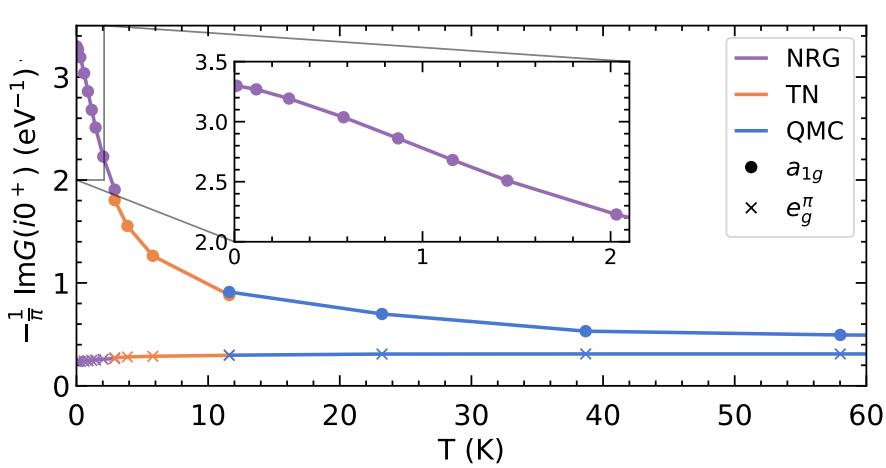

Fig. 13: *Zero-frequency limit of the spectral function of the a_{1g} band of lithium vanadate: NRG, TN and QMC solvers match seamlessly across the entire temperature range. From [21].*

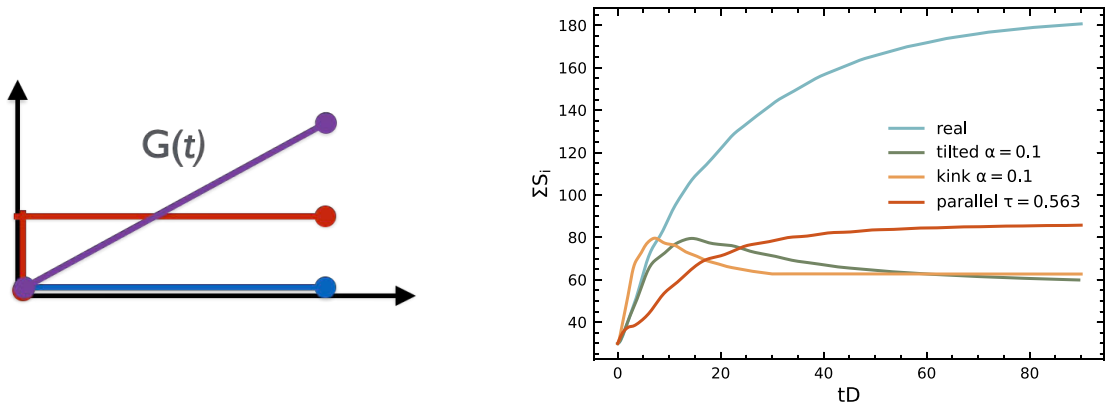

Fig. 14: *Left: Contours in the complex plane for the same final real time: real time (blue), parallel (red) and tilted (violet). Right: entanglement growth with time for various contours. All contours in the complex plane have largely lower entanglement growth than the real-time contour, saving substantial resources and allowing larger times. From [22].*

7 Real time and complex time evolution

The central weakness of TN-DMFT on the imaginary axis is, as in the case of QMC-DMFT, the necessity to perform some kind of analytical continuation from the imaginary to the real axis, an ill-conditioned problem. While a plethora of techniques has been developed (most prominently, maximum-entropy methods and the Padé approximation), none of them is fully convincing. On the other hand, real-frequency information is experimentally crucial, for instance for scattering, ARPES, and transport experiments. As we have already seen, we often require extremely high resolution, better than 1/1000. For this, analytic continuation methods are particularly bad.

The problem with real-time calculations is that they are limited in time due to entanglement growth, hence yield limited frequency resolution. One way forward is the use of time-evolution contours in the complex plane, where going into the complex plane suppresses high-frequency information, which we are usually less interested in, and strongly suppresses entanglement growth. The complex plane allows myriads of contours, and it is hard to name an "optimal" one. For instance, one contour might allow for faster time-evolutions because of less entanglement, but loses some accuracy. In the first self-consistency loops of a DMFT calculation, this may be an acceptable trade-off. In any case, the chosen contour must come with a numerically robust continuation scheme to the real axis.

Figure 14 (left) shows possible contours, in particular two that have been studied in depth: (i) a contour tilted by an angle α with respect to the real axis, and (ii) a contour which follows the imaginary axis up to some imaginary time τ and then continues parallel to the real axis ("parallel contour"). In all cases, entanglement, hence required tensor sizes, are much smaller than for a real time evolution (right panel); actual calculation times, which of course depend on α and τ, cannot easily be compared between contours as it is hard to define final identical accuracy on the real axis. In general, the parallel contour is slower, but ultimately yields the more accurate results for the moment. (All is of course much faster than on the real axis.)

Let us therefore consider the parallel contour in more detail. We introduce a generalized spectral function

$$\mathcal{A}_\tau(\omega) := \frac{1}{2\pi} \lim_{\eta\to 0^+} \int \mathrm{d}t \; \mathrm{e}^{\mathrm{i}(\omega+\mathrm{i}\eta)t} \left\langle \{\hat{c}_0(t+\mathrm{i}\tau), \hat{c}_0^\dagger\} \right\rangle_{\psi_0} \tag{37}$$

such that we can connect to the usual spectral function and the retarded Green function as

$$\mathcal{A}_0(\omega) = \mathcal{A}(\omega) = -\frac{1}{\pi} \mathrm{Im}\, G^R(\omega). \tag{38}$$

This is already the complete information required for $G^R(\omega)$: we can determine $\mathrm{Re}\, G^R(\omega)$ from $\mathrm{Im}\, G^R(\omega)$ by invoking the Kramers-Kronig relation, whose numerical implementation does not introduce unexpected problems.

The usefulness of this approach comes from the following analytically exact inversion formula

$$\mathcal{A}_\tau(\omega) = \mathcal{A}(\omega)\, \mathrm{e}^{-\tau|\omega|} \qquad (\eta \to 0). \tag{39}$$

This simple result, which only holds for the spectral function, not Green functions in general, can be derived easily by expressing $\mathcal{A}_\tau(\omega)$ in the Lehmann representation. We can therefore directly calculate $\mathcal{A}(\omega)$ from $\mathcal{A}_\tau(\omega)$!

Yet, a naive application of this formula runs into the following problem: because of $|\omega|$, $\mathrm{e}^{-\tau|\omega|}$ has a kink right at $\omega{=}0$, in the low-frequency regime that interests us most. As $\mathcal{A}(\omega)$ does *not* have such a kink, there must be one in $\mathcal{A}_\tau(\omega)$ to exactly compensate for that kink in the exponential. Numerically, one finds (not surprisingly) that this kink is rounded out due to finite times reached! There is, however, a surprisingly simple and cheap workaround: we can eliminate the kink to some desired order in τ by a superposition of several $\mathcal{A}_{\tau_k}(\omega)$, where the displacement from the real axis is by various τ_k. In fact, it is not even necessary to redo the costly time-evolutions; one may add small imaginary time-evolutions at the origin which cover the difference between the various τ_k. Hence, the additional cost is negligible.

In the superposition of n contours, we choose the weights a_k (which by construction must be both positive and negative) such that

$$\sum_{k=1}^n a_k = 1 \qquad \sum_{k=1}^n a_k \tau_k^l = 0 \tag{40}$$

where $l = 1, \ldots, n{-}1$. For $n{=}3$, then, we eliminate the kink up to ω^2 in

$$h(\omega) = \sum_{k=1}^n a_k\, \mathrm{e}^{-\tau_k|\omega|}. \tag{41}$$

Summing the inversion formula over the choices of τ_k weighted by a_k, we find

$$\mathcal{A}(\omega)h(\omega) = \sum_{k=1}^n a_k \mathcal{A}_{\tau_k}(\omega) \tag{42}$$

which we can very stably solve for $\mathcal{A}(\omega)$ also for small ω. A robust choice is for instance $n{=}3$ with $\tau_k = 1, 1.15, 1.3$.

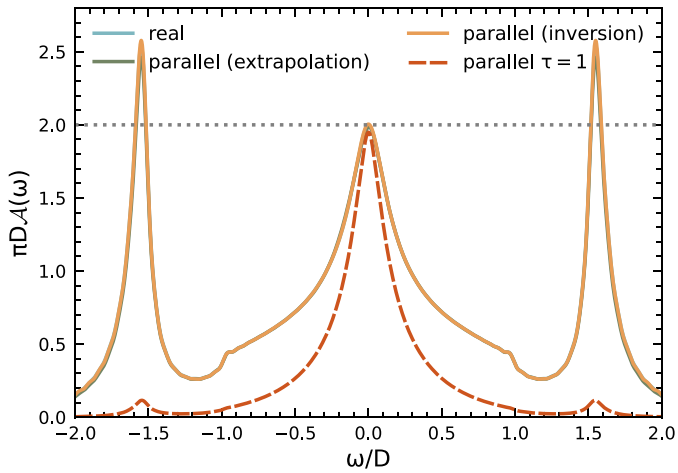

Fig. 15: *Spectral function of the SIAM for U/D=2: here, a quite precise real-time calculation is feasible. The raw data from a parallel contour is far away, but various schemes to bring us to the real axis work very well (here, we discuss only the "inversion" scheme). The sum rule $\pi D\mathcal{A}(0) = 2$ is well obeyed. From [22].*

As a simple test case, consider the single-impurity Anderson model $\hat{H} = \hat{H}_{\text{imp}} + \hat{H}_{\text{bath}}$ where

$$\hat{H}_{\text{imp}} = U\hat{n}_{0\uparrow}\hat{n}_{0\downarrow} + \varepsilon_0(\hat{n}_{0\uparrow} + \hat{n}_{0\downarrow}) \tag{43}$$

and

$$\hat{H}_{\text{bath}} = \sum_{k=1}^{N_b}\sum_{\sigma}\left(v_{0,k\sigma}\,\hat{c}^\dagger_{k\sigma}\hat{c}_{0\sigma} + v^*_{0,k\sigma}\,\hat{c}^\dagger_{0\sigma}\hat{c}_{k\sigma}\right) + \sum_{k=1}^{N_b}\sum_{\sigma}\varepsilon_{k\sigma}\hat{n}_{k\sigma} \tag{44}$$

with a semi-circular density of states

$$-\frac{1}{\pi}\text{Im}\,\Delta(\omega) = \frac{D}{2\pi}\sqrt{1-\left(\frac{\omega}{D}\right)^2} \tag{45}$$

at half-filling. This is a well-understood problem, hence a good benchmark. The gold-standard results here are provided by NRG. The bath contained 299 sites, the density of states was discretized linearly on the bandwidth $2D$. Figure 15 shows the results for the spectral function for U/D=2 also from a real-time calculation which is possible here. One recognizes that the "raw" parallel contour data is quite far from the real-frequency data, but that the inversion scheme (and also some extrapolation scheme, which we will not pursue here) very reliably put us onto the real-frequency data.

Yet, the spectral function is comparatively benign, and in this plot possible issues with the precise match of the sum rule $\pi D\mathcal{A}(\omega{=}0) = 2$ are obscured. Let us therefore turn to the imaginary part of the self-energy, which is numerically challenging, in the limit $\omega/D \to 0$ for various U/D. We observe (Figure 16) that a real-time calculation yields very bad results at low frequencies, but that the complex parallel-contour calculation, combined with the inversion method, is in excellent agreement with NRG results down to very low frequencies. The remaining discrepancies come from the discretization which limits resolution and also to some extent a quite naive use of Dyson's formula $\Sigma(\omega) = G_0^{-1}(\omega) - G^{-1}(\omega)$. Both can be improved (the extreme quality of the NRG results rests on such an improvement [41]).

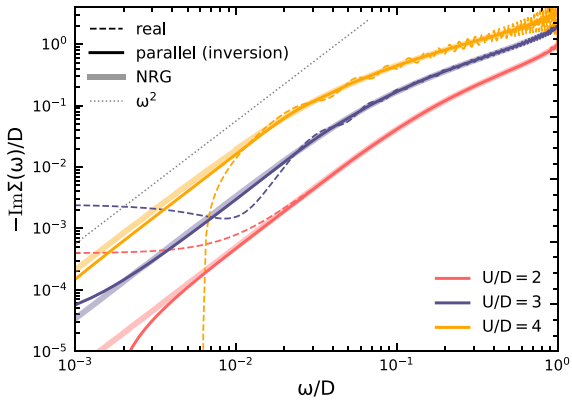

Fig. 16: *Low-frequency imaginary part of the self-energy of the SIAM: NRG vs. TN on the real axis and a complex parallel contour, combined with our inversion method. Pure real-time calculations fail; complex-time calculations agree very well with NRG down to very low frequencies. Remaining discrepancies are mainly due to the limited bath size. From [22].*

All methods are increasingly challenged for multi-orbital (multi-band) problems. For NRG, for instance, three bands are the limit under the proviso that the Hubbard-Kanamori Hamiltonian is simplified to the Dvorin-Narath Hamiltonian, where the pair-hopping term is eliminated. This increases the orbital symmetry from SO(3) to SU(3), making NRG almost exact. For the Hubbard-Kanamori Hamiltonian, results can be obtained, but are less accurate; still, they present the best currently available benchmark (in particular when combined with the methods in [41]). We again use a semi-circular DOS for the baths.

In both cases, calculations on the real-time axis are simply not sufficient. In the case of the Dvorin-Narath model, complex contours perform significantly better, but the parallel contour combined with inversion emerges as the clear winner, agreeing with NRG down to very low frequencies. As in the case of the SIAM, obvious improvements are still possible. For the Hubbard-Kanamori model, a similar picture emerges. NRG and the results from the parallel contour plus inversion do not agree as well, but this is probably rather due to the limited accuracy of NRG for this model. It can be shown, however, that improvements would shift the NRG curve downwards, closer to the TN result.

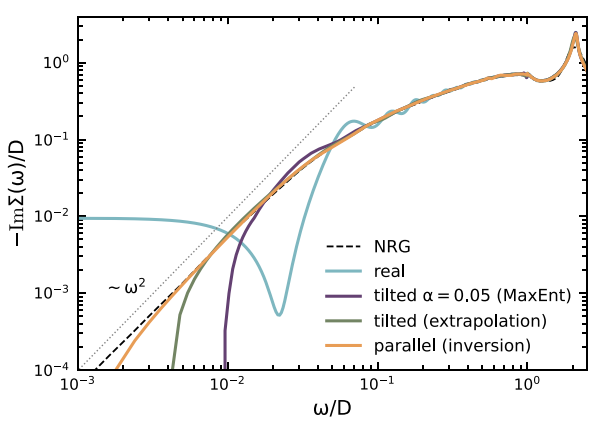

Fig. 17: *Low-frequency imaginary part of the self-energy of the Dvorin-Narath model: NRG vs. TN on the real axis and various complex contours, using different continuation methods. From [22].*

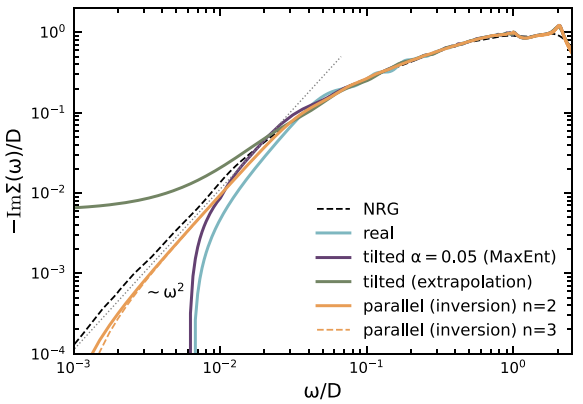

Fig. 18: *Low-frequency imaginary part of the self-energy of the HK model: NRG vs TN on the real axis and various complex contours, using different continuation methods. Note the robustness of the inversion method: the superposition of 2 and 3 contours respectively hardly affects the outcome. From [22].*

8 Outlook

Quo vadis, TN-DMFT? What is the future potential? On the one hand, we have by now a very reliable TN-DMFT working on the imaginary axis. It has shown its power in solving a long-debated problem, the "heavy fermion" behavior of lithium vanadate, by accessing very low temperatures reliably in a three-band problem. The setup of the problem and the quite short run-times indicate that 5-band and maybe even 7-band problems are within reach. Of course, all will depend on the size of the relevant temperature and energy scales to be resolved.

For TN-DMFT working on the real-time axis the jury is still out to some extent, but it is clear already that time evolution on complex contours is perfectly suited to obtain much more precise spectral functions and self-energies. What is still ongoing research is to go beyond "one-shot" calculations and to integrate the results into the self-consistency loops of DMFT. There is, however, no reason why this should present conceptual difficulties beyond the subtleties of a proper implementation. To be followed!

References

[1] S.R. White, Phys. Rev. Lett. **69**, 2863 (1992)

[2] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005)

[3] U. Schollwöck, Ann. Phys. **326**, 96 (2011)

[4] W. Metzner, D. Vollhardt, Phys. Rev. Lett. **62**, 324 (1989)

[5] A. Georges, G. Kotliar, Phys. Rev. **B 45**, 6479 (1992)

[6] A. Georges, G. Kotliar, W. Krauth, M.J. Rozenberg, Rev. Mod. Phys. **68**, 13 (1996)

[7] G. Kotliar, S.Y. Savrasov, K. Haule, V.S. Oudovenko, O. Parcollet, C.A. Marianetti, Rev. Mod. Phys. **78**, 865 (2006)

[8] E. Pavarini, E. Koch, A. Lichtenstein, and D. Vollhardt (Eds.): *Dynamical Mean-Field Theory of Correlated Electrons*, Reihe Modeling and Simulation, Vol. 12 (Forschungszentrum Jülich, 2022) `http://www.cond-mat.de/events/correl22`

[9] E. Pavarini, E. Koch, A. Lichtenstein, and D. Vollhardt (Eds.): *Understanding Correlated Materials with DMFT*, Reihe Modeling and Simulation, Vol. 15 (Forschungszentrum Jülich, 2025) `http://www.cond-mat.de/events/correl25`

[10] P. Werner, A. Comanac, L. De'Medici, M. Troyer, A.J. Millis, Phys. Rev. Lett. **97**, 076405 (2006)

[11] E. Gull, A.J. Millis, A.I. Lichtenstein, A.N. Rubtsov, M. Troyer, P. Werner, Rev. Mod. Phys. 83, 349 (2011)

[12] M. Caffarel, W. Krauth, Phys. Rev. Lett. **72**, 1545 (1994)

[13] R. Bulla, T.A. Costi, D. Vollhardt, Phys. Rev. **B 64**, 045103 (2001)

[14] C. Raas, G.S. Uhrig, Eur. Phys. J. **B 45**, 293 (2005)

[15] M. Ganahl, P. Thunström, F. Verstraete, K. Held, H.G. Evertz, Phys. Rev. **B 90**, 045144 (2014)

[16] F.A. Wolf, I.P. McCulloch, O. Parcollet, U. Schollwöck, Phys. Rev. **B 90**, 115124 (2014)

[17] M. Ganahl, M. Aichhorn, H.G. Evertz, P. Thunström, K. Held, F. Verstraete, Phys. Rev. **B 92**, 155132 (2015)

[18] F.A. Wolf, A. Go, I.P. McCulloch, A.J. Millis, U. Schollwöck, Phys. Rev. **X 5**, 041032 (2015)

[19] O. Parcollet, M. Ferrero, T. Ayral, H. Hafermann, I. Krivenko, L. Messio, P. Seth, Comput. Phys. Commun. **196**, 398 (2015)

[20] M. Aichhorn *et al.*, Comput. Phys. Commun. **204**, 200 (2016)

[21] M. Grundner, F.B. Kugler, O. Parcollet, U. Schollwöck, A. Georges, A. Hampel, Phys. Rev. B **112**, L041106 (2025)

[22] M. Grundner, P. Westhoff, F.B. Kugler, O. Parcollet, U. Schollwöck, Phys. Rev. B **109**, 155124 (2024)

[23] N. Schuch, M.M. Wolf, F. Verstraete, J.I. Cirac, Phys. Rev. Lett. **100**, 030504 (2008)

[24] J. Eisert, M. Cramer, M.B. Plenio, Rev. Mod. Phys. **82**, 277 (2010)

[25] S. Paeckel, T. Köhler, A. Swoboda, S.R. Manmana, U. Schollwöck, C. Hubig, Ann. Phys. **411**, 167998 (2019)

[26] G. Vidal, Phys. Rev. Lett. **93**, 040502 (2004)

[27] A.J. Daley, C. Kollath, U. Schollwöck, G. Vidal, J. Stat. Mech.: Theory Exp. P04005 (2004)

[28] S.R. White, A. Feiguin, Phys. Rev. Lett. **93**, 076401 (2004)

[29] F. Verstraete, J.J. Garica-Ripoll, J.I. Cirac, Phys. Rev. Lett. **93**, 207204 (2004)

[30] D. Bauernfeind, M. Zingl, R. Triebl, M. Aichhorn, H.G. Evertz, Phys. Rev. **X 7**, 031013 (2017)

[31] K.G. Wilson, Rev. Mod. Phys. **47**, 773 (1975)

[32] K. Gunst, F. Verstraete, S. Wouters, O. Legeza, D. Van Neck, J. Chem. Theory Comput. **14**, 2026 (2018)

[33] X. Cao, Y. Lu, P. Hansmann, M.W. Haverkort, Phys. Rev. B **104**, 115119 (2021)

[34] T. Barthel, U. Schollwöck, S.R. White, Phys. Rev. B **79**, 245101 (2009)

[35] D. Oh *et al.*, Proc. Nat. Acad. Sci. **122**, e2518213122 (2025)

[36] S. Kondo *et al.*, Phys. Rev. Lett. **78**, 3729 (1997)

[37] H. Takagi, C. Urano, S. Kondo, M. Nohara, Y. Ueda, T. Shiraki, T. Okubo, Mater. Sci. Eng.: B **63**, 147 (1999)

[38] C. Urano, M. Nohara, S. Kondo, F. Sakai, H. Takagi, T. Shiraki, T. Okubo, Phys. Rev. Lett. **85**, 1052 (2000)

[39] V.I. Anisimov, M.A. Korotin, M. Zölfl, T. Pruschke, K. Le Hur, T.M. Rice, Phys. Rev. Lett. **83**, 364 (1999)

[40] R. Arita, K. Held, A.V. Lukoyanov, V.I. Anisimov, Phys. Rev. Lett. **98**, 166402 (2007)

[41] F.B. Kugler, Phys. Rev. B **105**, 245132 (2022)